



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER**

**UNICAST ROUTING PROTOCOLS GUIDE
RELEASE 16.0.R4**

3HE 14140 AAAB TQZZA 02

Issue: 01

December 2018

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2018 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Getting Started.....	13
1.1	About This Guide.....	13
1.2	Router Configuration Process	15
2	RIP.....	17
2.1	RIP Overview.....	17
2.1.1	RIP Features	17
2.1.1.1	RIP Version Types.....	18
2.1.1.2	RIPv2 Authentication	18
2.1.1.3	RIP Packet Format	18
2.1.1.4	BFD Monitoring of RIP Neighbor Liveliness	20
2.2	RIPng.....	22
2.2.1	RIPng Protocol	22
2.3	Common Attributes	23
2.3.1	Metrics	23
2.3.2	Timers.....	23
2.3.3	Import and Export Policies.....	24
2.3.4	Hierarchical Levels	24
2.4	RIP Configuration Process Overview	25
2.5	Configuration Notes.....	26
2.5.1	General.....	26
2.6	Configuring RIP with CLI	27
2.6.1	RIP and RIPng Configuration Overview	27
2.6.1.1	Preconfiguration Requirements	27
2.6.1.2	RIP Hierarchy	27
2.6.2	Basic RIP Configuration	27
2.6.3	Common Configuration Tasks	28
2.6.3.1	Configuring Interfaces	29
2.6.3.2	Configuring a Route Policy	30
2.6.3.3	Configuring RIP Parameters.....	31
2.6.3.4	Configuring Global-Level Parameters	33
2.6.3.5	Configuring Group-Level Parameters	34
2.6.3.6	Configuring Neighbor-Level Parameters	35
2.7	RIP Configuration Management Tasks.....	37
2.7.1	Modifying RIP Parameters.....	37
2.7.2	Deleting a Group	37
2.7.3	Deleting a Neighbor.....	38
2.8	RIP Configuration Command Reference	39
2.8.1	Command Hierarchies.....	39
2.8.1.1	Configuration Commands.....	39
2.8.1.2	RIPng Configuration Commands.....	41
2.8.2	Command Descriptions	43
2.8.2.1	Generic Commands.....	43
2.8.2.2	RIP Commands	45
2.9	Show, Clear, and Debug Command Reference	59

2.9.1	Command Hierarchies	59
2.9.1.1	Show Commands	59
2.9.1.2	Clear Commands.....	59
2.9.1.3	Debug Commands.....	60
2.9.2	Command Descriptions	60
2.9.2.1	Show Commands	60
2.9.2.2	Clear Commands.....	85
2.9.2.3	Debug RIP Commands.....	86
3	OSPF	91
3.1	Configuring OSPF	91
3.1.1	OSPF Areas	92
3.1.1.1	Backbone Area	92
3.1.1.2	Stub Area.....	93
3.1.1.3	Not-So-Stubby Area	94
3.1.2	OSPFv3 Authentication	98
3.1.2.1	OSPFv3 Graceful Restart Helper	99
3.1.3	Virtual Links	100
3.1.4	Neighbors and Adjacencies	101
3.1.4.1	Broadcast and Point-to-Point Networks.....	101
3.1.4.2	Non-Broadcast Multi-Access (NBMA) Networks	101
3.1.5	Link-State Advertisements.....	102
3.1.6	Metrics	103
3.1.7	Authentication.....	103
3.1.8	IP Subnets	104
3.1.9	Preconfiguration Recommendations	104
3.1.10	Multiple OSPF Instances	105
3.1.10.1	Route Export Policies for OSPF	105
3.1.10.2	Preventing Route Redistribution Loops	105
3.1.11	Multi-Address Support for OSPFv3	106
3.1.12	IP Fast-reroute (IP FRR) For OSPF and IS-IS Prefixes	107
3.1.12.1	IP FRR Configuration	107
3.1.12.2	ECMP Considerations	108
3.1.12.3	IP FRR and RSVP Shortcut (IGP Shortcut).....	109
3.1.12.4	IP FRR and BGP Next Hop Resolution	109
3.1.12.5	OSPF and IS-IS Support for Loop-Free Alternate Calculation	109
3.2	Loop-Free Alternate Shortest Path First (LFA SPF) Policies	115
3.2.1	Configuring a Route Next Hop Policy Template	115
3.2.1.1	Configuring Affinity or Admin Group Constraints.....	116
3.2.1.2	Configuring SRLG Group Constraints	118
3.2.1.3	Interaction of IP and MPLS Admin Group and SRLG.....	119
3.2.1.4	Configuring Protection Type and Next Hop Type Preferences.....	119
3.2.2	Application of Route Next Hop Policy Template to an Interface.....	120
3.2.3	Excluding Prefixes from LFA SPF	121
3.2.4	Modification to LFA Next Hop Selection Algorithm.....	122
3.3	LFA Protection Using Segment Routing Backup Node SID	124
3.3.1	Configuring LFA Using Backup Node SID.....	125
3.3.2	Detailed Operation of LFA Protection Using Backup Node SID	125
3.3.3	Duplicate SID Handling	128

3.3.4	OSPF Control Plane Extensions	129
3.4	Segment Routing in Shortest Path Forwarding	132
3.4.1	OSPF Control Protocol Changes.....	132
3.4.2	Entropy Label for OSPF Segment Routing.....	135
3.4.3	IPv6 Segment Routing using MPLS Encapsulation in OSPFv3	135
3.4.3.1	OSPFv3 Extensions	135
3.5	Segment Routing Mapping Server	136
3.5.1	Segment Routing Mapping Server Prefix SID Resolution	137
3.6	SPF LSA Filtering	139
3.7	FIB Prioritization	140
3.8	Extended LSA Support in OSPFv3.....	141
3.9	Support of Multiple Instances of Router Information LSA in OSPFv2 and OSPFv3	142
3.10	OSPF Configuration Process Overview	143
3.11	Configuration Notes.....	144
3.11.1	General.....	144
3.11.1.1	OSPF Defaults.....	144
3.12	Configuring OSPF with CLI	145
3.12.1	OSPF Configuration Guidelines	145
3.12.2	Basic OSPF Configurations.....	146
3.12.2.1	Configuring the Router ID.....	147
3.12.3	Configuring OSPF Components	148
3.12.3.1	Configuring OSPF Parameters.....	148
3.12.3.2	Configuring OSPF3 Parameters.....	148
3.12.3.3	Configuring an OSPF or OSPF3 Area.....	149
3.12.3.4	Configuring a Stub Area	150
3.12.3.5	Configuring a Not-So-Stubby Area	152
3.12.3.6	Configuring a Virtual Link	153
3.12.3.7	Configuring an Interface	155
3.12.3.8	Configuring Authentication	157
3.12.3.9	Assigning a Designated Router.....	161
3.12.3.10	Configuring Route Summaries	163
3.12.3.11	Configuring Route Preferences	165
3.13	OSPF Configuration Management Tasks.....	168
3.13.1	Modifying a Router ID	168
3.13.2	Deleting a Router ID.....	169
3.13.3	Modifying OSPF Parameters	169
3.14	OSPF Configuration Command Reference	171
3.14.1	Command Hierarchies.....	171
3.14.1.1	Configuration Commands.....	171
3.14.2	Command Descriptions	178
3.14.2.1	Generic Commands.....	178
3.14.2.2	OSPF Global Commands.....	179
3.14.2.3	OSPF Area Commands.....	225
3.14.2.4	Interface/Virtual Link Commands	236
3.15	Show, Clear, and Debug Command Reference	253
3.15.1	Command Hierarchies.....	253
3.15.1.1	Show Commands	253
3.15.1.2	Clear Commands.....	254

3.15.1.3	Debug Commands.....	254
3.15.1.4	Tools Commands	256
3.15.2	Command Descriptions	256
3.15.2.1	Show Commands	256
3.15.2.2	Clear Commands.....	344
3.15.2.3	Debug Commands.....	346
3.15.2.4	Tools Commands	353
4	IS-IS	357
4.1	Configuring IS-IS	357
4.1.1	Routing	358
4.1.2	IS-IS Frequently Used Terms	359
4.1.3	ISO Network Addressing	360
4.1.3.1	IS-IS PDU Configuration.....	362
4.1.3.2	IS-IS Operations	362
4.1.4	IS-IS Route Summarization	363
4.1.4.1	Partial SPF Calculation.....	363
4.1.5	IS-IS MT-Topology Support.....	364
4.1.5.1	Native IPv6 Support.....	364
4.1.6	IS-IS Administrative Tags	364
4.1.6.1	Setting Route Tags.....	365
4.1.6.2	Using Route Tags.....	365
4.1.6.3	Unnumbered Interface Support	366
4.1.7	Segment Routing in Shortest Path Forwarding	366
4.1.7.1	Configuring Segment Routing in Shortest Path.....	367
4.1.7.2	Segment Routing Operational Procedures.....	371
4.1.7.3	Segment Routing Tunnel Management.....	379
4.1.7.4	Segment Routing Local Block	381
4.1.7.5	Entropy Label for IS-IS Segment Routing	385
4.1.7.6	Remote LFA with Segment Routing	385
4.1.7.7	Topology Independent LFA	389
4.1.7.8	Node Protection Support in TI-LFA and Remote LFA	397
4.1.7.9	IPv6 Segment Routing using MPLS Encapsulation.....	403
4.1.7.10	Data Path Support.....	405
4.1.7.11	IS-IS Control Protocol Changes	408
4.1.7.12	BGP Shortcut Using Segment Routing Tunnel.....	412
4.1.7.13	BGP Label Route Resolution Using Segment Routing Tunnel.....	413
4.1.7.14	Service Packet Forwarding with Segment Routing	413
4.1.7.15	Mirror Services and Lawful Intercept.....	414
4.1.8	Segment Routing Mapping Server Function for IPv4 Prefixes	416
4.1.8.1	Segment Routing Mapping Server	416
4.1.8.2	Prefix SID Resolution for Segment Routing Mapping Server	417
4.2	FIB Prioritization	420
4.3	IS-IS Configuration Process Overview	421
4.4	Configuration Notes.....	422
4.4.1	General.....	422
4.5	Configuring IS-IS with CLI	423
4.5.1	IS-IS Configuration Overview	423
4.5.1.1	Router Levels	423

4.5.1.2	Area Address Attributes.....	423
4.5.1.3	Interface Level Capacity.....	424
4.5.1.4	Route Leaking	425
4.5.2	Basic IS-IS Configuration	425
4.5.3	Common Configuration Tasks	428
4.5.4	Configuring IS-IS Components.....	428
4.5.4.1	Enabling IS-IS.....	429
4.5.4.2	Modifying Router-Level Parameters	429
4.5.4.3	Configuring ISO Area Addresses	430
4.5.4.4	Configuring Global IS-IS Parameters	431
4.5.4.5	Migration to IS-IS Multi-Topology	431
4.5.4.6	Configuring Interface Parameters.....	435
4.5.4.7	Configuring IS-IS Link Groups.....	439
4.6	IS-IS Configuration Management Tasks.....	440
4.6.1	Disabling IS-IS.....	440
4.6.2	Removing IS-IS	440
4.6.3	Modifying Global IS-IS Parameters	440
4.6.4	Modifying IS-IS Interface Parameters.....	441
4.6.5	Configuring Authentication using Keychains	443
4.6.6	Configuring Leaking.....	444
4.6.7	Redistributing External IS-IS Routers.....	446
4.6.8	Specifying MAC Addresses for All IS-IS Routers	447
4.7	IS-IS Configuration Command Reference	449
4.7.1	Command Hierarchies.....	449
4.7.1.1	Configuration Commands.....	449
4.7.2	Command Descriptions	454
4.7.2.1	Generic Commands.....	454
4.7.2.2	IS-IS Commands	455
4.8	Show, Clear, and Debug Command Reference	529
4.8.1	Command Hierarchies.....	529
4.8.1.1	Show Commands	529
4.8.1.2	Clear Commands.....	529
4.8.1.3	Debug Commands.....	530
4.8.1.4	Tools Commands	530
4.8.2	Command Descriptions	531
4.8.2.1	Show Commands	531
4.8.2.2	Clear Commands.....	567
4.8.2.3	Debug Commands.....	569
4.8.2.4	Tools Commands	574
5	BGP	577
5.1	BGP Overview	577
5.2	BGP Sessions	578
5.2.1	BGP Session States	580
5.2.2	Detecting BGP Session Failures	581
5.2.2.1	Peer Tracking	581
5.2.2.2	Bidirectional Forwarding Detection (BFD)	582
5.2.2.3	Fast External Failover.....	582
5.2.3	High Availability BGP Sessions	582

5.2.3.1	BGP Graceful Restart.....	583
5.2.3.2	BGP Long-Lived Graceful Restart.....	585
5.2.4	BGP Session Security	586
5.2.4.1	TCP MD5 Authentication	586
5.2.4.2	TTL Security Mechanism.....	587
5.2.5	BGP Address Family Support for Different Session Types	587
5.2.6	BGP Groups	589
5.3	BGP Design Concepts.....	590
5.3.1	Route Reflection	591
5.3.2	BGP Confederations.....	593
5.4	BGP Messages.....	595
5.4.1	Open Message	595
5.4.1.1	Changing the Autonomous System Number	596
5.4.1.2	Changing a Confederation Number.....	597
5.4.1.3	BGP Advertisement.....	597
5.4.2	Update Message	597
5.4.3	Keepalive Message	598
5.4.4	Notification Message	598
5.4.4.1	UPDATE Message Error Handling	598
5.4.5	Route Refresh Message.....	599
5.5	BGP Path Attributes	601
5.5.1	Origin	602
5.5.2	AS Path	602
5.5.2.1	AS Override	604
5.5.2.2	Using Local AS for ASN Migration.....	605
5.5.2.3	4-Octet Autonomous System Numbers	605
5.5.3	Next-Hop	607
5.5.3.1	Unlabeled IPv4 Unicast Routes.....	607
5.5.3.2	Unlabeled IPv6 Unicast Routes.....	608
5.5.3.3	VPN-IPv4 Routes	610
5.5.3.4	VPN-IPv6 Routes	611
5.5.3.5	Label-IPv4 Routes	613
5.5.3.6	Label-IPv6 Routes	614
5.5.3.7	Next-Hop Resolution	615
5.5.3.8	Next-Hop Tracking.....	621
5.5.3.9	Next-Hop Indirection.....	622
5.5.4	MED.....	622
5.5.4.1	Deterministic MED.....	622
5.5.5	Local Preference	623
5.5.6	Route Aggregation Path Attributes	623
5.5.7	Community Attributes	625
5.5.7.1	Standard Communities	625
5.5.7.2	Extended Communities	626
5.5.7.3	Large Communities	627
5.5.8	Route Reflection Attributes.....	628
5.5.9	Multi-Protocol BGP Attributes.....	628
5.5.10	4-Octet AS Attributes.....	630
5.5.11	AI GP Metric	630
5.6	BGP Routing Information Base (RIB).....	631

5.6.1	RIB-IN Features.....	631
5.6.1.1	BGP Import Policies	632
5.6.2	LOC-RIB Features.....	632
5.6.2.1	BGP Decision Process	633
5.6.2.2	BGP Route Installation in the Route Table.....	635
5.6.2.3	Weighted ECMP for BGP Routes.....	637
5.6.2.4	BGP Route Installation in the Tunnel Table	638
5.6.2.5	BGP Fast Reroute	639
5.6.2.6	QoS Policy Propagation via BGP (QPPB).....	641
5.6.2.7	BGP Policy Accounting.....	642
5.6.2.8	Route Flap Damping (RFD).....	643
5.6.3	RIB-OUT Features.....	644
5.6.3.1	BGP Export Policies	645
5.6.3.2	Outbound Route Filtering (ORF)	646
5.6.3.3	RT Constrained Route Distribution.....	648
5.6.3.4	Min Route Advertisement Interval (MRAI)	649
5.6.3.5	Advertise-Inactive	651
5.6.3.6	Best-External.....	651
5.6.3.7	Add-Paths	652
5.6.3.8	Split-Horizon	654
5.7	BGP Monitoring Protocol	655
5.8	BGP Applications.....	657
5.8.1	BGP Flow-Spec.....	657
5.8.1.1	Validating Received FlowSpec Routes	660
5.8.1.2	Using Flow Routes to Create Dynamic Filter Entries	661
5.8.2	Configuration of TTL Propagation for BGP Label Routes	661
5.8.2.1	TTL Propagation for RFC 3107 Label Route at Ingress LER	662
5.8.2.2	TTL Propagation for RFC 3107 Label Routes at LSR	662
5.8.3	BGP Prefix Origin Validation	664
5.8.4	BGP Route Leaking.....	666
5.8.5	BGP Optimal Route Reflection	668
5.8.6	LSP Tagging for BGP Next-Hops or Prefixes and BGP-LU	669
5.8.7	BGP-LS	670
5.8.7.1	Supported BGP-LS components	670
5.9	BGP Configuration Process Overview.....	673
5.10	Configuration Notes.....	674
5.10.1	General.....	674
5.10.1.1	BGP Defaults.....	674
5.10.1.2	BGP MIB Notes	674
5.11	Configuring BGP with CLI.....	677
5.11.1	Configuration Overview	677
5.11.1.1	Preconfiguration Requirements	677
5.11.1.2	BGP Hierarchy.....	677
5.11.1.3	Internal and External BGP Configuration	678
5.11.2	Basic BGP Configuration.....	678
5.11.3	Common Configuration Tasks	680
5.11.3.1	Creating an Autonomous System.....	681
5.11.3.2	Configuring a Router ID	681
5.11.3.3	BGP Confederations.....	683

5.11.3.4	BGP Router Reflectors	685
5.11.3.5	BGP Components	687
5.12	BGP Configuration Management Tasks	694
5.12.1	Modifying an AS Number	694
5.12.2	Modifying a Confederation Number	694
5.12.3	Modifying the BGP Router ID	695
5.12.4	Modifying the Router-Level Router ID	695
5.12.5	Deleting a Neighbor	696
5.12.6	Deleting Groups	697
5.13	BGP Command Reference	699
5.13.1	Command Hierarchies	699
5.13.1.1	Global BGP Commands	699
5.13.1.2	Group BGP Commands	704
5.13.1.3	Neighbor BGP Commands	707
5.13.1.4	BGP BMP Commands	711
5.13.1.5	Other BGP-Related Commands	712
5.13.2	Command Descriptions	713
5.13.2.1	Generic Commands	713
5.13.2.2	BGP Commands	714
5.13.2.3	BGP BMP Commands	798
5.13.2.4	Other BGP-Related Commands	805
5.14	Show, Clear, and Debug Command Reference	809
5.14.1	Command Hierarchies	809
5.14.1.1	Show Commands	809
5.14.1.2	Clear Commands	810
5.14.1.3	Debug Commands	811
5.14.2	Command Descriptions	812
5.14.2.1	Show Commands	812
5.14.2.2	Clear Commands	891
5.14.2.3	Debug Commands	894
6	Route Policies	903
6.1	Configuring Route Policies	903
6.1.1	Policy Statements	903
6.1.1.1	Policy Statement Chaining and Logical Expressions	904
6.1.1.2	Routing Policy Subroutines	906
6.1.1.3	Policy Evaluation Command	907
6.1.1.4	Exclusive Editing for Policy Configuration	907
6.1.1.5	Default Action Behavior	907
6.1.1.6	Denied IP Unicast Prefixes	908
6.1.1.7	Controlling Route Flapping	908
6.1.2	Regular Expressions	909
6.1.3	BGP and OSPF Route Policy Support	915
6.1.3.1	BGP Route Policies	915
6.1.3.2	Re-advertised Route Policies	916
6.1.3.3	Triggered Policies	916
6.1.3.4	Set MED to IGP Cost using Route Policies	917
6.1.3.5	BGP Policy Subroutines	918
6.1.3.6	Route Policies for BGP Next-Hop Resolution and Peer Tracking	918

6.1.4	Routing Policy Parameterization	918
6.1.5	When to Use Route Policies.....	925
6.2	Route Policy Configuration Process Overview	926
6.3	Configuration Notes.....	927
6.3.1	General.....	927
6.3.2	Policy Reference Checks	927
6.3.2.1	Known Limitations	929
6.4	Configuring Route Policies with CLI	931
6.4.1	Route Policy Configuration Overview	931
6.4.1.1	When to Create Routing Policies.....	931
6.4.1.2	Default Route Policy Actions	932
6.4.1.3	Policy Evaluation	933
6.4.1.4	Damping	935
6.4.2	Basic Configurations.....	936
6.4.3	Configuring Route Policy Components.....	937
6.4.3.1	Beginning the Policy Statement.....	938
6.4.3.2	Creating a Route Policy.....	938
6.4.3.3	Configuring a Default Action.....	939
6.4.3.4	Configuring an Entry.....	940
6.4.3.5	Configuring a Community List	941
6.4.3.6	Configuring Damping.....	941
6.4.3.7	Configuring a Prefix List	942
6.5	Route Policy Configuration Management Tasks	943
6.5.1	Editing Policy Statements and Parameters	943
6.5.2	Deleting an Entry.....	944
6.5.3	Deleting a Policy Statement	944
6.6	Route Policy Command Reference	945
6.6.1	Command Hierarchies.....	945
6.6.1.1	Route Policy Configuration Commands.....	945
6.6.1.2	Segment Routing Commands	950
6.6.2	Command Descriptions	950
6.6.2.1	Generic Commands.....	950
6.6.2.2	Segment Routing.....	952
6.6.2.3	Route Policy Options	959
6.6.2.4	Route Policy Damping Commands.....	967
6.6.2.5	Route Policy Prefix Commands.....	969
6.6.2.6	Route Policy Entry Match Commands.....	971
6.6.2.7	Route Policy Action Commands	992
6.7	Show, Clear, and Debug Command Reference	1011
6.7.1	Command Hierarchies.....	1011
6.7.1.1	Show Commands	1011
6.7.2	Command Descriptions	1011
6.7.2.1	Show Commands	1011
7	Standards and Protocol Support	1017

1 Getting Started

1.1 About This Guide

This guide describes routing protocols including multicast, RIP, OSPF, IS-IS, BGP, and route policies provided by the router and presents configuration and implementation examples.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e • 7750 SR-7s/14s • 7750 SR-1	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS 16.0.Rx Software Release Notes*, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 16.0.Rx content and may contain some content that will be released in later maintenance loads. Please refer to the *SR OS 16.0.Rx Software Release Notes*, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA, for information about features supported in each load of the Release 16.0.Rx software.

1.2 Router Configuration Process

[Table 2](#) lists the tasks necessary to configure RIP, OSPF, IS-IS, and BGP protocols, and route policies on the 7450 ESS, 7750 SR, and 7950 XRS. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 2 Configuration Process

Area	Task	Section
RIP configuration	Configure RIP	Configuring RIP with CLI
	RIP configuration management	RIP Configuration Management Tasks
OSPF configuration	Configure an LFA SPF policy	Loop-Free Alternate Shortest Path First (LFA SPF) Policies
	Configure LFA using backup node SID	LFA Protection Using Segment Routing Backup Node SID
	Configure OSPF	Configuring OSPF with CLI
	OSPF configuration management	OSPF Configuration Management Tasks
IS-IS configuration	Configure IS-IS	Configuring IS-IS with CLI
	Configure IS-IS components	Configuring IS-IS Components
	IS-IS configuration management	IS-IS Configuration Management Tasks
BGP configuration	Configure BGP	Configuring BGP with CLI
	BGP configuration management	BGP Configuration Management Tasks
Route Policies	Configure route policies	Configuring Route Policies with CLI
	Configure route policy components	Configuring Route Policy Components
	Route policy configuration management	Route Policy Configuration Management Tasks

2 RIP

2.1 RIP Overview

The Routing Information Protocol (RIP) is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using hop count as the metric. In order for the protocol to provide complete information on routing, every router in the domain must participate in the protocol.

RIP is a routing protocol based on a distance vector (Bellman-Ford) algorithm, which advertises network reachability by advertising prefix/mask and the metric (also known as hop count or cost). RIP selects the route with the lowest metric as the best route. RIP differs from link-state database protocols, such as OSPF and IS-IS, in that RIP advertises reachability information directly and link-state-database-based protocols advertise topology information. Each node is responsible for calculating the reachability information from the topology.

The router software supports RIPv1 and RIPv2. RIPv1, specified in RFC 1058, was written and implemented prior to the introduction of CIDR. It assumes the netmask information for non-local routes, based on the class the route belongs to:

- Class A – 8 bit mask
- Class B – 16 bit mask
- Class C – 24 bit mask

RIPv2 was written after CIDR was developed and transmits netmask information with every route. Because of the support for CIDR routes and other enhancements in RIPv2 such as triggered updates, multicast advertisements, and authentication, most production networks use RIPv2. However, there are some older systems (hosts and routers) that only support RIPv1, especially when RIP is used simply to advertise default routing information.

RIP is supported on all IP interfaces, including both network and access interfaces.

2.1.1 RIP Features

RIP, a UDP-based protocol, updates its neighbors, and the neighbors update their neighbors, and so on. Each host that uses RIP has a routing process that sends and receives datagrams on UDP port number 520.

Each RIP router advertises all RIP routes periodically via RIP updates. Each update can contain a maximum of 25 route advertisements. This limit is imposed by RIP specifications. RIP can sometimes be configured to send as many as 255 routes per update. The formats of the RIPv1 and RIPv2 updates are slightly different and are shown below. Additionally, RIPv1 updates are sent to a broadcast address, RIPv2 updates can be either sent to a broadcast or multicast address (224.0.0.9). RIPv2 supports subnet masks, a feature that was not available in RIPv1.

A network address of 0.0.0.0 is considered a default route. A default route is used when it is not convenient to list every possible network in the RIP updates, and when one or more closely-connected gateways in the system are prepared to handle traffic to the networks that are not listed explicitly. These gateways create RIP entries for the address 0.0.0.0, as if it were a network to which they are connected.

2.1.1.1 RIP Version Types

The router allows you to specify the RIP version that will be sent to RIP neighbors and RIP updates that will be accepted and processed. The router allows the following combinations:

- Send *only* RIPv1 or send *only* RIPv2 to either the broadcast or multicast address or send no messages.

The default sends RIPv2 formatted messages to the broadcast address.

- Receive *only* RIPv1, receive *only* RIPv2, or receive *both* RIPv1 and RIPv2, or receive none.

The default receives both.


2.1.1.2 RIPv2 Authentication

RIPv2 messages carry more information, which permit the use of a simple authentication mechanism to secure table updates. The router implementation enables the use of a simple password (plain text) or message digest (MD5) authentication.

2.1.1.3 RIP Packet Format

The RIP packet format is displayed in [Figure 1](#):

Figure 1 **RIP Packet Format**




A RIP packet consists of the following fields:

- **Command** — Indicates whether the packet is a request or a response message. The request asks the responding system to send all or part of its routing table. The response may be sent in response to a request, or it may be an unsolicited routing update generated by the sender.
- **Version** — The RIP version used. This field can signal different potentially incompatible versions.
- **Must be zero** — Not used in RIPv1. This field provides backward compatibility with pre-standard varieties of RIP. The default value is zero.
- **Address family identifier (AFI)** — The AFI is the type of address. RIP can carry routing information for several different protocols. Each entry in this field has an AFI to indicate the type of address being specified. The IP AFI is 2.
- **Address** — The IP address for the packet.
- **Metric** — Specifies the number of hops to the destination.
- **Mask** — Specifies the IP address mask.
- **Next hop** — Specifies the IP address of the next router along the path to the destination.

2.1.1.3.1 RIPv1 Format

There can be between 1 and 25 (inclusive) RIP entries. [Figure 2](#) displays RIPv1 format:


Figure 2 **RIPv1 Format**



2.1.1.3.2 **RIPv2 Format**

The RIP packet format is displayed in [Figure 3](#):

Figure 3 **RIPv2 Format**



The RIPv2 packets include the following fields:

- Subnet mask — The subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.
- Next hop —The IP address of the next hop to forward packets.

2.1.1.4 **BFD Monitoring of RIP Neighbor Liveliness**

BFD can be used to monitor the liveliness of the RIP neighbors. If a BFD session, associated with a RIP neighbor fails, that RIP neighbor is declared down and all routes learned from that RIP neighbor are withdrawn from the associated route tables.

BFD is enabled for RIP by configuring **enable-bfd** within the RIP neighbor configuration context.

BFD must be also be enabled on the interface associated with the RIP neighbor. The **bfd** command sets the necessary transmit and receive intervals as well as sets the optional multiplier.

2.2 RIPng

RIPng is the IPv6 form of the interior gateway protocol (IGP) Routing Information Protocol (RIP), originally implemented for IPv4 routing. This protocol is a distance vector routing protocol that periodically advertises IPv6 routing information to neighbors, typically through the use of UDP based multicast updates carrying a list of one or more entries, each containing an IPv6 prefix, prefix length, route metric and a possible route tag.

RIPng is supported in the base routing context and also as a PE-CE routing protocol within a VPRN context.

2.2.1 RIPng Protocol

RIPng packets are sent using the UDP protocol and the protocol port number 521. Unsolicited updates messages are sent with 521 as both the source and destination port.

- Source IP address
 - The Link-Local IPv6 address of the interface sending the RIPng packet is used as the source IP address of any RIPng update sent.
- Destination IP address
 - The destination IP for any periodic or triggered update should be sent to the multicast group FF02::9, (all-rip-routers multicast group).
 - When sending responses to an RIPng request, the RIPng response is sent to the unicast IP address of the requester.

Each route entry in an update message contains the following:

- IPv6 prefix
- Prefix length
- Route metric
- Route tag (optional)

2.3 Common Attributes

2.3.1 Metrics

By default, RIP advertises all RIP routes to each peer every 30 seconds. RIP uses a hop count metric to determine the distance between the packet's source and destination. The metric/cost values for a valid route is 1 through 15. A metric value of 16 (infinity) indicates that the route is no longer valid and should be removed from the router's routing table.

Each router along the path increments the hop count value by 1. When a router receives a routing update with new or different destination information, the metric increments by 1.

The maximum number of hops in a path is 15. If a router receives a routing update with a metric of 15 and contains a new or modified entry, increasing the metric value by 1 will cause the metric increment to 16 (infinity). Then, the destination is considered unreachable.

The router implementation of RIP uses *split horizon* with *poison reverse* to protect from such problems as "counting to infinity". Split horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

2.3.2 Timers

RIP uses numerous timers to determine how often RIP updates are sent and how long routes are maintained.

- Update — Times the interval between periodic routing updates.
- Timeout — This timer is initialized when a route is established and any time an update message is received for the route. When this timer expires, the route is no longer valid. It is retained in the table for a short time, so that neighbors can be notified that the route has been dropped.
- Flush — When the flush timer expires, the route is removed from the tables.

2.3.3 Import and Export Policies

Routing policies can control the content of the routing tables, the routes that are advertised and the best route to take to reach a destination. Import route policies determine which routes are accepted from RIP neighbors. Export route policies determine which routes are exported from the route table to RIP. By default, RIP does not export routes it has learned to its neighbors.

There are no default routing policies. A policy must be created explicitly and applied to a RIP import or export command.

2.3.4 Hierarchical Levels

The minimum RIP configuration must define one group and one neighbor. The parameters configured on the global level are inherited by the group and neighbor levels. Parameters can be modified and overridden on a level-specific basis. RIP command hierarchy consists of three levels:


- Global
- Group
- Neighbor

Many of the hierarchical RIP commands can be modified on different levels. The most specific value is used. That is, a RIP group-specific command takes precedence over a global RIP command. A neighbor-specific statement takes precedence over a global RIP and group-specific command; for example, if you modify a RIP neighbor-level command default, the new value takes precedence over group- and global-level settings.

2.4 RIP Configuration Process Overview

Figure 4 displays the process to configure RIP parameters.

Figure 4 RIP Configuration and Implementation Flow



2.5 Configuration Notes

This section describes RIP configuration caveats.

2.5.1 General

Before RIP neighbor parameters can be configured, router interfaces must be configured.

RIP must be explicitly created for each router interface. There are no default RIP instances on a router.

2.6 Configuring RIP with CLI

This section provides information to configure Routing Information Protocol (RIP) using the command line interface.

2.6.1 RIP and RIPng Configuration Overview

2.6.1.1 Preconfiguration Requirements

Configure the following entities before beginning the RIP configuration:

- (Optional) Policy statements should be defined in the **config>router>policy-options** context.

2.6.1.2 RIP Hierarchy

RIP is configured in the **config>router>rip** context. RIP is not enabled by default. Three hierarchical levels are included in RIP configurations:

- Global
- Group
- Neighbor

Commands and parameters configured on the global level are inherited by the group and neighbor levels although parameters configured on the group and neighbor levels take precedence over global configurations.

2.6.2 Basic RIP Configuration

This section provides information to configure RIP and examples of common configuration tasks. For a router to accept RIP updates, in the **config>router>rip** context, you must define at least one group and one neighbor. A router will ignore updates received from routers on interfaces not configured for RIP. Configuring other RIP commands and parameters are optional.

By default, the local router imports all routes from this neighbor and does not advertise routes. The router receives both RIPv1 and RIPv2 update messages with 25 to 255 route entries per message.

The RIP configuration commands have three primary configuration levels: **rip** for global configurations, **group group-name** for RIP group configurations, and **neighbor ip-int-name** for RIP neighbor configurations. Within the different levels, the configuration commands are identical. For the repeated commands, the command that is most specific to the neighboring router is in effect; that is, neighbor settings have precedence over group settings which have precedence over RIP global settings.

The minimal RIP parameters that need to be configured in the **config>router>rip** context are:

- Group
- Neighbor

The following example displays a basic RIP configuration.

```
ALA-A>config>router>rip# info
-----
group "RIP-ALA-A"
    neighbor "to-ALA-4"
    exit
exit
-----
ALA-A>config>router>rip#
```

2.6.3 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure RIP and provides the CLI commands.

Configure RIP hierarchically using the global level (applies to all peers), the group level (applies to all peers in peer-group), or the neighbor level (only applies to the specified interface). By default, group members inherit the group's configuration parameters although a parameter can be modified on a per-member basis without affecting the group-level parameters.

Many of the hierarchical RIP commands can be used on different levels. The most specific value is used. That is, a RIP group-specific command takes precedence over a global RIP command. A neighbor-specific statement takes precedence over a global RIP or group-specific command.

All RIP instances must be explicitly created on each device. Once created, RIP is administratively enabled.

To configure RIP, perform the following tasks:

- Step 1.** Configure interfaces
- Step 2.** Configure policy statements (optional)
- Step 3.** Enable RIP
- Step 4.** Configure group parameters
- Step 5.** Configure neighbor parameters

2.6.3.1 Configuring Interfaces

The following command sequences create a logical IP interface. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG), or the system. For more information about configuring interfaces, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

To configure a network interface:

CLI Syntax:

```
config router
    interface ip-int-name
        address ip-addr{/mask-length | mask}
        [broadcast {all-ones | host-ones}]
        port port-id
```

The following example displays router interface configuration command usage:

Example:

```
config>router> interface "to-ALA-4"
config>router>if$ address 10.10.12.1/24
config>router>if# port 1/1/1
config>router>if# exit
```

The following example displays the IP configuration output showing the interface information.

```
ALA-3>config>router# info
#-----
echo "IP Configuration "
#-----
    interface "system"
        address 10.10.10.103/32
    exit
    interface "to-ALA-4"
        address 10.10.12.1/24
        port 1/1/1
```

```

                                exit
#-----
ALA-3>config>router#

```

2.6.3.2 Configuring a Route Policy

The import route policy command allows you to filter routes being imported by the local router from its neighbors. If no match is found, the local router does not import any routes.

The export route policy command allows you to determine which routes are exported from the route table to RIP. By default, RIP does not export routes it has learned to its neighbors. If no export policy is specified, non-RIP routes will not be exported from the routing table manager to RIP.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

This section only provides brief instructions to configure route policies. For more details, refer to the *Route Policy Overview* chapter.

To enter the mode to create or edit route policies, you must enter the begin keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves and enables changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

Use the following CLI syntax to configure a policy to use for the RIP global, group, and neighbor **import** and **export** commands.

```

CLI Syntax:  config>router>policy-options
                  begin
                  commit
                  abort
                  policy-statement name
                      description text
                      default-action {accept|next-entry|next-
                        policy|drop|reject}
                      entry entry-id
                          description text
                          action {accept|next-entry|next-
                            policy|drop|reject}

```

from
to

Use the following CLI syntax to enter the edit mode:

CLI Syntax: config>router>policy-options
 begin

The following example displays some commands to configure a policy statement. Policy option commands are configured in the config>router context. Use the commit command to save the changes.

Example: config>router>policy-options# begin
 policy-options# policy-statement "RIP policy"
 policy-options>policy-statement\$ description "this is a
 test RIP policy"
 policy-options>policy-statement>default# entry 1
 policy-options>policy-statement>entry\$ action accept
 policy-options>policy-statement>entry# exit
 policy-options>policy-statement# default-action reject
 policy-options>policy-statement# exit
 policy-options# commit

```
ALA-A>config>router>policy-options# info
-----
policy-statement "RIP-policy"
description "this is a test RIP policy"
entry 1
action accept
exit
exit
default-action reject
exit
-----
ALA-A>config>router>policy-options>policy-statement#
```

2.6.3.3 Configuring RIP Parameters

Use the CLI syntax displayed below for:

- [Configuring Global-Level Parameters](#)
- [Configuring Group-Level Parameters](#)
- [Configuring Neighbor-Level Parameters](#)

CLI Syntax: config>router
 rip

```
authentication-key [authentication-key | hash-  
key [hash|hash2]  
authentication-type {none | password | message-  
digest | message-digest-20}  
check-zero {enable | disable}  
description string  
export policy-name [policy-name ...up to 5 max]  
import policy-name [policy-name ...up to 5 max]  
message-size number  
metric-in metric  
metric-out metric  
preference number  
receive {both | none | version-1 | version-2}  
send {broadcast | multicast | none | version-1  
| both}  
no shutdown  
split-horizon {enable | disable}  
timers update timeout flush  
  
group group-name  
    authentication-key [authentication-key |  
        hash-key [hash | hash2]  
    authentication-type {none | password |  
        message-digest | message-digest-20}  
    check-zero {enable | disable}  
    description string  
    export policy-name [policy-name ...up to 5  
        max]  
    import policy-name [policy-name ...up to 5  
        max]  
    message-size number  
    metric-in metric  
    metric-out metric  
    preference number  
    receive {both | none | version-1 |  
        version-2}  
    send {broadcast | multicast | none |  
        version-1}  
    no shutdown  
    split-horizon {enable | disable}  
    timers update timeout flush  
  
neighbor ip-int-name  
    authentication-key [authentication-  
        key | hash-key [hash | hash2]  
    authentication-type {none | password  
        | message-digest | message-  
        digest-20}  
    check-zero {enable | disable}
```



```

description string
export policy-name [policy-name ...up
    to 5 max]]
import policy-name [policy-name ...up
    to 5 max]]
message-size number
metric-in metric
metric-out metric
preference number
receive {both | none | version-1 |
    version-2}
send {broadcast | multicast | none |
    version-1}
split-horizon {enable | disable}
timers update timeout flush
no shutdown

```

2.6.3.4 Configuring Global-Level Parameters

After the RIP protocol instance is created, the no shutdown command is not required because RIP is administratively enabled upon creation. To enable RIP on a router, at least one group and one neighbor must be configured. There are no default groups or neighbors. Each group and neighbor must be explicitly configured.



Note: Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor-levels. Because the RIP commands are hierarchical, analyze the values that can disable features on a particular level.

Use the following CLI syntax to configure global-level RIP parameters.

CLI Syntax:

```

config>router
    rip
        authentication-key [authentication-key | hash-
            key [hash | hash2]
        authentication-type {password | message-
            digest}
        check-zero {enable | disable}
        export policy-name [policy-name ...up to 5 max]
        import policy-name [policy-name ...up to 5 max]
        message-size number
        metric-in metric
        metric-out metric
        preference number
        receive {both | none | version-1 | version-2}

```

```

send {broadcast | multicast | none | version-1
    | both}
no shutdown
split-horizon {enable | disable}
timers update timeout flush

```

The following example displays global RIP configuration command usage:

Example:

```

config>router# rip
config>router>rip# authentication-type password
config>router>rip# authentication-key test123
config>router>rip# receive both
config>router>rip# split-horizon enable
config>router>rip# timers 300 600 600
config>router>rip>group# exit

```

The following example displays the RIP group configuration:

```

ALA-A>config>router>rip# info
-----
authentication-type simple
authentication-key "ac1865lvz1d" hash
timers 300 600 600
-----
ALA-A>config>router>rip#

```

2.6.3.5 Configuring Group-Level Parameters

A group is a collection of related RIP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis. Use the following CLI syntax to configure a group.

CLI Syntax:

```

config>router# rip
group group-name
authentication-key [authentication-key | hash-
    key [hash | hash2]
authentication-type {password | message-digest}
check-zero {enable | disable}
description string
export policy-name [policy-name ...]
import policy-name [policy-name ...]
message-size number
metric-in metric

```

```
metric-out metric
preference number
receive {both|none|version-1|version-2}
        send {broadcast|multicast|none|version-
            1|both}
no shutdown
split-horizon {enable|disable}
timers update timeout flush
```

The following example displays group configuration command usage:

Example:

```
config>router# rip
config>router>rip# group headquarters
config>router>rip>group$ description "Mt. View"
config>router>rip>group# no shutdown
```

The following example displays the RIP group configuration:

```
ALA-A>config>router>rip# info
-----
authentication-type simple
authentication-key "ac1865lvz1d" hash
timers 300 600 600
group "headquarters"
    description "Mt. View"
exit
-----
ALA-A>config>router>rip#
```

2.6.3.6 Configuring Neighbor-Level Parameters

After you create a group name and assign options, add neighbor interfaces within the same group. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

Use the following CLI syntax to add a neighbor to a group and define options that override the same group-level command value.

CLI Syntax:

```
config>router# rip
group group-name
neighbor ip-int-name
    authentication-key [authentication-
        key|hash-key [hash|hash2]
    authentication-type {password|message-
        digest}
    check-zero {enable|disable}
    description string
    export policy-name [policy-name ...]
```

```
import policy-name [policy-name ...]
message-size number
metric-in metric
metric-out metric
preference number
receive {both|none|version-1|version-2}
send {broadcast|multicast|none|version-1}
split-horizon {enable|disable}
timers update timeout flush
no shutdown
```

The following example displays neighbor configuration command usage:

Example:

```
config>router# rip
config>router>rip# group headquarters1
config>router>rip>group# neighbor ferguson-274
config>router>rip>group>neighbor$ preference 255
config>router>rip>group>neighbor# send both
config>router>rip>group>neighbor# split-horizon enable
config>router>rip>group>neighbor# message-size 255
```

The following example displays the neighbor configured in group “headquarters”.

```
ALA-A>config>router>rip>group>neighbor# info
-----
message-size 255
preference 255
split-horizon enable
no timers
-----
ALA-A>config>router>rip>group>neighbor#
```

2.7 RIP Configuration Management Tasks

This section provides information about RIP configuration management tasks.

2.7.1 Modifying RIP Parameters

Modify, add or remove RIP parameters in the CLI. The changes are applied immediately. For the complete list of CLI commands, refer to [Configuring RIP Parameters](#).

CLI Syntax:

```
config>router# rip
      group group-name
      ...
      neighbor ip-int-name
      ...
```

Example:

```
config>router>rip# group "headquarters"
config>router>rip>group# neighbor "ferguson-274"
config>router>rip>group>neighbor# import RIPpolicy
config>router>rip>group>neighbor# message-size 150
```

The following example displays the updated parameters:

```
ALA-A>config>router>rip# info
-----
authentication-type simple
authentication-key "ac1865lvz1d" hash
timers 300 600 600
group "headquarters"
  description "Mt. View"
  neighbor "ferguson-274"
  import "RIPpolicy"
  message-size 150
  preference 255
  split-horizon enable
  no timers
exit
exit
-----
ALA-A>config>router>rip#
```

2.7.2 Deleting a Group

A group must be shut down before it can be deleted.

Use the following CLI syntax to shut down and then delete a group.

CLI Syntax: config>router# rip
 [no] group *group-name*
 shutdown

Example: config>router# rip
 config>router>rip# group "RIP-ALA-3"
 config>router>rip>group# shutdown
 config>router>rip>group# exit
 config>router>rip# no group "RIP-ALA-33"

Deleting the group without first shutting it down causes the following message to appear:

```
INFO: RIP #1204 group should be administratively down - virtual router index 1,group  
RIP-ALA-4
```

2.7.3 Deleting a Neighbor

The neighbor must be shut down before it can be deleted.

Use the following CLI syntax to delete a neighbor.

CLI Syntax: config>router# rip
 [no] group *group-name*
 [no] neighbor *ip-int-name*
 shutdown

Example: config>router# rip
 config>router>rip# group "RIP-ALA-4"
 config>router>rip>group# neighbor "to-ALA-3"
 config>router>rip>group>neighbor# shutdown
 config>router>rip>group>neighbor# exit
 config>router>rip>group# no neighbor "to-ALA-3"

Deleting the neighbor without first shutting it down causes the following message to appear:

```
INFO: RIP #1101 neighbor should be administratively down - virtual router index
```

2.8 RIP Configuration Command Reference

2.8.1 Command Hierarchies

- Configuration Commands
 - Group Commands
 - Neighbor Commands
- RIPv6 Configuration Commands
 - Group Commands
 - Neighbor Commands

2.8.1.1 Configuration Commands

```

config
  — router router-name
    — [no] rip
      — authentication-key {authentication-key | hash-key} [{hash | hash2}]
      — no authentication-key
      — authentication-type {none | password | message-digest | message-digest-20}
      — no authentication-type
      — [no] bfd-enable
      — check-zero {enable | disable}
      — no check-zero
      — description description-string
      — no description
      — export policy-name [policy-name]
      — no export
      — export-limit number [log percentage]
      — no export-limit
      — import policy-name [policy-name]
      — no import
      — message-size max-num-of-routes
      — no message-size
      — metric-in metric
      — no metric-in
      — metric-out metric
      — no metric-out
      — preference preference
      — no preference
      — receive receive-type
      — no receive
      — send send-type
      — no send

```

- [no] **shutdown**
- **split-horizon** {enable | disable}
- no **split-horizon**
- **timers** *update timeout flush*
- no **timers**

2.8.1.1.1 Group Commands

```

config
  — router router-name
    — [no] rip
      — [no] group group-name
        — authentication-key {authentication-key | hash-key} [{hash | hash2}]
        — no authentication-key
        — authentication-type {none | password | message-digest | message-
          digest-20}
        — no authentication-type
        — [no] bfd-enable
        — check-zero {enable | disable}
        — no check-zero
        — description description-string
        — no description
        — export policy-name [policy-name]
        — no export
        — import policy-name [policy-name]
        — no import
        — message-size max-num-of-routes
        — no message-size
        — metric-in metric
        — no metric-in
        — metric-out metric
        — no metric-out
        — preference preference
        — no preference
        — receive receive-type
        — no receive
        — send send-type
        — no send
        — [no] shutdown
        — split-horizon {enable | disable}
        — no split-horizon
        — timers update timeout flush
        — no timers

```

2.8.1.1.2 Neighbor Commands

```

config
  — router router-name
    — [no] rip

```


- [no] **group** *group-name*
 - [no] **neighbor** *ip-int-name*
 - **authentication-key** {*authentication-key* | *hash-key*} [{*hash* | *hash2*}]
 - no **authentication-key**
 - **authentication-type** {*none* | *password* | *message-digest* | *message-digest-20*}
 - no **authentication-type**
 - [no] **bfd-enable**
 - **check-zero** {*enable* | *disable*}
 - no **check-zero**
 - **description** *description-string*
 - no **description**
 - **export** *policy-name* [*policy-name*]
 - no **export**
 - **import** *policy-name* [*policy-name*]
 - no **import**
 - **message-size** *max-num-of-routes*
 - no **message-size**
 - **metric-in** *metric*
 - no **metric-in**
 - **metric-out** *metric*
 - no **metric-out**
 - **preference** *preference*
 - no **preference**
 - **receive** *receive-type*
 - no **receive**
 - **send** *send-type*
 - no **send**
 - [no] **shutdown**
 - **split-horizon** {*enable* | *disable*}
 - no **split-horizon**
 - **timers** *update timeout flush*
 - no **timers**
 - [no] **unicast-address** *ip-address*

2.8.1.2 RIPng Configuration Commands

- config
 - **router** *router-name*
 - [no] **ripng**
 - [no] **bfd-enable**
 - **check-zero** {*enable* | *disable*}
 - no **check-zero**
 - **description** *description-string*
 - no **description**
 - **export** *policy-name* [*policy-name*]
 - no **export**
 - **export-limit** *number* [*log percentage*]
 - no **export-limit**
 - **import** *policy-name* [*policy-name*]

- **no import**
- **message-size** *max-num-of-routes*
- **no message-size**
- **metric-in** *metric*
- **no metric-in**
- **metric-out** *metric*
- **no metric-out**
- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- **[no] shutdown**
- **split-horizon** {**enable** | **disable**}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**

2.8.1.2.1 Group Commands

- ```
config
 — router router-name
 — [no] ripng
 — [no] group group-name
 — [no] bfd-enable
 — check-zero {enable | disable}
 — no check-zero
 — description description-string
 — no description
 — export policy-name [policy-name]
 — no export
 — import policy-name [policy-name]
 — no import
 — message-size max-num-of-routes
 — no message-size
 — metric-in metric
 — no metric-in
 — metric-out metric
 — no metric-out
 — preference preference
 — no preference
 — receive receive-type
 — no receive
 — send send-type
 — no send
 — [no] shutdown
 — split-horizon {enable | disable}
 — no split-horizon
 — timers update timeout flush
 — no timers
```

### 2.8.1.2.2 Neighbor Commands

```

config
 — router router-name
 — [no] ripng
 — [no] group group-name
 — [no] neighbor ip-int-name
 — [no] bfd-enable
 — check-zero {enable | disable}
 — no check-zero
 — description description-string
 — no description
 — export policy-name [policy-name]
 — no export
 — import policy-name [policy-name]
 — no import
 — message-size max-num-of-routes
 — no message-size
 — metric-in metric
 — no metric-in
 — metric-out metric
 — no metric-out
 — preference preference
 — no preference
 — receive receive-type
 — no receive
 — send send-type
 — no send
 — [no] shutdown
 — split-horizon {enable | disable}
 — no split-horizon
 — timers update timeout flush
 — no timers
 — [no] unicast-address ipv6-address

```

## 2.8.2 Command Descriptions

### 2.8.2.1 Generic Commands

#### description

|                |                                                                       |
|----------------|-----------------------------------------------------------------------|
| <b>Syntax</b>  | <b>description</b> <i>description-string</i><br><b>no description</b> |
| <b>Context</b> | config>router>rip<br>config>router>rip>group                          |

---

|                    |                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <pre>config&gt;router&gt;rip&gt;group&gt;neighbor config&gt;router&gt;ripng config&gt;router&gt;ripng&gt;group config&gt;router&gt;ripng&gt;group&gt;neighbor</pre>                                                                                                                                                                                        |
| <b>Description</b> | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The <b>no</b> form of the command removes any description string from the context.</p> |
| <b>Default</b>     | no description                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p>                                                  |

## shutdown

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | <pre>config&gt;router&gt;rip config&gt;router&gt;rip&gt;group config&gt;router&gt;rip&gt;group&gt;neighbor config&gt;router&gt;ripng config&gt;router&gt;ripng&gt;group config&gt;router&gt;ripng&gt;group&gt;neighbor</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command administratively disables an entity. Downing an entity does not change, reset or remove any configuration settings or statistics. Many objects must be shutdown before they may be deleted.</p> <p>The <b>shutdown</b> command administratively downs an entity. Administratively downing an entity changes the operational state of the entity to down and the operational state of any entities contained within the administratively down entity.</p> <p>Unlike other commands and parameters where the default state will not be indicated in the configuration file, <b>shutdown</b> and <b>no shutdown</b> are always indicated in system generated configuration files.</p> <p>The <b>no</b> form of the command puts an entity into the administratively enabled state.</p> |

- Special Cases**
- RIP Global** — In the **config>router>rip** context, the **shutdown** command administratively enables/disables the RIP protocol instance. If RIP is globally shutdown, then all RIP group and neighbor interfaces transition to the operationally down state. Routes learned from a neighbor that is shutdown are immediately removed from the RIP database and route table manager (RTM). A RIP protocol instance is administratively enabled by default.
- RIP Group** — In the **config>router>rip>group group-name** context, the **shutdown** command administratively enables/disables the RIP group. If a RIP group is shutdown, all member neighbor interfaces transition to the operationally down state. Routes learned from a neighbor that is shutdown are immediately removed from the RIP database and route table manager (RTM). A RIP group is administratively enabled by default.
- RIP Neighbor** — In the **config>router>rip>group group-name>neighbor ip-int-name** context, the **shutdown** command administratively enables/disables the RIP neighbor interface. If a RIP neighbor is shutdown, the neighbor interface transitions to the operationally down state. Routes learned from a neighbor that is shutdown are immediately removed from the RIP database and route table manager (RTM). A RIP neighbor interface is administratively enabled by default.

## 2.8.2.2 RIP Commands

### rip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] rip                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command creates the context to configure the RIP protocol instance.</p> <p>When a RIP instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the <b>[no] shutdown</b> command.</p> <p>The <b>no</b> form of the command deletes the RIP protocol instance removing all associated configuration parameters.</p> |
| <b>Default</b>     | no rip                                                                                                                                                                                                                                                                                                                                                                                                               |

### ripng

|                |               |
|----------------|---------------|
| <b>Syntax</b>  | [no] ripng    |
| <b>Context</b> | config>router |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command creates the context to configure the RIPv6 protocol instance.</p> <p>When a RIPv6 instance is created, the protocol is enabled by default. To start or suspend execution of the RIPv6 protocol without affecting the configuration, use the <b>[no] shutdown</b> command.</p> <p>The <b>no</b> form of the command deletes the RIPv6 protocol instance removing all associated configuration parameters.</p> |
| <b>Default</b>     | no rIPv6                                                                                                                                                                                                                                                                                                                                                                                                                     |

## authentication-key

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-key</b> { <i>authentication-key</i>   <i>hash-key</i> } [{ <b>hash</b>   <b>hash2</b> }]<br><b>no authentication-key</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>rIPv6<br>config>router>rIPv6>group<br>config>router>rIPv6>group>neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command sets the authentication password to be passed between RIPv6 neighbors.</p> <p>The authentication type and authentication key must match exactly for the RIPv6 message to be considered authentic and processed.</p> <p>The <b>no</b> form of the command removes the authentication password from the configuration and disables authentication.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>     | no authentication-key                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>authentication-key</i> — Specifies the authentication key. Allowed values are any string up to 16 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>hash-key</i> — Specifies the hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> |

**hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

## authentication-type

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-type</b> { <b>none</b>   <b>password</b>   <b>message-digest</b>   <b>message-digest-20</b> }<br><b>no authentication-type</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command sets the type of authentication to be used between RIP neighbors.</p> <p>The type and password must match exactly for the RIP message to be considered authentic and processed.</p> <p>The <b>no</b> form of the command removes the authentication type from the configuration and effectively disables authentication.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>     | no authentication-type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><b>none</b> — The <b>none</b> parameter explicitly disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.</p> <p><b>password</b> — Specifies that the password enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple <b>password</b> authentication is enabled.</p> <p><b>message-digest</b> — Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one <b>message-digest-key</b> must be configured.</p> <p><b>message-digest-20</b> — Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, <i>RIP-2 MD5 Authentication</i>. If this option is configured, then at least one <b>message-digest-key</b> must be configured.</p> |

## bfd-enable

|                |                                                                                                                    |
|----------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | [no] <b>bfd-enable</b>                                                                                             |
| <b>Context</b> | config>router>rip<br>config>router>rip>group>neighbor<br>config>router>ripng<br>config>router>ripng>group>neighbor |

**Description** This command enables bi-directional forwarding (BFD) to control the state of the associated protocol adjacency. By enabling BFD on a given protocol interface, the state of the RIP neighbor is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set using the **bfd** command under the IP interface configuration context.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

## check-zero

**Syntax** **check-zero {enable | disable}**  
**no check-zero**

**Context** config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor  
config>router>ripng  
config>router>ripng>group  
config>router>ripng>group>neighbor

**Description** This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

The **check-zero enable** command enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages.

The **check-zero disable** command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (**no check-zero**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command removes the **check-zero** command from the configuration.

**Special Cases** **RIP Global** — By default, check-zero is disabled at the global RIP instance level.

**Parameters** **enable** — Specifies to reject RIP messages which do not have zero in the RIPv1 and RIPv2 mandatory fields.

**disable** — Specifies allows receipt of RIP messages which do not have the mandatory zero fields reset.



## export

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export</b> <i>policy-name</i> [ <i>policy-name</i> ]<br><b>no export</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor<br>config>router>ripng<br>config>router>ripng>group<br>config>router>ripng>group>neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command specifies the export route policies used to determine which routes are exported to RIP.</p> <p>If no export policy is specified, non-RIP routes will not be exported from the routing table manager to RIP. RIP-learned routes will be exported to RIP neighbors.</p> <p>If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.</p> <p>The <b>no</b> form of the command removes all policies from the configuration.</p> |
| <b>Default</b>     | no export                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>policy-name</i> — Specifies up to five export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on.), the entire string must be enclosed within double quotes.</p> <p>The specified names must already be defined.</p>                                                                                                                                                                                                                                                                                                                       |

## export-limit

|                    |                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-limit</b> <i>number</i> [ <i>log percentage</i> ]<br><b>no export-limit</b>                                                                                                                              |
| <b>Context</b>     | config>router>rip<br>config>router>ripng                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.</p> <p>The <b>no</b> form of the command removes the parameters from the configuration.</p> |
| <b>Default</b>     | no export-limit                                                                                                                                                                                                    |

---

|                   |                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.                 |
| <b>Values</b>     | 1 to 4294967295                                                                                                                       |
|                   | <i>percentage</i> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent. |
| <b>Values</b>     | 1 to 100                                                                                                                              |

## group

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] group</b> <i>group-name</i>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>rip<br>config>router>ripng                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command creates a context for configuring a RIP group of neighbor interfaces.</p> <p>RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.</p> <p>The <b>no</b> form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.</p> |
| <b>Default</b>     | no group                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>group-name</i> — Specifies the RIP group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.                                                                                                                                                        |

## import

|                    |                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>import</b> <i>policy-name</i> [ <i>policy-name</i> ]<br><b>no import</b>                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor<br>config>router>ripng<br>config>router>ripng>group<br>config>router>ripng>group>neighbor                                                                                                                                          |
| <b>Description</b> | This command configures import route policies to determine which routes are accepted from RIP neighbors. If no import policy is specified, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics. |

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

|                   |                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no import                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b> | <p><i>policy-name</i> — Specifies up to five import route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p>The specified names must already be defined.</p> |

## message-size

|                    |                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>message-size</b> <i>max-num-of-routes</i></p> <p><b>no message-size</b></p>                                                                                                                                                                           |
| <b>Context</b>     | <p>config&gt;router&gt;rip</p> <p>config&gt;router&gt;rip&gt;group</p> <p>config&gt;router&gt;rip&gt;group&gt;neighbor</p> <p>config&gt;router&gt;ripng</p> <p>config&gt;router&gt;ripng&gt;group</p> <p>config&gt;router&gt;ripng&gt;group&gt;neighbor</p> |
| <b>Description</b> | <p>This command configures the maximum number of routes per RIP update message.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>                                                                                                  |
| <b>Default</b>     | message-size 25                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>max-num-of-routes</i> — The maximum number of RIP routes per RIP update message expressed as a decimal integer.</p> <p><b>Values</b> 25 to 255</p>                                                                                                    |

## metric-in

|                |                                                                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <p><b>metric-in</b> <i>metric</i></p> <p><b>no metric-in</b></p>                                                                                                                                      |
| <b>Context</b> | <p>config&gt;router&gt;rip</p> <p>config&gt;router&gt;rip&gt;group</p> <p>config&gt;router&gt;rip&gt;group&gt;neighbor</p> <p>config&gt;router&gt;ripng</p> <p>config&gt;router&gt;ripng&gt;group</p> |

---

```
config>router>ripng>group>neighbor
```

|                    |                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures the metric added to routes received from a RIP neighbor.</p> <p>When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the <b>metric-in</b> and <b>metric-out</b> values.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | metric-in 1                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>metric</i> — Specifies the value added to the metric of routes received from a RIP neighbor expressed as a decimal integer.</p> <p><b>Values</b>      1 to 16</p>                                                                                                                                                                                          |

## metric-out

|                    |                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>metric-out</b> <i>metric</i><br><b>no metric-out</b>                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor<br>config>router>ripng<br>config>router>ripng>group<br>config>router>ripng>group>neighbor                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures the metric assigned to routes exported into RIP and advertised to RIP neighbors.</p> <p>When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the <b>metric-in</b> and <b>metric-out</b> values.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | metric-out 1                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>metric</i> — Specifies the value added to the metric for routes exported into RIP and advertised to RIP neighbors expressed as a decimal integer.</p> <p><b>Values</b>      1 to 16</p>                                                                                                                                                                                            |

## neighbor

|                |                                                      |
|----------------|------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] neighbor</b> <i>ip-int-name</i>              |
| <b>Context</b> | config>router>rip>group<br>config>router>ripng>group |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command creates a context for configuring a RIP neighbor interface.</p> <p>By default, interfaces are not activated in any interior gateway protocol, such as RIP, unless explicitly configured.</p> <p>The <b>no</b> form of the command deletes the RIP interface configuration for this interface. The <b>shutdown</b> command in the <b>config&gt;router&gt;rip&gt;group group-name&gt;neighbor ip-int-name</b> context can be used to disable an interface without removing the configuration for the interface.</p>                                                                                                                          |
| <b>Default</b>     | no neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>ip-int-name</i> — Specifies the IP interface name. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service ies interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p>If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.</p> |

## preference

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>preference preference</b></p> <p><b>no preference</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | <p>config&gt;router&gt;rip</p> <p>config&gt;router&gt;rip&gt;group</p> <p>config&gt;router&gt;rip&gt;group&gt;neighbor</p> <p>config&gt;router&gt;ripng</p> <p>config&gt;router&gt;ripng&gt;group</p> <p>config&gt;router&gt;ripng&gt;group&gt;neighbor</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command configures the preference for RIP routes.</p> <p>A route can be learned by the router from different protocols in which case the costs are not comparable. When this occurs the preference is used to decide which route will be used.</p> <p>Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in <a href="#">Table 3</a>. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.</p> <p>If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the <b>ecmp</b> in the <b>config&gt;router</b> context.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | preference 100                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Parameters** *preference* — Specifies the preference for RIP routes expressed as a decimal integer. Defaults for different route types are listed in [Table 3](#).

**Table 3 Route Preference Defaults by Route Type**

| Route Type             | Preference | Configurable |
|------------------------|------------|--------------|
| Direct attached        | 0          | No           |
| Static routes          | 5          | Yes          |
| OSPF internal          | 10         | Yes          |
| IS-IS level 1 internal | 15         | Yes          |
| IS-IS level 2 internal | 18         | Yes          |
| RIP                    | 100        | Yes          |
| OSPF external          | 150        | Yes          |
| IS-IS level 1 external | 160        | Yes          |
| IS-IS level 2 external | 165        | Yes          |
| BGP                    | 170        | Yes          |

**Values** 0 to 255

## receive

**Syntax** **receive {both | none | version-1 | version-2}**  
**no receive**

**Context** config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor  
config>router>ripng  
config>router>ripng>group  
config>router>ripng>group>neighbor

**Description** This command configures the types of RIP updates that will be accepted and processed.

If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accept packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior is to accept and process both RIPv1 and RIPv2 messages.

The **no** form of the command reverts to the default value.

|                   |                                                                                                    |
|-------------------|----------------------------------------------------------------------------------------------------|
| <b>Default</b>    | receive both – in the config>router>rip context                                                    |
|                   | receive version-1 – in the config>router>ripng context                                             |
| <b>Parameters</b> | <b>both</b> — Specifies that RIP updates in either version 1 or version 2 format will be accepted. |
|                   | <b>none</b> — Specifies that RIP updates will not be accepted.                                     |
|                   | <b>version-1</b> — Specifies that RIP updates in version 1 format only will be accepted.           |
|                   | <b>version-2</b> — Specifies that RIP updates in version 2 format only will be accepted.           |

## send

|                    |                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>send {broadcast   multicast   none   version-1}</b><br><b>no send</b>                                            |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor                                    |
| <b>Description</b> | This command specifies the type of RIP messages sent to RIP neighbors.                                              |
|                    | If <b>version-1</b> is specified, the router need only listen for and accept packets sent to the broadcast address. |
|                    | This control can be issued at the global, group or interface level.                                                 |
|                    | The <b>no</b> form of the command reverts to the default value.                                                     |
| <b>Default</b>     | send version-1                                                                                                      |
| <b>Parameters</b>  | <b>broadcast</b> — Specifies send RIPv2 formatted messages to the broadcast address.                                |
|                    | <b>multicast</b> — Specifies send RIPv2 formatted messages to the multicast address.                                |
|                    | <b>none</b> — Specifies not to send any RIP messages (i.e. silent listener).                                        |
|                    | <b>version-1</b> — Specifies send RIPv1 formatted messages to the broadcast address.                                |

## send

|                |                                                                                        |
|----------------|----------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>send {none   ripng   unicast}</b><br><b>no send</b>                                 |
| <b>Context</b> | config>router>ripng<br>config>router>ripng>group<br>config>router>ripng>group>neighbor |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command specifies if RIPng are sent to RIP neighbors or not and what type of IPv6 address is to be used to deliver the messages.</p> <p>This control can be issued at the global, group or interface level.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>                                                                                                                          |
| <b>Default</b>     | send ripng                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><b>ripng</b> — Specifies RIPng messages to be sent to the standard multicast address (FF02::9).</p> <p><b>none</b> — Specifies not to send any RIPng messages (i.e. silent listener).</p> <p><b>unicast</b> — Specifies to send RIPng updates as unicast messages to the defined unicast address configured through the <b>unicast-address</b> command. This option is only allowed within the neighbor context.</p> |

## split-horizon

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>split-horizon {enable   disable}</b><br><b>no split-horizon</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>rip<br>config>router>rip>group<br>config>router>rip>group>neighbor<br>config>router>ripng<br>config>router>ripng>group<br>config>router>ripng>group>neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables the use of split-horizon.</p> <p>RIP uses split-horizon with poison-reverse to protect from such problems as “counting to infinity”. Split-horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).</p> <p>The <b>split-horizon disable</b> command enables split horizon without poison reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.</p> <p>This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (<b>no split-horizon</b>), the setting from the less specific level is inherited by the lower level.</p> <p>The <b>no</b> form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.</p> |
| <b>Default</b>     | enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



- Parameters**
- enable** — Specifies enable split horizon and poison reverse.
  - disable** — Specifies disable split horizon allowing routes to be re-advertised on the same interface on which they were learned with the advertised metric incremented by the **metric-in** value.

## timers

- Syntax** **timers** *update timeout flush*  
**no timers**
- Context** config>router>rip  
config>router>rip>group  
config>router>rip>group>neighbor  
config>router>ripng  
config>router>ripng>group  
config>router>ripng>group>neighbor
- Description** This command configures values for the update, timeout and flush RIP timers.
- The RIP update timer determines how often RIP updates are sent.
- If the route is not updated by the time the RIP timeout timer expires, the route is declared invalid but is maintained in the RIP database.
- The RIP flush timer determines how long a route is maintained in the RIP database after it has been declared invalid. Once the flush timer expires, the route is removed from the RIP database.
- The **no** form of the command reverts to the default values.
- Default** timers 30 180 120
- Parameters** *update* — Specifies the RIP update timer value in seconds expressed as a decimal integer.
- Values** 1 to 600
- timeout* — Specifies the RIP timeout timer value in seconds expressed as a decimal integer.
- Values** 1 to 1200
- flush* — Specifies the RIP flush timer value in seconds expressed as a decimal integer.
- Values** 1 to 1200

## unicast-address

- Syntax** [**no**] **unicast-address** *ipv6-address*

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>rip>group>neighbor<br>config>router>ripng>group>neighbor                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the unicast IPv6 address that RIP and RIPng update messages will be sent to if the <b>send</b> command is set to <b>send unicast</b>.</p> <p>Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address.</p> <p>The <b>no</b> form of the command deletes the specified IPv6 unicast address from the configuration.</p> |
| <b>Parameters</b>  | <i>ipv6-address</i> — Specifies the IPv6 unicast address to which unicast RIP or RIPng updates should be sent.                                                                                                                                                                                                                                                                                                                   |

## 2.9 Show, Clear, and Debug Command Reference

### 2.9.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

#### 2.9.1.1 Show Commands

```
show
 — router
 — bfd
 — session
 — rip
 — database [ip-prefix [/mask] [longer]] [peer ip-address] [detail [qos]]
 — group [name] [detail]
 — neighbor [ip-int-name | ip-address] [detail] [advertised-routes]
 — peer [interface-name]
 — statistics [ip-int-name | ip-address]
 — ripng
 — database [peer ipv6-address] [detail]
 — database [peer ipv6-address] detail [qos]
 — database ipv6-address[/prefix-length] [longer] [peer ipv6-address] [detail]
 — database ipv6-address[/prefix-length] [longer] [peer ipv6-address] [detail]
 [qos]
 — group [name] [detail]
 — neighbor [ip-int-name | ipv6-address] [detail]
 — neighbor [detail] advertised-routes [ip-int-name | ipv6-link-local-address]
 — peer [interface-name]
 — statistics [ip-int-name | ip-address]
```

#### 2.9.1.2 Clear Commands

```
clear
 — router
 — rip
 — database
 — export
 — statistics [neighbor ip-int-name | ip-address]
 — ripng
 — database
```

- **export**
- **export** [neighbor *ip-int-name* | *ipv6-address*]

### 2.9.1.3 Debug Commands

```

debug
 — router
 — rip
 — [no] auth [neighbor ip-int-name | ip-address]
 — [no] error [neighbor ip-int-name | ip-address]
 — [no] events [neighbor ip-int-name | ip-address]
 — [no] holddown [neighbor ip-int-name | ip-address]
 — [no] packets [neighbor ip-int-name | ip-address]
 — [no] requests [neighbor ip-int-name | ip-address]
 — [no] trigger [neighbor ip-int-name | ip-address]
 — [no] updates [neighbor ip-int-name | ip-address]

debug
 — router
 — ripng
 — [no] error [neighbor ip-int-name | ipv6-address]
 — [no] events [neighbor ip-int-name]
 — [no] holddown [neighbor ip-int-name | ipv6-address]
 — [no] packets [neighbor ip-int-name | ipv6-address]
 — [no] requests [neighbor ip-int-name | ipv6-address]
 — [no] trigger [neighbor ip-int-name | ipv6-address]
 — [no] updates [neighbor ip-int-name | ipv6-address]

```

## 2.9.2 Command Descriptions

### 2.9.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

#### database

|                    |                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database</b> [ <i>ip-prefix</i> [/mask] [longer]] [peer <i>ip-address</i> ] [detail [qos]] |
| <b>Context</b>     | show>router>rip                                                                               |
| <b>Description</b> | This command displays the routes in the RIP database.                                         |

**Output** Table 4 describes the route database output fields.

**Table 4 Database Output Fields**

| Label       | Description                                                                                                                                   |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Destination | The RIP destination for the route.                                                                                                            |
| Peer        | The router ID of the peer router.                                                                                                             |
| NextHop     | The IP address of the next hop.                                                                                                               |
| Metric      | The hop count to rate the value of different hops.                                                                                            |
| Tag         | The value to distinguish between internal routes (learned by RIP) and external routes (learned from other protocols).                         |
| TTL         | Displays how many seconds the specific route will remain in the routing table. When an entry reaches 0, it is removed from the routing table. |
| Valid       | No — The route is not valid.<br>Yes — The route is valid.                                                                                     |

The following output is an example of database information.

### Sample Output

```
A:ALA-A# show rip database
=====
RIP Route Database
=====
Destination Peer NextHop Metric Tag TTL Valid

180.0.0.10/32 180.1.7.15 0.0.0.0 2 0 163 No
180.0.0.10/32 180.1.8.14 0.0.0.0 2 0 179 No
180.0.0.14/32 180.1.8.14 0.0.0.0 1 0 179 Yes
180.0.6.0/24 180.1.7.15 0.0.0.0 11 8194 163 No
180.0.6.0/24 180.1.8.14 0.0.0.0 11 8194 179 No
180.0.7.0/24 180.1.7.15 0.0.0.0 11 8194 163 No
180.1.5.0/24 180.1.7.15 0.0.0.0 2 0 151 Yes
180.1.5.0/24 180.1.8.14 0.0.0.0 1 0 167 No
180.100.17.16/30 180.1.7.15 0.0.0.0 2 0 151 No
180.100.17.16/30 180.1.8.14 0.0.0.0 2 0 167 No

No. of Routes: 10
=====
A:ALA-A#
```

---

## group

- Syntax** `group [group-name] [detail]`
- Context** `show>router>rip`
- Description** This command displays RIP group information.
- Parameters** *group-name* — Displays RIP group information for the specified group.  
*detail* — Displays detailed RIP group information.
- Output** [Table 5](#) describes the standard command output fields for a RIP group.

**Table 5**      **RIP Group Fields**

| Label     | Description                                                                                                                                                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group     | The RIP group name.                                                                                                                                                                                                                                                                                                                   |
| Adm       | Down — The RIP group is administratively down.<br>Up — The RIP group is administratively up.                                                                                                                                                                                                                                          |
| Opr       | Down — The RIP group is operationally down.<br>Up — The RIP group is operationally up.                                                                                                                                                                                                                                                |
| Send Mode | Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address.<br>Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address.<br>None — Specifies that no RIP messages are sent (silent listener)<br>RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address.asp |
| Recv Mode | Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted.<br>None — Specifies that RIP updates will not be accepted.<br>RIPv1 — Specifies that RIP updates in version 1 format only will be accepted.<br>RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.              |
| Metric In | The metric value added to routes received from a RIP neighbor.                                                                                                                                                                                                                                                                        |

The following output is an example of RIP group information.

### Sample Standard RIP Group Output

```
A:ALA-A# show router rip group
```

```
=====
RIP Groups
=====
Group Adm Opr Send Recv Metric
 Mode Mode Mode Mode In

rip-group Up Down BCast Both 1
=====
A:ALA-A#
```

### Sample Detailed Output

```
A:ALA-A# show router rip group detail
=====
RIP groups (Detail)
=====

Group "rip-group"

Description : No Description Available
Admin State : Up
Oper State : Down
Send Mode : Broadcast
Receive Mode : Both
Metric In : 1
Metric Out : 1
Split Horizon : Enabled
Check Zero : Disabled
Message Size : 25
Preference : 100
Auth. Type : None
Update Timer : 30
Timeout Timer : 180
Flush Timer : 120
Export Policies:
 None
Import Policies:
 None
=====
A:ALA-A#
```

## neighbor

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>neighbor</b> [ <i>ip-addr</i>   <i>ip-int-name</i> ] [ <b>advertised-routes</b>   <b>detail</b> ]                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | show>router>rip                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command displays RIP neighbor interface information.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>ip-addr</i>   <i>ip-int-name</i> — Displays information for the specified IP interface.<br><br><b>Default</b> all neighbor interfaces<br><br><b>advertised-routes</b> — Displays the routes advertised to RIP neighbors. If no neighbors are specified, then all routes advertised to all neighbors are displayed. If a specific neighbor is given then only routes advertised to the given neighbor/interface are displayed.<br><br><b>Default</b> display RIP information |
| <b>Output</b>      | <a href="#">Table 6</a> describes the standard command output fields for a RIP group.                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 6**      **RIP Neighbor Standard Output Fields**

| Label      | Description                                                                                                                                                                                                                                                                                                                                   |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Neighbor   | The RIP neighbor interface name.                                                                                                                                                                                                                                                                                                              |
| Adm        | Down — The RIP neighbor interface is administratively down.<br>Up — The RIP neighbor interface is administratively up.                                                                                                                                                                                                                        |
| Opr        | Down — The RIP neighbor interface is operationally down.<br>Up — The RIP neighbor interface is operationally up.                                                                                                                                                                                                                              |
| Primary IP | The Primary IP address of the RIP neighbor interface.                                                                                                                                                                                                                                                                                         |
| Send Mode  | Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address.<br>Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address.<br>None — Specifies that no RIP messages are sent (such as a silent listener).<br>RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address. |
| Recv Mode  | Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted.<br>None — Specifies that RIP updates will not be accepted.<br>RIPv1 — Specifies that RIP updates in version 1 format only are accepted.<br>RIPv2v — Specifies that RIP updates in version 2 format only are accepted.                             |
| Metric In  | The metric added to routes received from a RIP neighbor.                                                                                                                                                                                                                                                                                      |

The following output is an example of RIP neighbor information.

**Sample Output**

```

A:ALA-A# show router rip neighbor
=====
RIP Neighbors
=====
Interface Adm Opr Primary IP Send Recv Metric
 Mode Mode In

router-2/1 Up Up 10.0.3.12 None Both 1
router-2/2 Up Up 10.0.5.12 BCast Both 1
router-2/3 Up Up 10.0.6.12 BCast Both 1
router-2/5 Up Up 10.0.9.12 BCast Both 1
router-2/6 Up Up 10.0.17.12 None Both 1
router-2/7 Up Up 10.0.16.12 None Both 1

```



=====

A:ALA-A#

[Table 7](#) describes the standard command output fields for a RIP group.

**Table 7 Detailed RIP Neighbor Output Fields**

| Label         | Description                                                                                                                                                                                                                                                                                                                                   |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Neighbor      | The RIP neighbor name.                                                                                                                                                                                                                                                                                                                        |
| Description   | The RIP neighbor description. No Description Available indicates no description is configured.                                                                                                                                                                                                                                                |
| Primary IP    | The RIP neighbor interface primary IP address.                                                                                                                                                                                                                                                                                                |
| Group         | The RIP group name of the neighbor interface.                                                                                                                                                                                                                                                                                                 |
| Admin State   | Down — The RIP neighbor interface is administratively down.<br>Up — The RIP neighbor interface is administratively up.                                                                                                                                                                                                                        |
| Oper State    | Down — The RIP neighbor interface is operationally down.<br>Up — The RIP neighbor interface is operationally up.                                                                                                                                                                                                                              |
| Send Mode     | Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address.<br>Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address.<br>None — Specifies that no RIP messages are sent (such as a silent listener).<br>RIPv1 — Specifies that RIPv1 formatted messages are sent to the broadcast address. |
| Recv Mode     | Both — Specifies that RIP updates in either version 1 or version 2 format will be accepted.<br>None — Specifies that RIP updates will not be accepted.<br>RIPv1 — Specifies that RIP updates in version 1 format only will be accepted.<br>RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.                      |
| Metric In     | The metric value added to routes received from a RIP neighbor.                                                                                                                                                                                                                                                                                |
| Metric Out    | The value added to routes exported into RIP and advertised to RIP neighbors.                                                                                                                                                                                                                                                                  |
| Split Horizon | Disabled — Split horizon disabled for the neighbor.<br>Enabled — Split horizon and poison reverse enabled for the neighbor.                                                                                                                                                                                                                   |

**Table 7 Detailed RIP Neighbor Output Fields (Continued)**

| Label           | Description                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Check Zero      | Disabled — Checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications are not checked allowing receipt of RIP messages even if mandatory zero fields are non-zero for the neighbor.<br>Enabled — checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages is enabled for the neighbor. |
| Message Size    | The maximum number of routes per RIP update message.                                                                                                                                                                                                                                                                                                                |
| Preference      | The preference of RIP routes from the neighbor.                                                                                                                                                                                                                                                                                                                     |
| Auth. Type      | Specifies the authentication type.                                                                                                                                                                                                                                                                                                                                  |
| Update Timer    | The current setting of the RIP update timer value expressed in seconds.                                                                                                                                                                                                                                                                                             |
| Timeout Timer   | The current RIP timeout timer value expressed in seconds.                                                                                                                                                                                                                                                                                                           |
| Export Policies | The export route policy that is used to determine routes advertised to all peers.                                                                                                                                                                                                                                                                                   |
| Import Policies | The import route policy that is used to determine which routes are accepted from RIP neighbors.                                                                                                                                                                                                                                                                     |

The following output is an example of detailed RIP neighbor information.

**Sample Detailed Output**

```
A:ALA-A# show router rip neighbor detail
=====
RIP Neighbors (Detail)
=====
Neighbor "router-2/7"

Description : No Description Available
Primary IP : 10.0.16.12 Group : seven
Admin State : Up Oper State : Up
Send Mode : None Receive Mode : Both
Metric In : 1 Metric Out : 1
Split Horizon : Enabled Check Zero : Disabled
Message Size : 25 Preference : 100
Auth. Type : None Update Timer : 3
Timeout Timer : 6 Flush Timer : 6
BFD: Enabled|Disabled
Export Policies:
 Rip2Rip
 direct2Rip
Import Policies:
```

```
None
=====
A:ALA-A#

Sample Output

A:ALA-A# show router rip neighbors interface advertised-routes
=====
RIP Advertised Routes
=====
Destination Interface NextHop Metric Tag TTL

180.0.0.2/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.5/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.8/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.9/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.10/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.11/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.12/32 180.1.8.12 0.0.0.0 1 0 n/a
180.0.0.13/32 180.1.8.12 0.0.0.0 10 8194 n/a
180.0.0.14/32 180.1.8.12 0.0.0.0 16 0 n/a
180.0.0.15/32 180.1.8.12 0.0.0.0 2 0 n/a
180.0.0.16/32 180.1.8.12 0.0.0.0 3 0 n/a

No. of Advertised Routes: 11
=====
A:ALA-A#
```

## peer

|                    |                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> [ <i>ip-int-name</i> ]                                                                                                         |
| <b>Context</b>     | show>router>rip                                                                                                                            |
| <b>Description</b> | This command displays RIP peer information.                                                                                                |
| <b>Parameters</b>  | <i>ip-int-name</i> — Displays peer information for peers on the specified IP interface.<br><b>Default</b> display peers for all interfaces |
| <b>Output</b>      | <a href="#">Table 8</a> describes the command output fields for a RIP peer.                                                                |

**Table 8** RIP Peer Output Fields

| Label          | Description                               |
|----------------|-------------------------------------------|
| Peer IP Addr   | The IP address of the peer router.        |
| Interface Name | The peer interface name.                  |
| Version        | The version of RIP running on the peer.   |
| Last Update    | The number of days since the last update. |

**Table 8** RIP Peer Output Fields (Continued)

| Label        | Description              |
|--------------|--------------------------|
| No. of Peers | The number of RIP peers. |

The following output is an example of RIP peer information.

**Sample Output**

```
A:ALA-A# show router rip peers
=====
RIP Peers
=====
Peer IP Addr Interface Name Version Last Update

10.0.5.13 router-2/2 RIPv2 0
10.0.6.16 router-2/3 RIPv2 2
10.0.9.14 router-2/5 RIPv2 8
10.0.10.15 router-2/4 RIPv2 0

No. of Peers: 4
=====
A:ALA-A#
```

**statistics**

- Syntax** **statistics** [*ip-addr* | *ip-int-name*]
- Context** show>router>rip
- Description** This command displays interface level statistics for RIP.
- If no IP address or interface name is specified, then all configured RIP interfaces are displayed.
- If an IP address or interface name is specified, then only data regarding the specified RIP interface is displayed.
- Parameters** *ip-addr* | *ip-int-name* — Displays statistics for the specified RIP IP interface.
- Output** [Table 9](#) describes the output fields for RIP statistics.

**Table 9** RIP Statistics Output Fields

| Label            | Description                                                               |
|------------------|---------------------------------------------------------------------------|
| Learned Routes   | The number of RIP-learned routes were exported to RIP or RIPv2 neighbors. |
| Timed Out Routes | The number of routes that have been timed out.                            |

**Table 9**      **RIP Statistics Output Fields (Continued)**

| Label                  | Description                                                                                                                                                                                                                                                                                                                                      |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Current Memory         | The amount of memory used by this RIP router instance.                                                                                                                                                                                                                                                                                           |
| Maximum Memory         | The amount of memory allocated for this RIP router instance.                                                                                                                                                                                                                                                                                     |
| Interface              | Displays the name of each interface configured in RIP and associated RIP statistics.                                                                                                                                                                                                                                                             |
| Primary IP             | The interface IP address.                                                                                                                                                                                                                                                                                                                        |
| Update Timer           | The current setting of the RIP update timer value expressed in seconds.                                                                                                                                                                                                                                                                          |
| Timeout Timer          | The current RIP timeout timer value expressed in seconds.                                                                                                                                                                                                                                                                                        |
| Flush Timer            | The number of seconds after a route has been declared invalid that it is flushed from the route database.                                                                                                                                                                                                                                        |
| Updates Sent           | Total — The total number of updates that were sent.<br>Last 5 Min — The number of updates that were sent in the last 5 minutes.<br>Last 1 Min — The number of updates that were sent in the last 1 minute.                                                                                                                                       |
| Triggered Updates      | Total — The total number of triggered updates sent. These updates are sent before the entire routing table is sent.<br>Last 5 Min — The number of triggered updates that were sent in the last 5 minutes.<br>Last 1 Min — The number of triggered updates that were sent in the last 1 minute.                                                   |
| Bad Packets Received   | Total — The total number of RIP updates received on this interface that were discarded as invalid.<br>Last 5 Min — The number of updates received on this interface that were discarded as invalid the last 5 minutes.<br>Last 1 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 1 minute. |
| RIPv1 Updates Received | Total — The total number of updates received.<br>Last 5 Min — The number of updates received in the last 5 minutes.<br>Last 1 Min — The number of updates received in the last 1 minute.                                                                                                                                                         |

**Table 9**      **RIP Statistics Output Fields (Continued)**

| Label                   | Description                                                                                                                                                                                                                                                                                                               |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RIPv1 Updates Ignored   | Total — The total number of updates ignored.<br>Last 5 Min — The number of updates ignored in the last 5 minutes.<br>Last 1 Min — The number of updates ignored in the last 1 minute.                                                                                                                                     |
| RIPv1 Bad Routes        | Total — The total number of bad routes received from the peer.<br>Last 5 Min — The number of bad routes received from the peer in the last 5 minutes.<br>Last 1 Min — The number of bad routes received from the peer in the last minute.                                                                                 |
| RIPv1 Requests Received | Total — The total number of times the router received route requests from other routers.<br>Last 5 Min — The number of times the router received route requests from other routers in the last 5 minutes.<br>Last 1 Min — The number of times the router received route requests from other routers in the last 1 minute. |
| RIPv1 Requests Ignored  | Total — The total number of times the router ignored route requests from other routers.<br>Last 5 Min — The number of times the router ignored route requests from other routers in the last 5 minutes.<br>Last 1 Min — The number of times the router ignored route requests from other routers in the last 1 minute.    |
| RIPv2 Updates Received  | Total — The total number of updates received.<br>Last 5 Min — The number of updates received in the last 5 minutes.<br>Last 1 Min — The number of updates received in the last minute.                                                                                                                                    |
| RIPv2 Updates Ignored   | Total — The total number of updates ignored.<br>Last 5 Min — The number of updates ignored in the last 5 minutes.<br>Last 1 Min — The number of updates ignored in the last minute.                                                                                                                                       |
| RIPv2 Bad Routes        | Total — The total number of bad routes received from the peer.<br>Last 5 Min — The number of bad routes received from the peer in the last 5 minutes.<br>Last 1 Min — The number of bad routes received from the peer in the last minute.                                                                                 |

**Table 9 RIP Statistics Output Fields (Continued)**

| Label                   | Description                                                                                                                                                                                                                                                                                                             |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RIPv2 Requests Received | Total — The total number of times the router received route requests from other routers.<br>Last 5 Min — The number of times the router received route requests from other routers in the last 5 minutes.<br>Last 1 Min — The number of times the router received route requests from other routers in the last minute. |
| RIPv2 Requests Ignored  | Total — The total number of times the router ignored route requests from other routers.<br>Last 5 Min — The number of times the router ignored route requests from other routers in the last 5 minutes.<br>Last 1 Min — The number of times the router ignored route requests from other routers in the last minute.    |
| Authentication Errors   | Total — The total number of authentication errors to secure table updates.<br>Last 5 Min — The number of authentication errors to secure table updates in the last 5 minutes.<br>Last 1 Min — The number of authentication errors to secure table updates in the last minute.                                           |

The following output is an example of RIP statistics information.

### Sample Output

```
A:ALA-A# show router rip statistics
=====
RIP Statistics
=====
Learned Routes : 0 Timed Out Routes : 0
Current Memory : 120624 Maximum Memory : 262144

Interface "to-web"

Primary IP : 10.1.1.3 Update Timer : 30
Timeout Timer : 180 Flush Timer : 120

Counter Total Last 5 Min Last 1 Min

Updates Sent 0 0 0
Triggered Updates 0 0 0
Bad Packets Received 0 0 0
RIPv1 Updates Received 0 0 0
RIPv1 Updates Ignored 0 0 0
RIPv1 Bad Routes 0 0 0
RIPv1 Requests Received 0 0 0
```

```
RIPv1 Requests Ignored 0 0 0
RIPv2 Updates Received 0 0 0
RIPv2 Updates Ignored 0 0 0
RIPv2 Bad Routes 0 0 0
RIPv2 Requests Received 0 0 0
RIPv2 Requests Ignored 0 0 0
Authentication Errors 0 0 0
=====
A:ALA-A#
```

database

**Syntax**     **database** [**peer** *ipv6-address*] [**detail**]  
              **database** [**peer** *ipv6-address*] **detail** [**qos**]  
              **database** *ipv6-address*[/*prefix-length*] [**longer**] [**peer** *ipv6-address*] [**detail**]  
              **database** *ipv6-address*[/*prefix-length*] [**longer**] [**peer** *ipv6-address*]    **detail** [**qos**]

**Context**     show>router>ripng

**Description**   This command displays the routes in the RIPng database.

**Parameters**   *ipv6-address* — Specifies the RIPng IPv6 address.

**Values**  
                  *ipv6-address*:    x:x:x:x:x:x:x [-interface]  
                                      x:x:x:x:x:d.d.d.d [-interface]  
                                      x: [0..FFFF]H  
                                      d: [0..255]D  
                                      interface — link local addresses up to 32  
                                      characters

*prefix-length* — The IP prefix in dotted decimal notation for the range used by the ABR  
                                      to advertise that summarizes the area into another area.

**Values**    *ipv6-prefix*:  
                              • x:x:x:x:x:x:x (eight 16-bit pieces)  
                              • x:x:x:x:x:d.d.d.d  
                              • x: [0 to FFFF]H  
                              • d: [0 to 255]D  
                              *prefix-length*: 0 to 128

**detail** — Displays detailed information.

**qos** — Displays QoS information.

**longer** — Displays more detailed information.

**Output**       The following are examples of router RIPng database information.



### Sample Output

```
*A:Dut-C>config>router>if# show router ripng database
=====
RIP-NG Route Database
=====
Destination Interface
Peer Metric TTL Valid

3ffe::10:10:3:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:3:0/120 two
 fe80::6629:ffff:fe00:0 1 3 Yes
3ffe::10:10:3:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:3:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 two
 fe80::6629:ffff:fe00:0 1 3 Yes
3ffe::10:10:12:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 two
 fe80::6629:ffff:fe00:0 1 3 Yes
3ffe::10:10:103:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 two
 fe80::6629:ffff:fe00:0 1 3 Yes
3ffe::10:10:112:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 two
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 four
 fe80::6629:ffff:fe00:0 1 3 Yes

No. of Routes: 20
=====
show router ripng database [peer <ipv6-address>]
*A:Dut-C>config>router>if# show router ripng database peer fe80::6629:ffff:fe00:0
=====
RIP-NG Route Database
=====
Destination Interface
Peer Metric TTL Valid

```

```

3ffe::10:10:3:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:3:0/120 two
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:3:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:3:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 two
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:12:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 two
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:103:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 two
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:10:112:0/120 four
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 one
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 two
 fe80::6629:ffff:fe00:0 1 3 Yes
3ffe::10:20:1:2/128 three
 fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 four
 fe80::6629:ffff:fe00:0 1 4 Yes

```

No. of Routes: 20

```
=====
show router ripng database [peer <ipv6-address>] detail
```

```
*A:DutC>config>router>if# show router ripng database peer fe80::6629:ffff:fe00:0 detail
```

```
=====
RIP-NG Route Database (Detail)
```

```
=====
Destination : 3ffe::10:10:3:0/120 Metric : 1
Interface : one TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:10:3:0/120 Metric : 1
Interface : two TTL : 3
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

```

---

|             |                          |        |       |
|-------------|--------------------------|--------|-------|
| Destination | : 3ffe::10:10:3:0/120    | Metric | : 1   |
| Interface   | : three                  | TTL    | : 3   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:3:0/120    | Metric | : 1   |
| Interface   | : four                   | TTL    | : 3   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:12:0/120   | Metric | : 1   |
| Interface   | : one                    | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:12:0/120   | Metric | : 1   |
| Interface   | : two                    | TTL    | : 3   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:12:0/120   | Metric | : 1   |
| Interface   | : three                  | TTL    | : 3   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:12:0/120   | Metric | : 1   |
| Interface   | : four                   | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:103:0/120  | Metric | : 1   |
| Interface   | : one                    | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:103:0/120  | Metric | : 1   |
| Interface   | : two                    | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:103:0/120  | Metric | : 1   |
| Interface   | : three                  | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:103:0/120  | Metric | : 1   |
| Interface   | : four                   | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:112:0/120  | Metric | : 1   |
| Interface   | : one                    | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |
| Peer        | : fe80::6629:ffff:fe00:0 | Valid  | : Yes |
|             |                          |        |       |
| Destination | : 3ffe::10:10:112:0/120  | Metric | : 1   |
| Interface   | : two                    | TTL    | : 4   |
| Next Hop    | : ::                     | Tag    | : 0   |

```
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:10:112:0/120 Metric : 1
Interface : three TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:10:112:0/120 Metric : 1
Interface : four TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:20:1:2/128 Metric : 1
Interface : one TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:20:1:2/128 Metric : 1
Interface : two TTL : 3
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:20:1:2/128 Metric : 1
Interface : three TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes

Destination : 3ffe::10:20:1:2/128 Metric : 1
Interface : four TTL : 4
Next Hop : :: Tag : 0
Peer : fe80::6629:ffff:fe00:0 Valid : Yes
```

```
=====
show router ripng database <ipv6-address[/prefix-length]>
*A:Dut-C>config>router>if# show router ripng database 3ffe::10:20:1:2/128
=====
RIP-NG Route Database
=====
Destination Interface
Peer Metric TTL Valid

3ffe::10:20:1:2/128 one
fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 two
fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 three
fe80::6629:ffff:fe00:0 1 4 Yes
3ffe::10:20:1:2/128 four
fe80::6629:ffff:fe00:0 1 3 Yes

No. of Routes: 4
=====
```

group

**Syntax**    **group** [*group-name*] [*detail*]

|                    |                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>router>ripng                                                                                                                           |
| <b>Description</b> | This command displays RIPng group information.                                                                                              |
| <b>Parameters</b>  | <i>group-name</i> — Displays RIPng group information for the specified group.<br><b>detail</b> — Displays detailed RIPng group information. |
| <b>Output</b>      | The following output is an example of RIPng group information.                                                                              |

### Sample Output

```
*A:Dut-C>config>router>if# show router ripng group "A"
=====
RIP-NG Groups
=====
Group Adm Opr Send Recv Metric
 Mode Mode Mode Mode In

A Up Up RipNg RipNg 1

Groups : 1
=====
*A:Dut-C>config>router>if# show router ripng group "A" detail
=====
RIP-NG groups (Detail)
=====
Group "A"

Description : No Description Available
Admin State : Up Oper State : Up
Send Mode : RipNg Receive Mode : RipNg
Metric In : 1 Metric Out : 1
Split Horizon : Enabled Check Zero : Disabled
Message Size : 25 Preference : 100
 Update Timer : 1
Timeout Timer : 5 Flush Timer : 5
BFD : Enabled
Export Policies:
 directToRipng
Import Policies:
 None
=====
```

## neighbor

|                    |                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>neighbor</b> [ <i>ip-int-name</i>   <i>ipv6-address</i> ] [ <b>detail</b> ]<br><b>neighbor</b> [ <b>detail</b> ] <b>advertised-routes</b> [ <i>ip-int-name</i>   <i>ipv6-link-local-address</i> ] |
| <b>Context</b>     | show>router>ripng                                                                                                                                                                                    |
| <b>Description</b> | This command displays RIPng neighbor interface information.                                                                                                                                          |

- Parameters** *ip-int-name* | *ipv6-address* — Displays information for the specified IPv6 interface and address.
- Default** all neighbor interfaces
- advertised-routes** — Displays the routes advertised to RIPng neighbors. If no neighbors are specified, then all routes advertised to all neighbors are displayed. If a specific neighbor is given then only routes advertised to the given neighbor/interface are displayed.
- Default** displays RIPng information
- ipv6-link-local-address* — Displays information about the of the IPv6 link local address for the neighbor.
- detail** — Displays detailed information.
- Output** The following output is an example of RIPng neighbor information.

### Sample Output

```
*A:Dut-C>config>router>if# show router ripng neighbor "one"
=====
RIP-NG Neighbors
=====
Interface Adm Opr Primary IP Send Recv Metric
 Mode Mode In

one Up Up 3ffe::10:10:3:* RipNg RipNg 1

No. of Neighbors: 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C>config>router>if# show router ripng neighbor "one" detail
=====
RIP-NG Neighbors (Detail)
=====
Neighbor "one"

Description : n/a
Primary IP : 3ffe::10:10:3:3/1* Group : A
Admin State : Up Oper State : Up
Send Mode : RipNg Receive Mode : RipNg
Metric In : 1 Metric Out : 1
Split Horizon : Enabled Check Zero : Disabled
Message Size : 25 Preference : 100
 Update Timer : 1
 Flush Timer : 5
Timeout Timer : 5
BFD : Enabled
Export Policies:
 directToRipng
Import Policies:
 None
=====
* indicates that the corresponding row element may have been truncated.
```

```
*A:Dut-C>config>router>if# show router ripng neighbor 3ffe::10:10:12:3
=====
RIP-NG Neighbors
=====
Interface Adm Opr Primary IP Send Recv Metric
 Mode Mode In

three Up Up 3ffe::10:10:12* RipNg RipNg 1

No. of Neighbors: 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C>config>router>if# show router ripng neighbor 3ffe::10:10:12:3 detail
=====
RIP-NG Neighbors (Detail)
=====
Neighbor "three"

Description : n/a
Primary IP : 3ffe::10:10:12:3/* Group : B
Admin State : Up Oper State : Up
Send Mode : RipNg Receive Mode : RipNg
Metric In : 1 Metric Out : 1
Split Horizon : Enabled Check Zero : Disabled
Message Size : 25 Preference : 100
 Update Timer : 1
 Flush Timer : 5
Timeout Timer : 5
BFD : Disabled
Export Policies:
 directToRipng
Import Policies:
 None
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C>config>router>if# show router ripng neighbor advertised-routes
=====
RIP-NG Advertised Routes
=====
Destination Interface
Link-local-address Metric TTL

3ffe::103:800/120 one
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::103:800/120 two
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::103:800/120 three
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::103:800/120 four
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:10:3:0/120 one
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:10:3:0/120 two
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:10:3:0/120 three
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:10:3:0/120 four
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:10:12:0/120 one
```

|                        |       |     |
|------------------------|-------|-----|
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:12:0/120   | two   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:12:0/120   | three |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:12:0/120   | four  |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:103:0/120  | one   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:103:0/120  | two   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:103:0/120  | three |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:103:0/120  | four  |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:112:0/120  | one   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:112:0/120  | two   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:112:0/120  | three |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:10:112:0/120  | four  |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:20:1:2/128    | one   |     |
| fe80::662d:ffff:fe00:0 | 16    | n/a |
| 3ffe::10:20:1:2/128    | two   |     |
| fe80::662d:ffff:fe00:0 | 16    | n/a |
| 3ffe::10:20:1:2/128    | three |     |
| fe80::662d:ffff:fe00:0 | 16    | n/a |
| 3ffe::10:20:1:2/128    | four  |     |
| fe80::662d:ffff:fe00:0 | 16    | n/a |
| 3ffe::10:20:1:3/128    | one   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:20:1:3/128    | two   |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:20:1:3/128    | three |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |
| 3ffe::10:20:1:3/128    | four  |     |
| fe80::662d:ffff:fe00:0 | 1     | n/a |

-----  
No. of Advertised Routes: 28  
=====

\*A:Dut-C>config>router>if# show router ripng neighbor advertised-  
routes "one"

=====

| Destination            | Interface |     |
|------------------------|-----------|-----|
| Link-local-address     | Metric    | TTL |
| -----                  |           |     |
| 3ffe::103:800/120      | one       |     |
| fe80::662d:ffff:fe00:0 | 1         | n/a |
| 3ffe::10:10:3:0/120    | one       |     |
| fe80::662d:ffff:fe00:0 | 1         | n/a |
| 3ffe::10:10:12:0/120   | one       |     |
| fe80::662d:ffff:fe00:0 | 1         | n/a |
| 3ffe::10:10:103:0/120  | one       |     |
| fe80::662d:ffff:fe00:0 | 1         | n/a |
| 3ffe::10:10:112:0/120  | one       |     |



```

fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:20:1:2/128 one
fe80::662d:ffff:fe00:0 16 n/a
3ffe::10:20:1:3/128 one
fe80::662d:ffff:fe00:0 1 n/a

```

-----  
No. of Advertised Routes: 7

=====

```

*A:Dut-C>config>router>if# show router ripng neighbor advertised-
routes fe80::662d:ffff:fe00:0

```

=====

RIP-NG Advertised Routes

=====

| Destination<br>Link-local-address | Interface<br>Metric TTL |
|-----------------------------------|-------------------------|
| -----                             | -----                   |
| 3ffe::103:800/120                 | one                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::103:800/120                 | two                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::103:800/120                 | three                   |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::103:800/120                 | four                    |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:3:0/120               | one                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:3:0/120               | two                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:3:0/120               | three                   |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:3:0/120               | four                    |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:12:0/120              | one                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:12:0/120              | two                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:12:0/120              | three                   |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:12:0/120              | four                    |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:103:0/120             | one                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:103:0/120             | two                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:103:0/120             | three                   |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:103:0/120             | four                    |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:112:0/120             | one                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:112:0/120             | two                     |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:112:0/120             | three                   |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:10:112:0/120             | four                    |
| fe80::662d:ffff:fe00:0            | 1 n/a                   |
| 3ffe::10:20:1:2/128               | one                     |
| fe80::662d:ffff:fe00:0            | 16 n/a                  |
| 3ffe::10:20:1:2/128               | two                     |

```

 fe80::662d:ffff:fe00:0 16 n/a
3ffe::10:20:1:2/128 three
 fe80::662d:ffff:fe00:0 16 n/a
3ffe::10:20:1:2/128 four
 fe80::662d:ffff:fe00:0 16 n/a
3ffe::10:20:1:3/128 one
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:20:1:3/128 two
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:20:1:3/128 three
 fe80::662d:ffff:fe00:0 1 n/a
3ffe::10:20:1:3/128 four
 fe80::662d:ffff:fe00:0 1 n/a

No. of Advertised Routes: 28
=====
```

peer

- Syntax

peer [ip-int-name]
- Context

show>router>ripng
- Description

This command displays RIPng peer information.
- Parameters

ip-int-name — Displays peer information for peers on the specified IPv6 interface.

Default display peers for all interfaces
- Output

The following is an example of RIPng peer information.

Sample Output

```
*A:Dut-C>config>router>if# show router ripng peer
=====
RIP-NG Peers
=====
Peer IP Addr Version Last Update
Interface Name

fe80::6629:ffff:fe00:0 ripNg-v1 0
one
fe80::6629:ffff:fe00:0 ripNg-v1 0
two
fe80::6629:ffff:fe00:0 ripNg-v1 0
three
fe80::6629:ffff:fe00:0 ripNg-v1 0
four

No. of Peers: 4
=====
*A:Dut-C>config>router>if# show router ripng peer "one"
=====
RIP-NG Peers
=====
```

| Peer IP Addr           | Version  | Last Update |
|------------------------|----------|-------------|
| Interface Name         |          |             |
| -----                  |          |             |
| fe80::6629:ffff:fe00:0 | ripNg-v1 | 0           |
| one                    |          |             |
| -----                  |          |             |
| No. of Peers: 1        |          |             |
| =====                  |          |             |

statistics

|                    |                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b> [ <i>ipv6-address</i>   <i>ip-int-name</i> ]                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | show>router>ripng                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command displays the interface level statistics for the RIPng protocol</p> <p>If no IPv6 address or interface name is specified, then all configured RIPng interfaces are displayed.</p> <p>If an IPv6 address or interface name is specified, then only data regarding the specified RIPng interface is displayed.</p> |
| <b>Parameters</b>  | <i>ipv6-address</i>   <i>ip-int-name</i> — Displays statistics for the specified IPv6 interface.                                                                                                                                                                                                                                |
| <b>Output</b>      | The following output are examples of RIPng statistics information.                                                                                                                                                                                                                                                              |

Sample Output

|                                                         |                      |                  |             |
|---------------------------------------------------------|----------------------|------------------|-------------|
| *A:Dut-C>config>router>if# show router ripng statistics |                      |                  |             |
| =====                                                   |                      |                  |             |
| RIP-NG Statistics                                       |                      |                  |             |
| =====                                                   |                      |                  |             |
| Learned Routes                                          | : 40                 | Timed Out Routes | : 0         |
| Current Memory                                          | : 1,199,016          | Maximum Memory   | : 2,097,152 |
| -----                                                   |                      |                  |             |
| Interface "one"                                         |                      |                  |             |
| -----                                                   |                      |                  |             |
| Primary IP                                              | : 3ffe::10:10:3:3/1* | Update Timer     | : 1         |
| Timeout Timer                                           | : 5                  | Flush Timer      | : 5         |
| -----                                                   |                      |                  |             |
| Counter                                                 | Total                | Last 5 Min       | Last 1 Min  |
| -----                                                   |                      |                  |             |
| Updates Sent                                            | 221                  | 180              | 5           |
| Triggered Updates                                       | 0                    | 0                | 0           |
| Bad Packets Received                                    | 0                    | 0                | 0           |
| Updates Received                                        | 225                  | 185              | 3           |
| Updates Ignored                                         | 0                    | 0                | 0           |
| Bad Routes                                              | 0                    | 0                | 0           |
| Requests Received                                       | 0                    | 0                | 0           |
| Requests Ignored                                        | 0                    | 0                | 0           |
| -----                                                   |                      |                  |             |
| Interface "two"                                         |                      |                  |             |
| -----                                                   |                      |                  |             |

```

Primary IP : 3ffe::10:10:103:3* Update Timer : 1
Timeout Timer : 5 Flush Timer : 5

```

| Counter              | Total | Last 5 Min | Last 1 Min |
|----------------------|-------|------------|------------|
| Updates Sent         | 220   | 183        | 5          |
| Triggered Updates    | 0     | 0          | 0          |
| Bad Packets Received | 0     | 0          | 0          |
| Updates Received     | 229   | 183        | 5          |
| Updates Ignored      | 0     | 0          | 0          |
| Bad Routes           | 0     | 0          | 0          |
| Requests Received    | 0     | 0          | 0          |
| Requests Ignored     | 0     | 0          | 0          |

```
Interface "three"
```

```

Primary IP : 3ffe::10:10:12:3/* Update Timer : 1
Timeout Timer : 5 Flush Timer : 5

```

| Counter              | Total | Last 5 Min | Last 1 Min |
|----------------------|-------|------------|------------|
| Updates Sent         | 224   | 184        | 6          |
| Triggered Updates    | 0     | 0          | 0          |
| Bad Packets Received | 0     | 0          | 0          |
| Updates Received     | 221   | 182        | 4          |
| Updates Ignored      | 0     | 0          | 0          |
| Bad Routes           | 0     | 0          | 0          |
| Requests Received    | 0     | 0          | 0          |
| Requests Ignored     | 0     | 0          | 0          |

```
Interface "four"
```

```

Primary IP : 3ffe::10:10:112:3* Update Timer : 1
Timeout Timer : 5 Flush Timer : 5

```

| Counter              | Total | Last 5 Min | Last 1 Min |
|----------------------|-------|------------|------------|
| Updates Sent         | 220   | 182        | 8          |
| Triggered Updates    | 0     | 0          | 0          |
| Bad Packets Received | 0     | 0          | 0          |
| Updates Received     | 227   | 185        | 4          |
| Updates Ignored      | 0     | 0          | 0          |
| Bad Routes           | 0     | 0          | 0          |
| Requests Received    | 0     | 0          | 0          |
| Requests Ignored     | 0     | 0          | 0          |

```
* indicates that the corresponding row element may have been truncated.
```

```
*A:Dut-C>config>router>if# show router ripng statistics "one"
```

```
RIP-NG Statistics
```

```

Learned Routes : 40 Timed Out Routes : 0
Current Memory : 1,199,016 Maximum Memory : 2,097,152

```

```
Interface "one"
```

```

Primary IP : 3ffe::10:10:3:3/1* Update Timer : 1
Timeout Timer : 5 Flush Timer : 5

```

---

| Counter                                                                 | Total | Last 5 Min | Last 1 Min |
|-------------------------------------------------------------------------|-------|------------|------------|
| -----                                                                   | ----- | -----      | -----      |
| Updates Sent                                                            | 224   | 183        | 8          |
| Triggered Updates                                                       | 0     | 0          | 0          |
| Bad Packets Received                                                    | 0     | 0          | 0          |
| Updates Received                                                        | 230   | 190        | 8          |
| Updates Ignored                                                         | 0     | 0          | 0          |
| Bad Routes                                                              | 0     | 0          | 0          |
| Requests Received                                                       | 0     | 0          | 0          |
| Requests Ignored                                                        | 0     | 0          | 0          |
| =====                                                                   |       |            |            |
| * indicates that the corresponding row element may have been truncated. |       |            |            |

## 2.9.2.2 Clear Commands

### database

|                    |                                                      |
|--------------------|------------------------------------------------------|
| <b>Syntax</b>      | <b>database</b>                                      |
| <b>Context</b>     | clear>router>rip                                     |
| <b>Description</b> | This command flushes all routes in the RIP database. |

### export

|                    |                                                |
|--------------------|------------------------------------------------|
| <b>Syntax</b>      | <b>export</b>                                  |
| <b>Context</b>     | clear>router>rip                               |
| <b>Description</b> | This command re-evaluates RIP export policies. |

### statistics

|                    |                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b> [ <i>neighbor ip-int-name</i>   <i>ip-address</i> ]                           |
| <b>Context</b>     | clear>router>rip                                                                                |
| <b>Description</b> | This command clears statistics for RIP neighbors.                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Clears the statistics for the specified RIP interface. |
| <b>Default</b>     | clears statistics for all RIP interfaces                                                        |

---

## database

|                    |                                                        |
|--------------------|--------------------------------------------------------|
| <b>Syntax</b>      | <b>database</b>                                        |
| <b>Context</b>     | clear>router>ripng                                     |
| <b>Description</b> | This command flushes all routes in the RIPng database. |

## export

|                    |                                                  |
|--------------------|--------------------------------------------------|
| <b>Syntax</b>      | <b>export</b>                                    |
| <b>Context</b>     | clear>router>rip                                 |
| <b>Description</b> | This command re-evaluates RIPng export policies. |

## statistics

|                    |                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b> [ <b>neighbor</b> <i>ip-int-name</i>   <i>ip-address</i> ]                                                                                 |
| <b>Context</b>     | clear>router>rip                                                                                                                                             |
| <b>Description</b> | This command clears statistics for RIPng neighbors.                                                                                                          |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Clears the statistics for the specified RIPng interface.<br><b>Default</b> clears statistics for all RIP interfaces |

### 2.9.2.3 Debug RIP Commands

## auth

|                    |                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>auth</b> [ <b>neighbor</b> <i>ip-int-name</i>   <i>ip-address</i> ]                                      |
| <b>Context</b>     | debug>router>rip                                                                                                 |
| <b>Description</b> | This command enables debugging for RIP authentication.                                                           |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP authentication for the neighbor IP address or interface. |

---

## error

|                    |                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] error [neighbor <i>ip-int-name</i>   <i>ip-address</i>]</b>                                          |
| <b>Context</b>     | debug>router>rip                                                                                             |
| <b>Description</b> | This command enables debugging for RIP errors.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP errors sent on the neighbor IP address or interface. |

## events

|                    |                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] events [neighbor <i>ip-int-name</i>   <i>ip-addr</i>]</b>                                            |
| <b>Context</b>     | debug>router>rip                                                                                             |
| <b>Description</b> | This command enables debugging for RIP events.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP events sent on the neighbor IP address or interface. |

## holddown

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] holddown [neighbor <i>ip-int-name</i>   <i>ip-address</i>]</b>                                          |
| <b>Context</b>     | debug>router>rip                                                                                                |
| <b>Description</b> | This command enables debugging for RIP holddowns.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP holddowns sent on the neighbor IP address or interface. |

## packets

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] packets [neighbor <i>ip-int-name</i>   <i>ip-addr</i>]</b>                                            |
| <b>Context</b>     | debug>router>rip                                                                                              |
| <b>Description</b> | This command enables debugging for RIP packets.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP packets sent on the neighbor IP address or interface. |

---

## requests

|                    |                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] requests [neighbor <i>ip-int-name</i>   <i>ip-address</i>]</b>                                         |
| <b>Context</b>     | debug>router>rip                                                                                               |
| <b>Description</b> | This command enables debugging for RIP requests.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP requests sent on the neighbor IP address or interface. |

## trigger

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] trigger [neighbor <i>ip-int-name</i>   <i>ip-address</i>]</b>                                         |
| <b>Context</b>     | debug>router>rip                                                                                              |
| <b>Description</b> | This command enables debugging for RIP trigger updates.                                                       |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP updates sent on the neighbor IP address or interface. |

## updates

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] updates [neighbor <i>ip-int-name</i>   <i>ip-address</i>]</b>                                         |
| <b>Context</b>     | debug>router>rip                                                                                              |
| <b>Description</b> | This command enables debugging for RIP updates.                                                               |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ip-address</i> — Debugs the RIP updates sent on the neighbor IP address or interface. |

## error

|                    |                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] error [neighbor <i>ip-int-name</i>   <i>ipv6-address</i>]</b>                                            |
| <b>Context</b>     | debug>router>ripng                                                                                               |
| <b>Description</b> | This command enables debugging for RIPng errors.                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIPng errors sent on the neighbor IP address or interface. |



---

## events

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] events [neighbor <i>ip-int-name</i>]</b>                                |
| <b>Context</b>     | debug>router>ripng                                                              |
| <b>Description</b> | This command enables debugging for RIPng events.                                |
| <b>Parameters</b>  | <i>ip-int-name</i> — Debugs the RIPng events sent on the neighbor IP interface. |

## holddown

|                    |                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] holddown [neighbor <i>ip-int-name</i>   <i>ipv6-address</i>]</b>                                            |
| <b>Context</b>     | debug>router>ripng                                                                                                  |
| <b>Description</b> | This command enables debugging for RIPng holddowns.                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIPng holddowns sent on the neighbor IP address or interface. |

## packets

|                    |                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] packets [neighbor <i>ip-int-name</i>   <i>ipv6-address</i>]</b>                                           |
| <b>Context</b>     | debug>router>ripng                                                                                                |
| <b>Description</b> | This command enables debugging for RIPng packets.                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIPng packets sent on the neighbor IP address or interface. |

## requests

|                    |                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] requests [neighbor <i>ip-int-name</i>   <i>ipv6-address</i>]</b>                                         |
| <b>Context</b>     | debug>router>ripng                                                                                               |
| <b>Description</b> | This command enables debugging for RIP requests.                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIP requests sent on the neighbor IP address or interface. |

## trigger

|               |                                                                         |
|---------------|-------------------------------------------------------------------------|
| <b>Syntax</b> | <b>[no] trigger [neighbor <i>ip-int-name</i>   <i>ipv6-address</i>]</b> |
|---------------|-------------------------------------------------------------------------|

---

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | debug>router>ripng                                                                                              |
| <b>Description</b> | This command enables debugging for RIP trigger updates.                                                         |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIP updates sent on the neighbor IP address or interface. |

## updates

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] updates [neighbor <i>ip-int-name</i>   <i>ipv6-address</i> ]                                               |
| <b>Context</b>     | debug>router>ripng                                                                                              |
| <b>Description</b> | This command enables debugging for RIP updates.                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i>   <i>ipv6-address</i> — Debugs the RIP updates sent on the neighbor IP address or interface. |

---

## 3 OSPF

### 3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

When a router is started with OSPF configured, OSPF, along with the routing-protocol data structures, is initialized and waits for indications from lower-layer protocols that its interfaces are functional. Nokia's implementation of OSPF conforms to OSPF Version 2 specifications presented in RFC 2328, OSPF Version 2 and OSPF Version 3 specifications presented in RFC 2740, OSPF for IPv6. Routers running OSPF can be enabled with minimal configuration. All default and command parameters can be modified.

Changes between OSPF for IPv4 and OSPF3 for IPv6 include the following:

- Addressing semantics have been removed from OSPF packets and the basic link-state advertisements (LSAs). New LSAs have been created to carry IPv6 addresses and prefixes.
- OSPF3 runs on a per-link basis, instead of on a per-IP-subnet basis.
- Flooding scope for LSAs has been generalized.
- Unlike OSPFv2, OSPFv3 authentication relies on IPV6's authentication header and encapsulating security payload.
- Most packets in OSPF for IPv6 are almost as compact as those in OSPF for IPv4, even with the larger IPv6 addresses.
- Most field and packet-size limitations present in OSPF for IPv4 have been relaxed.
- Option handling has been made more flexible.

Key OSPF features are:

- Backbone areas
- Stub areas
- Not-So-Stubby areas (NSSAs)

- Virtual links
- Authentication
- Route redistribution
- Routing interface parameters
- OSPF-TE extensions (Nokia's implementation allows MPLS fast reroute)

### 3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.


Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

#### 3.1.1.1 Backbone Area

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see area 0.0.0.5 in [Figure 5](#)) then the ABRs (such as routers Y and Z) must be connected via a virtual link. The two ABRs form a point-to-point-like adjacency across the transit area (see area 0.0.0.4).

**Figure 5 Backbone Area**



### 3.1.1.2 Stub Area

A stub area is a designated area that does not allow external route advertisements. Routers in a stub area do not maintain external routes. A single default route to an ABR replaces all external routes. This OSPF implementation supports the optional summary route (type-3) advertisement suppression from other areas into a stub area. This feature further reduces topological database sizes and OSPF protocol traffic, memory usage, and CPU route calculation time.

In [Figure 5](#), areas 0.0.0.1, 0.0.0.2 and 0.0.0.5 could be configured as stub areas. A stub area cannot be designated as the transit area of a virtual link and a stub area cannot contain an AS boundary router. An AS boundary router exchanges routing information with routers in other ASs.

### 3.1.1.3 Not-So-Stubby Area

Another OSPF area type is called a Not-So-Stubby area (NSSA). NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. External routes learned by OSPF routers in the NSSA area are advertised as type-7 LSAs within the NSSA area and are translated by ABRs into type-5 external route advertisements for distribution into other areas of the OSPF domain. An NSSA area cannot be designated as the transit area of a virtual link.

In [Figure 5](#), area 0.0.0.3 could be configured as a NSSA area.

#### 3.1.1.3.1 OSPF Super Backbone

The 77x0 PE routers have implemented a version of the BGP/OSPF interaction procedures as defined in RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs). Features included in this RFC includes:

- Loop prevention
- Handling LSAs received from the CE
- Sham links
- Managing VPN-IPv4 routes received by BGP

VPN routes can be distributed among the PE routers by BGP. If the PE uses OSPF to distribute routes to the CE router, the standard procedures governing BGP/OSPF interactions causes routes from one site to be delivered to another in type 5 LSAs, as AS-external routes.

The MPLS VPN super backbone behaves like an additional layer of hierarchy in OSPF. The PE-routers that connect the respective OSPF areas to the super backbone function as OSPF Area Border Routers (ABR) in the OSPF areas to which they are attached. In order to achieve full compatibility, they can also behave as AS Boundary Routers (ASBR) in non-stub areas.


The PE-routers insert inter-area routes from other areas into the area where the CE-router is present. The CE-routers are not involved at any level, nor are they aware of the super backbone or of other OSPF areas present beyond the MPLS VPN super backbone.

The CE always assumes the PE is an ABR:

- If the CE is in the backbone, then the CE router assumes that the PE is an ABR linking one or more areas to the backbone.

- If the CE is not in the backbone, then the CE believes that the backbone is on the other side of the PE.
- As such, the super backbone looks like another area to the CE.

**Figure 6 PEs Connected to an MPLS-VPN Super Backbone**



In [Figure 6](#), the PEs are connected to the MPLS-VPN super backbone. In order to be able to distinguish if two OSPF instances are in fact the same and require Type 3 LSAs to be generated, or are two separate routing instances where type 5 external LSAs need to be generated, the concept of a domain-id is introduced.

The domain ID is carried with the MP-BGP update and indicates the source OSPF Domain. When the routes are being redistributed into the same OSPF Domain, the concepts of super backbone described above apply and Type 3 LSAs are generated. If the OSPF domain does not match, then the route type will be external.

Configuring the super backbone (not the sham links) makes all destinations learned by PEs with matching domain IDs inter-area routes.

When configuring sham links, these links become intra-area routes if they are present in the same area.

### 3.1.1.3.2 Sham Links

**Figure 7 Sham Links**




Figure 7 displays the red link between CE-3 and CE-4 could be a low speed OC-3/STM-1 link but because it establishes an intra-area route connection between the CE-3 and CE-4 the potentially high-speed PE-1 to PE-2 connection will not be utilized. Even with a super backbone configuration, it is regarded as an inter-area connection.

The establishment of the (green) sham-link is also constructed as an intra-area link between PE routers, a normal OSPF adjacency is formed and the link-state database is exchanged across the MPLS-VPN. As a result, the desired intra-area connectivity is created, at this time the cost of the green and red links can be managed such that the red link becomes a standby link only in case the VPN fails.

As the sham-link forms an adjacency over the MPLS-VPN backbone network, be aware that when protocol-protection is enabled in the **config>sys>security>cpu-protection>protocol-protection** context, the operator must explicitly allow the OSPF packets to be received over the backbone network. This is performed using the **allow-sham-links** parameter of the protocol-protection command.

### 3.1.1.3.3 Implementing the OSPF Super Backbone

With the OSPF super backbone architecture, the continuity of OSPF routing is preserved:

- The OSPF intra-area LSAs (type-1 and type-2) advertised by the CE are inserted into the MPLS-VPN super backbone by redistributing the OSPF route into MP-BGP by the PE adjacent to the CE.



- The MP-BGP route is propagated to other PE-routers and inserted as an OSPF route into other OSPF areas. Considering the PEs across the super backbone always act as ABRs they will generate inter area route OSPF summary LSAs, Type 3.
- The inter-area route can now be propagated into other OSPF areas by other customer owned ABRs within the customer site.
- Customer Area 0 (backbone) routes when carried across the MPLS-VPRN using MPBGP will appear as Type 3 LSAs even if the customer area remains area 0 (backbone).

A BGP extended community (OSPF domain ID) provides the source domain of the route. This domain ID is not carried by OSPF but carried by MP-BGP as an extended community attribute.

If the configured extended community value matches the receiving OSPF domain, then the OSPF super backbone is implemented.

From a BGP perspective, the cost is copied into the MED attribute.

#### **3.1.1.3.4 Loop Avoidance**

If a route sent from a PE router to a CE router could then be received by another PE router from one of its own CE routers then it is possible for routing loops to occur. RFC 4577 specifies several methods of loop avoidance.

#### **3.1.1.3.5 DN-BIT**

When a Type 3 LSA is sent from a PE router to a CE router, the DN bit in the LSA options field is set. This is used to ensure that if any CE router sends this Type 3 LSA to a PE router, the PE router will not redistribute it further.

When a PE router needs to distribute to a CE router a route that comes from a site outside the latter's OSPF domain, the PE router presents itself as an ASBR (Autonomous System Border Router), and distributes the route in a type 5 LSA. The DN bit must be set in these LSAs to ensure that they will be ignored by any other PE routers that receive them.

DN-BIT loop avoidance is also supported.

### 3.1.1.3.6 Route Tag

If a particular VRF in a PE is associated with an instance of OSPF, then by default it is configured with a special OSPF route tag value called the VPN route tag. This route tag is included in the Type 5 LSAs that the PE originates and sends to any of the attached CEs. The configuration and inclusion of the VPN Route Tag is required for backward compatibility with deployed implementations that do not set the DN bit in Type 5 LSAs.

### 3.1.1.3.7 Sham Links

A sham link is only required if a backdoor link (shown as the red link in [Figure 7](#)) is present, otherwise configuring an OSPF super backbone will probably suffice.

## 3.1.2 OSPFv3 Authentication

OSPFv3 authentication requires IPv6 IPsec and supports the following:

- IPsec transport mode
- AH and ESP
- Manual keyed IPsec Security Association (SA)
- Authentication Algorithms MD5 and SHA1

To pass OSPFv3 authentication, OSPFv3 peers must have matching inbound and outbound SAs configured using the same SA parameters (SPI, keys, etc.). The implementation must allow the use of one SA for both inbound and outbound directions.

This feature is supported on IES and VPRN interfaces as well as on virtual links.

The re-keying procedure defined in RFC 4552, *Authentication/Confidentiality for OSPFv3*, supports the following:

- For every router on the link, create an additional inbound SA for the interface being re-keyed using a new SPI and the new key.
- For every router on the link, replace the original outbound SA with one using the new SPI and key values. The SA replacement operation should be atomic with respect to sending OSPFv3 packet on the link so that no OSPFv3 packets are sent without authentication or encryption.
- For every router on the link, remove the original inbound SA.

The key rollover procedure automatically starts when the operator changes the configuration of the inbound static-sa or bi-directional static-sa under an interface or virtual link. Within the KeyRolloverInterval time period, OSPF3 accepts packets with both the previous inbound static-sa and the new inbound static-sa, and the previous outbound static-sa should continue to be used. When the timer expires, OSPF3 will only accept packets with the new inbound static-sa and for outgoing OSPF3 packets, the new outbound static-sa will be used instead.

### 3.1.2.1 OSPFv3 Graceful Restart Helper


This feature extends the Graceful Restart helper function supported under other protocols to OSPFv3.

The primary difference between graceful restart helper for OSPFv2 and OSPFv3 is in OSPFv3 a different grace-LSA format is used.

As SR OS platforms can support a fully non-stop routing model for control plane high availability, SR OSs have no need for graceful restart as defined by the IETF in various RFCs for each routing protocol. However, since the router does need to co-exist in multi-vendor networks and other routers do not always support a true non-stop routing model with stateful failover between routing control planes, there is a need to support a Graceful Restart Helper function.

Graceful restart helper mode allows SR OS-based systems to provide a grace period to other routers which have requested it, during which the SR OS systems will continue to use routes authored by or transiting the router requesting the grace period. This is typically used when another router is rebooting the control plane but the forwarding plane is expected to continue to forward traffic based on the previously available FIB.


The format of the Graceful OSPF restart (GRACE) LSA format is:



See section 2.2 of RFC 5187, *OSPFv3 Graceful Restart*.

The Link State ID of a grace-LSA in OSPFv3 is the Interface ID of the interface originating the LSA.

The format of each TLV is:



TLV Format

Grace-LSA TLVs are formatted according to section 2.3.2 of RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*. The Grace-LSA TLVs are used to carry the Grace period (type 1) and the reason the router initiated the graceful restart process (type 2).

Other information in RFC 5187 is directed to routers that require the full graceful restart mechanism as they do not support a stateful transition from primary or backup control plane module (CPM).

### 3.1.3 Virtual Links

The backbone area in an OSPF AS must be contiguous and all other areas must be connected to the backbone area. Sometimes, this is not possible. You can use virtual links to connect to the backbone through a non-backbone area.

[Figure 5](#) depicts routers Y and Z as the start and end points of the virtual link while area 0.0.0.4 is the transit area. In order to configure virtual links, the router must be an ABR. Virtual links are identified by the router ID of the other endpoint, another ABR. These two endpoint routers must be attached to a common area, called the transit area. The area through which you configure the virtual link must have full routing information.

Transit areas pass traffic from an area adjacent to the backbone or to another area. The traffic does not originate in, nor is it destined for, the transit area. The transit area cannot be a stub area or a NSSA area.

Virtual links are part of the backbone, and behave as if they were unnumbered point-to-point networks between the two routers. A virtual link uses the intra-area routing of its transit area to forward packets. Virtual links are brought up and down through the building of the shortest-path trees for the transit area.

## **3.1.4 Neighbors and Adjacencies**

### **3.1.4.1 Broadcast and Point-to-Point Networks**

A router uses the OSPF Hello protocol to discover neighbors. The router sends hello packets to a multicast address and receives hello packets in return.

In broadcast networks, a designated router and a backup designated router are elected. The designated router is responsible for sending link-state advertisements (LSAs) describing the network, which reduces the amount of network traffic.

The routers attempt to form adjacencies. An adjacency is a relationship formed between a router and the designated or backup designated router. For point-to-point networks, no designated or backup designated router is elected. An adjacency must be formed with the neighbor.

To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast medium such as Ethernet.

When the link-state databases of two neighbors are synchronized, the routers are considered to be fully adjacent. When adjacencies are established, pairs of adjacent routers synchronize their topological databases. Not every neighboring router forms an adjacency. Routing protocol updates are only sent to and received from adjacencies. Routers that do not become fully adjacent remain in the two-way neighbor state.

### **3.1.4.2 Non-Broadcast Multi-Access (NBMA) Networks**

In addition to point-to-point and broadcast networks, OSPF can operate in non-broadcast multi-access (NBMA) mode.

An NBMA segment emulates the function of a broadcast network. Every router on the segment must be configured with the IP addresses of each of its neighbors, and may need to be configured with the MAC address of its neighbor if the network does not support Layer 2 broadcast. OSPF Hello packets are transmitted individually as unicast packets to each adjacent neighbor. Because an NBMA network has no broadcast or multicast capabilities, the routing device cannot discover its neighbors dynamically, so all neighbors must be configured statically.

As in a broadcast network, a designated router and a backup designated router are elected when OSPF is operating in NBMA mode. The designated router is similarly responsible for sending link-state advertisements (LSAs) for the network.

OSPF does not support NBMA interfaces that are part of a multi-area adjacency. An interface can either be in multiple areas or in NBMA mode.

**Note:**

- OSPFv3 sends Hello traffic to IPv6 link-local addresses and neighbors must be statically configured using their IPv6 link-local address.
- OSPF NBMA is supported on Ethernet ports only.

### 3.1.5 Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

### 3.1.6 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

### 3.1.7 Authentication

All OSPF protocol exchanges can be authenticated. This means that only trusted routers can participate in autonomous system routing. Nokia's implementation of OSPF supports plain text and Message Digest 5 (MD5) authentication (also called simple password).

MD5 allows an authentication key to be configured per network. Routers in the same routing domain must be configured with the same key. When the MD5 hashing algorithm is used for authentication, MD5 is used to verify data integrity by creating a 128-bit message digest from the data input. It is unique to that data. Nokia's implementation of MD5 allows the migration of an MD5 key by using a key ID for each unique key.

By default, authentication is not enabled on an interface.

### 3.1.8 IP Subnets

OSPF enables the flexible configuration of IP subnets. Each distributed OSPF route has a destination and mask. A network mask is a 32-bit number that indicates the range of IP addresses residing on a single IP network/subnet. This specification displays network masks as hexadecimal numbers; for example, the network mask for a class C IP network is displayed as 0xfffff00. Such a mask is often displayed as 255.255.255.0.

Two different subnets with same IP network number have different masks, called variable length subnets. A packet is routed to the longest or most specific match. Host routes are considered to be subnets whose masks are all ones (0xffffffff).

### 3.1.9 Preconfiguration Recommendations

Prior to configuring OSPF, the router ID must be available. The router ID is a 32-bit number assigned to each router running OSPF. This number uniquely identifies the router within an AS. OSPF routers use the router IDs of the neighbor routers to establish adjacencies. Neighbor IDs are learned when Hello packets are received from the neighbor.

Before configuring OSPF parameters, ensure that the router ID is derived by one of the following methods:

- Define the value in the **config>router** *router-id* context.
- Define the system interface in the **config>router>interface** *ip-int-name* context (used if the router ID is not specified in the **config>router** *router-id* context).  
A system interface must have an IP address with a 32-bit subnet mask. The system interface is used as the router identifier by higher-level protocols such as OSPF and IS-IS. The system interface is assigned during the primary router configuration process when the interface is created in the logical IP interface context.
- If you do not specify a router ID, then the last four bytes of the MAC address are used.



**Note:** On the BGP protocol level, a BGP router ID can be defined in the **config>router>bgp** *router-id* context and is only used within BGP.



---

## 3.1.10 Multiple OSPF Instances

The main route table manager (RTM) can create multiple instances of OSPF by extending the current creation of an instance. A given interface can only be a member of a single OSPF instance. When an interface is configured in a given domain and needs to be moved to another domain the interface must first be removed from the old instance and re-created in the new instance.

### 3.1.10.1 Route Export Policies for OSPF

Route policies allow specification of the source OSPF process ID in the **from** and **to** parameters in the **config>router>policy-options>policy-statement>entry>from** context, for example from **protocol ospf instance-id**.

If an instance-id is specified, only routes installed by that instance are picked up for announcement. If no instance-id is specified, then only routes installed by the base instance will be announced. The **all** keyword announces routes installed by all instances of OSPF.

When announcing internal (intra/inter-area) OSPF routes from another process, the default type should be type-1, and metric set to the route metric in RTM. For AS-external routes, by default the route type (type-1/2) should be preserved in the originated LSA, and metric set to the route metric in RTM. By default, the tag value should be preserved when an external OSPF route is announced by another process. All these can be changed with explicit action statements.

Export policy should allow a match criteria based on the OSPF route hierarchy, e.g. only intra-area, only inter-area, only external, only internal (intra/inter-area). There must also be a possibility to filter based on existing tag values.

### 3.1.10.2 Preventing Route Redistribution Loops

The legacy method for this was to assign a tag value to each OSPF process and mark each external route originated within that domain with that value. However, since the tag value must be preserved throughout different OSPF domains, this only catches loops that go back to the originating domain and not where looping occurs in a remote set of domains. To prevent this type of loop, the route propagation information in the LSA must be accumulative. The following method has been implemented:

- The OSPF tag field in the AS-external LSAs is treated as a bit mask, rather than a scalar value. In other words, each bit in the tag value can be independently checked, set or reset as part of the routing policy.
- When a set of OSPF domains are provisioned in a network, each domain is assigned a specific bit value in the 32-bit tag mask. When an external route is originated by an ASBR using an internal OSPF route in a given domain, a corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy--if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

From the CLI perspective, this involves adding a set of **from tag** and **action tag** commands that allow for bit operations.

### 3.1.11 Multi-Address Support for OSPFv3

While OSPFv3 was originally designed to carry only IPv6 routing information, the protocol has been extended to add support for other address families through work within the IETF (RFC 5838). These extensions within SR OS allow separate OSPFv3 instances to be used for IPv6 or IPv4 routing information.

To configure an OSPFv3 instance to distribute IPv4 routing information, a specific OSPFv3 instance must be configured using an instance ID within the range specified by the RFC. For unicast IPv4, the range is 64 to 95.

The following shows the basic configuration steps needed to create the OSPFv3 (ospf3) instance to carry IPv4 routing information. Once the instance is created, the OSPFv3 instance can be configured as needed for the associated network areas, interfaces, and other protocol attributes as you would for OSPFv2.

**Example:**

```
config
 router
 ospf3 64 10.20.1.3
```

---

### 3.1.12 IP Fast-reroute (IP FRR) For OSPF and IS-IS Prefixes

This feature provides for the use of the Loop-Free Alternate (LFA) backup next hop for forwarding in-transit and CPM generated IP packets when the primary next hop is not available. This means that a node resumes forwarding IP packets to a destination prefix without waiting for the routing convergence.

When any of the following events occurs, IGP instructs in the fast path the IOM or the forwarding engine to enable the LFA backup next hop:

- OSPF/IS-IS interface goes operationally down: physical or local admin shutdown.
- Timeout of a BFD session to a next hop when BFD is enabled on the OSPF/IS-IS interface.

IP FRR is supported on IPv4 and IPv6 OSPF/IS-IS prefixes forwarded in the base router instance to a network IP interface or to an IES SAP interface or spoke interface. It is also supported for VPRN VPN-IPv4 OSPF prefixes and VPN-IPv6 OSPF prefixes forwarded to a VPRN SAP interface or spoke interface.

IP FRR also provides a LFA backup next hop for the destination prefix of a GRE tunnel used in an SDP or in VPRN auto-bind.

The LFA next hop pre-computation by IGP is described in RFC 5286 *Basic Specification for IP Fast Reroute: Loop-Free Alternates*.

#### 3.1.12.1 IP FRR Configuration

The user first enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level:

**CLI Syntax:**

```
config>router>isis>loopfree-alternate
config>router>ospf>loopfree-alternate
config>service>vprn>ospf>loopfree-alternate
```

The above commands instruct the IGP SPF to attempt to pre-compute both a primary next hop and an LFA next hop for every learned prefix. When found, the LFA next hop is populated into the RTM along with the primary next hop for the prefix.

Next the user enables IP FRR to cause RTM to download to the IOM or the forwarding engine a LFA next hop, when found by SPF, in addition to the primary next hop for each prefix in the FIB.

---

**CLI Syntax:**    `config>router>ip-fast-reroute`

### 3.1.12.1.1 Reducing the Scope of the LFA Calculation by SPF

The user can instruct IGP to not include all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

**CLI Syntax:**    `config>router>isis>level>loopfree-alternate-exclude`  
                  `config>router>ospf>area>loopfree-alternate-exclude`

The user can also exclude a specific IP interface from being included in the LFA SPF computation by IS-IS or OSPF:

**CLI Syntax:**    `config>router>isis>if>loopfree-alternate-exclude`  
                  `config>router>ospf>area>if>loopfree-alternate-exclude`

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When the user excludes an interface from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

Finally, the user can apply the same above commands for an OSPF instance within a VPRN service:

**CLI Syntax:**    `config>service>vprn>ospf>area>loopfree-alternate-exclude`  
                  `config>service>vprn>ospf>area>if>loopfree-alternate-exclude`

### 3.1.12.2 ECMP Considerations

Whenever the SPF computation determined there is more than one primary next hop for a prefix, it will not program any LFA next hop in RTM. Thus, IP prefixes will resolve to the multiple primary next hops in this case which provides the required protection.

---

### 3.1.12.3 IP FRR and RSVP Shortcut (IGP Shortcut)

When both IGP shortcut and LFA are enabled in IS-IS or OSPF, and IP FRR is also enabled, then the following additional IP FRR are supported:

- A prefix which is resolved to a direct primary next hop can be backed up by a tunneled LFA next hop.
- A prefix which is resolved to a tunneled primary next hop will not have an LFA next hop. It will rely on RSVP FRR for protection.

The LFA SPF is extended to use IGP shortcuts as LFA next hops as explained in [OSPF and IS-IS Support for Loop-Free Alternate Calculation](#).


### 3.1.12.4 IP FRR and BGP Next Hop Resolution

An LFA backup next hop will be able to protect the primary next hop to reach a prefix advertised by a BGP neighbor. The BGP next hop will remain up when the FIB switches from the primary IGP next hop to the LFA IGP next hop.

### 3.1.12.5 OSPF and IS-IS Support for Loop-Free Alternate Calculation

SPF computation in IS-IS and OSPF is enhanced to compute LFA alternate routes for each learned prefix and populate it in RTM.

[Figure 8](#) illustrates a simple network topology with point-to-point (P2P) interfaces and highlights three routes to reach router R5 from router R1.

**Figure 8 Example Topology with Primary and LFA Routes**


OSSG712

The primary route is via R3. The LFA route via R2 has two equal cost paths to reach R5. The path by way of R3 protects against failure of link R1-R3. This route is computed by R1 by checking that the cost for R2 to reach R5 by way of R3 is lower than the cost by way of routes R1 and R3. This condition is referred to as the “loop-free criterion”.

The path by way of R2 and R4 can be used to protect against the failure of router R3. However, with the link R2-R3 metric set to 5, R2 sees the same cost to forward a packet to R5 by way of R3 and R4. Thus R1 cannot guarantee that enabling the LFA next hop R2 will protect against R3 node failure. This means that the LFA next hop R2 provides link-protection only for prefix R5. If the metric of link R2-R3 is changed to 8, then the LFA next hop R2 provides node protection since a packet to R5 will always go over R4. In other words it is required that R2 becomes loop-free with respect to both the source node R1 and the protected node R3.

Consider now the case where the primary next hop uses a broadcast interface as illustrated in [Figure 9](#).

**Figure 9 Example Topology with Broadcast Interfaces**



OSSG713

In order for next hop R2 to be a link-protect LFA for route R5 from R1, it must be loop-free with respect to the R1-R3 link Pseudo-Node (PN). However, since R2 has also a link to that PN, its cost to reach R5 by way of the PN, or router R4 are the same. Thus R1 cannot guarantee that enabling the LFA next hop R2 will protect against a failure impacting link R1-PN since this may cause the entire subnet represented by the PN to go down. If the metric of link R2-PN is changed to 8, then R2 next hop will be an LFA providing link protection.

The following are the detailed equations for this criterion as provided in RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*:

- **Rule 1:** Link-protect LFA backup next hop (primary next hop R1-R3 is a P2P interface):  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$   
and,  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) \geq \text{Distance\_opt}(\text{R2}, \text{R3}) + \text{Distance\_opt}(\text{R3}, \text{R5})$
- **Rule 2:** Node-protect LFA backup next hop (primary next hop R1-R3 is a P2P interface):  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$   
and,  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R3}) + \text{Distance\_opt}(\text{R3}, \text{R5})$

- **Rule 3:** Link-protect LFA backup next hop (primary next hop R1-R3 is a broadcast interface):

$\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$

and,

$\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{PN}) + \text{Distance\_opt}(\text{PN}, \text{R5})$

where; PN stands for the R1-R3 link Pseudo-Node.

For the case of P2P interface, if SPF finds multiple LFA next hops for a given primary next hop, it follows the following selection algorithm:

- It will pick the node-protect type in favor of the link-protect type.
- If there is more than one LFA next hop within the selected type, then it will pick one based on the least cost.
- If more than one LFA next hop with the same cost results from step **b**, then SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

For the case of a broadcast interface, a node-protect LFA is not necessarily a link protect LFA if the path to the LFA next hop goes over the same PN as the primary next hop. Similarly, a link protect LFA may not guarantee link protection if it goes over the same PN as the primary next hop. The selection algorithm when SPF finds multiple LFA next hops for a given primary next hop is modified as follows:

- The algorithm splits the LFA next hops into two sets:
  - The first set consists of LFA next hops which do not go over the PN used by primary next hop.
  - The second set consists of LFA next hops which do go over the PN used by the primary next hop.
- If there is more than one LFA next hop in the first set, it will pick the node-protect type in favor of the link-protect type.
- If there is more than one LFA next hop within the selected type, then it will pick one based on the least cost.
- If more than one LFA next hop with equal cost results from Step C, SPF will select the first one from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
- If no LFA next hop results from Step D, SPF will rerun Steps B-D using the second set.

This algorithm is more flexible than strictly applying Rule 3 above; i.e., the link protect rule in the presence of a PN and specified in RFC 5286. A node-protect LFA which does not avoid the PN; i.e., does not guarantee link protection, can still be selected as a last resort. The same thing, a link-protect LFA which does not avoid the PN may still be selected as a last resort.



Both the computed primary next hop and LFA next hop for a given prefix are programmed into RTM.

### 3.1.12.5.1 Loop-Free Alternate Calculation in the Presence of IGP Shortcuts

In order to expand the coverage of the LFA backup protection in a network, RSVP LSP based IGP shortcuts can be placed selectively in parts of the network and be used as an LFA backup next hop.

When IGP shortcut is enabled in IS-IS or OSPF on a given node, all RSVP LSP originating on this node and with a destination address matching the router-id of any other node in the network are included in the main SPF by default.

In order to limit the time it takes to compute the LFA SPF, the user must explicitly enable the use of an IGP shortcut as LFA backup next hop using one of a couple of new optional argument for the existing LSP level IGP shortcut command:

**config>router>mpls>lsp>igp-shortcut [lfa-protect | lfa-only]**

The **lfa-protect** option allows an LSP to be included in both the main SPF and the LFA SPFs. For a given prefix, the LSP can be used either as a primary next hop or as an LFA next hop but not both. If the main SPF computation selected a tunneled primary next hop for a prefix, the LFA SPF will not select an LFA next hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next hop, then the LFA SPF will select an LFA next hop for this prefix but will prefer a direct LFA next hop over a tunneled LFA next hop.

The **lfa-only** option allows an LSP to be included in the LFA SPFs only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next hop. The LFA SPF will select a an LFA next hop for this prefix but will prefer a direct LFA next hop over a tunneled LFA next hop.

Thus, with the selection algorithm when SPF finds multiple LFA next hops for a given primary next hop is modified as follows:

- a. The algorithm splits the LFA next hops into two sets:
  - the first set consists of direct LFA next hops
  - the second set consists of tunneled LFA next hops. after excluding the LSPs which use the same outgoing interface as the primary next hop.
- b. The algorithms continues with first set if not empty, otherwise it continues with second set.

- 
- c. If the second set is used, the algorithm selects the tunneled LFA next hop which endpoint corresponds to the node advertising the prefix.
    - If more than one tunneled next hop exists, it selects the one with the lowest LSP metric.
    - If still more than one tunneled next hop exists, it selects the one with the lowest tunnel-id.
    - If none is available, it continues with rest of the tunneled LFAs in second set.
  - d. Within the selected set, the algorithm splits the LFA next hops into two sets:
    - The first set consists of LFA next hops which do not go over the PN used by primary next hop.
    - The second set consists of LFA next hops which go over the PN used by the primary next hop.
  - e. If there is more than one LFA next hop in the selected set, it will pick the node-protect type in favor of the link-protect type.
  - f. If there is more than one LFA next hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
  - g. If there is more than one LFA next hop within the selected type (ecmp-case) in the first set, it will select the first direct next hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
  - h. If there is more than one LFA next hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next hop with the lowest tunnel-id.

#### **3.1.12.5.2 Loop-Free Alternate Calculation for Inter-Area/inter-Level Prefixes**

When SPF resolves OSPF inter-area prefixes or IS-IS inter-level prefixes, it will compute an LFA backup next hop to the same exit area/border router as used by the primary next hop.

## 3.2 Loop-Free Alternate Shortest Path First (LFA SPF) Policies

An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of a LFA backup next hop for a subset of prefixes that resolve to a specific primary next hop. The feature introduces the concept of route next hop template to influence LFA backup next hop selection.

### 3.2.1 Configuring a Route Next Hop Policy Template

The LFA SPF policy consists of applying a route next hop policy template to a set of prefixes.

The user first creates a route next hop policy template under the global router context:

**CLI Syntax:** `config>router>route-next-hop-policy>template template-name`

A policy template can be used in both IS-IS and OSPF to apply the specific criteria described in the next sub-sections to prefixes protected by LFA. Each instance of IS-IS or OSPF can apply the same policy template to one or more prefix lists and to one or more interfaces.

The commands within the route next hop policy use the **begin-commit-abort** model introduced with BFD templates. The following are the steps to create and modify the template:

- To create a template, the user enters the name of the new template directly under **route-next-hop-policy** context.
- To delete a template which is not in use, the user enters the no form for the template name under the **route-next-hop-policy** context.
- The user enters the editing mode by executing the **begin** command under **route-next-hop-policy** context. The user can then edit and change any number of route next hop policy templates. However, the parameter value will still be stored temporarily in the template module until the **commit** is executed under the **route-next-hop-policy** context. Any temporary parameter changes will be lost if the user enters the **abort** command before the commit command.

- The user is allowed to create or delete a template instantly once in the editing mode without the need to enter the **commit** command. Furthermore, the **abort** command if entered will have no effect on the prior deletion or creation of a template.

After the **commit** command is issued, IS-IS or OSPF will re-evaluate the templates and if there are any net changes, it will schedule a new LFA SPF to re-compute the LFA next hop for the prefixes associated with these templates.

### 3.2.1.1 Configuring Affinity or Admin Group Constraints

Administrative groups (admin groups), also known as affinity, are used to tag IP interfaces which share a specific characteristic with the same identifier. For example, an admin group identifier could represent all links which connect to core routers, or all links which have bandwidth higher than 10G, or all links which are dedicated to a specific service.

The user first configures locally on each router the name and identifier of each admin group:

**CLI Syntax:** `config>router>if>if-attribute>admin-group group-name`  
`value group-value`

A maximum of 32 admin groups can be configured per system.

Next the user configures the admin group membership of the IP interfaces used in LFA. The user can apply admin groups to IES, VPRN, or network IP interface.

**CLI Syntax:**

```
config>router> interface>if-attribute>admin-group group-
name [group-name...(up to 5 max)]
config>service>ies>if>if-attribute>admin-group group-
name [group-name...(up to 5 max)]
config>service>vprn >if>if-attribute>admin-group group-
name [group-name...(up to 5 max)]
```

The user can add as many admin groups as configured to a given IP interface. The same above command can be applied multiple times.



**Note:** the configured **admin-group** membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

The **no** form of the **admin-group** command under the interface deletes one or more of the admin-group memberships of the interface. It deletes all memberships if no group name is specified.

Finally, the user adds the admin group constraint into the route next hop policy template:

**CLI Syntax:**

```
configure router route-next-hop-template template
 template-name
 include-group group-name [pref 1]
 include-group group-name [pref 2]
 exclude-group group-name
```

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next hops among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in a **include-group** statement but also belongs to other groups which are not part of any **include-group** statement in the route next hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select a LFA backup next hop which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, i.e., numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next hop selection for a prefix.

If the same group name is part of both **include** and **exclude** statements, the **exclude** statement will win. In other words, the **exclude** statement can be viewed as having an implicit preference value of 0.



**Note:** the admin-group criterion is applied before running the LFA next hop selection algorithm.

### 3.2.1.2 Configuring SRLG Group Constraints

Shared Risk Loss Group (SRLG) is used to tag IP interfaces which share a specific fate with the same identifier. For example, an SRLG group identifier could represent all links which use separate fibers but are carried in the same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all IP interfaces using these fiber links will fail. Thus the user can enable the SRLG constraint to select a LFA next hop for a prefix which avoids all interfaces that share fate with the primary next.

The user first configures locally on each router the name and identifier of each SRLG group:

**CLI Syntax:** `config>router>if>if-attribute>srlg-group group-name value group-value`

A maximum of 1024 SRLGs can be configured per system.

Next the user configures the admin group membership of the IP interfaces used in LFA. The user can apply SRLG groups to IES, VPRN, or network IP interface.

**CLI Syntax:** `config>router>if>if-attribute>srlg-group group-name [group-name...(up to 5 max)]`  
`config>service>vprn>if>if-attribute>srlg-group group-name [group-name...(up to 5 max)]`  
`config>service>ies>if>if-attribute>srlg-group group-name [group-name...(up to 5 max)]`

The user can add a maximum of 64 SRLG groups to a given IP interface. The same above command can be applied multiple times.



**Note:** the configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

The **no** form of the **srlg-group** command under the interface deletes one or more of the SRLG memberships of the interface. It deletes all SRLG memberships if no group name is specified.

Finally, the user adds the SRLG constraint into the route next hop policy template:

**CLI Syntax:** `configure router route-next-hop-template template template-name srlg-enable`

When this command is applied to a prefix, the LFA SPF will select a LFA next hop, among the computed ones, which uses an outgoing interface that does not participate in any of the SRLGs of the outgoing interface used by the primary next hop.



**Note:** the SRLG and admin-group criteria are applied before running the LFA next hop selection algorithm.

### 3.2.1.3 Interaction of IP and MPLS Admin Group and SRLG

The LFA SPF policy feature generalizes the use of admin-group and SRLG to other types of interfaces. To that end, it is important that the IP admin groups and SRLGs be compatible with the ones already supported in MPLS. The following rules are implemented:

- The binding of an MPLS interface to a group, that is, configuring membership of an MPLS interface in a group, is performed under **config>router>mpls>interface** context.
- The binding of a local or remote MPLS interface to an SRLG in the SRLG database is performed under the **config>router>mpls>srlg-database** context.
- The binding of an IS-IS/OSPF interface to a group is performed in the **config>router>if>if-attribute** or **config>service>vprn>if>if-attribute** or **config>service>ies>if>if-attribute** contexts. This is used by IS-IS or OSPF in route next hop policies.
- Only the admin groups and SRLGs bound to an MPLS interface context or the SRLG database context are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

### 3.2.1.4 Configuring Protection Type and Next Hop Type Preferences

The user can select if link protection or node protection is preferred in the selection of a LFA next hop for all IP prefixes and LDP FEC prefixes to which a route next hop policy template is applied. The default in SR OS implementation is node protection. The implementation will fall back to the other type if no LFA next hop of the preferred type is found.

The user can also select if tunnel backup next hop or IP backup next hop is preferred. The default in SR OS implementation is to prefer IP next hop over tunnel next hop. The implementation will fall back to the other type if no LFA next hop of the preferred type is found.

The following options are thus added into the Route next hop policy template:

**CLI Syntax:**

```
configure router route-nh-template template template-name
 protection-type {link | node}
 nh-type {ip | tunnel}
```

When the route next hop policy template is applied to an IP interface, all prefixes using this interface as a primary next hop will follow the protection type and next hop type preference specified in the template.

### 3.2.2 Application of Route Next Hop Policy Template to an Interface

Once the route next hop policy template is configured with the desired policies, the user can apply it to all prefixes which primary next hop uses a specific interface name. The following command achieves that:

**CLI Syntax:**

```
config>router>isis>if>lfa-policy-map route-nh-template template-name
config>router>ospf(3)>area>if>lfa-policy-map route-nh-template template-name
config>service>vprn>ospf(3)>area>if>lfa-policy-map route-nh-template template-name
```

When a route next hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the above CLI command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy if applied to the interface has no effect.

Finally, if the user applied a route next hop policy template to a loopback interface or to the system interface, the command will not be rejected but it will result in no action taken.



### 3.2.3 Excluding Prefixes from LFA SPF

In the current SR OS implementation, the user can exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

This feature adds the ability to exclude prefixes from a prefix policy which matches on prefixes or on IS-IS tags:

**CLI Syntax:**

```
config>router>isis>loopfree-alternate-exclude prefix-
policy prefix-policy1 [prefix-policy2...up to 5]
config>router>ospf(3)>loopfree-alternate-exclude
prefix-policy prefix-policy1 [prefix-policy2...up to 5]
config>service>vprn>ospf(3)>loopfree-alternate-exclude
prefix-policy prefix-policy1 [prefix-policy2...up to 5]
```

The prefix policy is configured as in existing SR OS implementation:

**CLI Syntax:**

```
config
router
 policy-options
 [no] prefix-list prefix-list1
 prefix 62.225.16.0/24 prefix-length-
 range 32-32
 [no] policy-statements prefix-policy1
 entry 10
 from
 prefix-list "prefix-list1"
 exit
 action accept
 exit
 exit
 exit
 default-action reject
exit
```

If the user enabled the IS-IS prefix prioritization based on tag, it will also apply to SPF LFA. If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority; however, the prefix tag will be used in the main SPF.



**Note:** Prefix tags are not defined for OSPF protocol.

The default action of the above **loopfree-alternate-exclude** command when not explicitly specified by the user in the prefix policy is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix which did not match any entry in the policy will be accepted into LFA SPF.

### 3.2.4 Modification to LFA Next Hop Selection Algorithm

This feature modifies the LFA next-hop selection algorithm. The SRLG and admin-group criteria are applied before running the LFA next-hop selection algorithm. In other words, links which do not include one or more of the admin-groups in the **include-group** statements and links which belong to admin-groups which have been explicitly excluded using the **exclude-group** statement, and the links which belong to the SRLGs used by the primary next hop of a prefix are first pruned.

This pruning applies only to IP next hops. Tunnel next hops can have the admin-group or SRLG constraint applied to them under MPLS. For example, if a tunnel next hop is using an outgoing interface which belongs to given SRLG ID, the user can enable the **srlg-frr** option under the **config>router>mpls** context to be sure the RSVP LSP FRR backup LSP will not use an outgoing interface with the same SRLG ID. A prefix which is resolved to a tunnel next hop is protected by the RSVP FRR mechanism and not by the IP FRR mechanism. Similarly, the user can include or exclude admin-groups for the RSVP LSP and its FRR bypass backup LSP in MPLS context. The admin-group constraints will, however, be applied to the selection of the outgoing interface of both the LSP primary path and its FRR bypass backup path.

The following is the modified LFA selection algorithm which is applied to prefixes resolving to a primary next hop which uses a given route next hop policy template.

- Split the LFA next hops into two sets:
  - IP or direct next hops.
  - Tunnel next hops after excluding the LSPs which use the same outgoing interface as the primary next hop.
- Prune the IP LFA next hops which use the following links:
  - links which do not include one or more of the admin-groups in the **include-group** statements in the route next hop policy template.
  - links which belong to admin-groups which have been explicitly excluded using the **exclude-group** statement in the route next hop policy template.
  - links which belong to the SRLGs used by the primary next hop of a prefix.
- Continue with the set indicated in the **nh-type** value in the route next hop policy template if not empty; otherwise continue with the other set.

- 
- Within IP next hop set:
    - prefer LFA next hops which do not go over the Pseudo-Node (PN) used by the primary next hop
    - Within selected subset prefer the node-protect type or the link-protect type according to the value of the **protection-type** option in the route next hop policy template.
    - Within the selected subset, select the best admin-group(s) according to the preference specified in the value of the **include-group** option in the route next hop policy template.
    - Within selected subset, select lowest **total cost** of a prefix.
    - If same **total cost**, select lowest **router-id**.
    - If same **router-id**, select lowest **interface-index**.
  - Within tunnel next hop set:
    - Select tunnel next hops which endpoint corresponds to the node owning or advertising the prefix.
      - Within selected subset, select the one with the lowest cost (lowest LSP metric).
      - If same lowest cost, select tunnel with lowest **tunnel-index**.
    - If none is available, continue with rest of the tunnel LFA next hop set.
    - Prefer LFA next hops which do not go over the Pseudo-Node (PN) used by the primary next hop.
    - Within selected subset prefer the node-protect type or the link-protect type according to the value of the **protection-type** in the route next hop policy template.
    - Within selected subset, select lowest **total cost** of a prefix. For a tunnel next hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
    - If same **total cost**, select lowest **endpoint to destination cost**
    - If same **endpoint to destination cost**, select lowest **router-id**,
    - If same **router-id**, select lowest **tunnel-index**.

### 3.3 LFA Protection Using Segment Routing Backup Node SID


One of the challenges in MPLS deployments across multiple IGP areas or domains, such as in seamless MPLS design, is the provisioning of FRR local protection in access and metro domains that make use of a ring, a square, or a partial mesh topology. In order to implement IP, LDP, or SR FRR in these topologies, the remote LFA feature must be implemented. Remote LFA provides a Segment Routing (SR) tunneled LFA next hop for an IP prefix, an LDP tunnel, or an SR tunnel. For prefixes outside of the area or domain, the access or aggregation router must push four labels: service label, BGP label for the destination PE, LDP/RSVP/SR label to reach the exit ABR/ASBR, and one label for the remote LFA next hop. Small routers deployed in these parts of the network have limited MPLS label stack size support.

Figure 10 illustrates the label stack required for the primary next hop and the remote LFA next hop computed by aggregation node AGN2 for the inter-area prefix of a remote PE. For an inter-area BGP label unicast route prefix for which ABR1 is the primary exit ABR, AGN2 resolves the prefix to the transport tunnel of ABR1 and therefore, uses the remote LFA next hop of ABR1 for protection. The primary next hop uses two transport labels plus a service label. The remote LFA next hop for ABR1 uses PQ node AGN5 and pushes three transport labels plus a service label.

Seamless MPLS with Fast Restoration requires up to four labels to be pushed by AGN2, as shown in Figure 10.

Figure 10 Label Stack for Remote LFA in Ring Topology

| Label Location   | Label Name             | Assigned By | Protocol      | Use                             |
|------------------|------------------------|-------------|---------------|---------------------------------|
| Label 1 (Bottom) | Service (PW, VC) Label | Remote PE   | MP-BGP, T-LDP | Identifies Service on Remote PE |
| Label 2          | Inter-Area Label       | ABR1        | BGP-LU        | Identifies Path to Remote PE    |
| Label 3          | Intra-Area             | AGN1        | LDP, RSVP, SR | Identifies Path to ABR1         |
| Label 4 (Top)    | R-LFA Label            | AGN3        | LDP, RSVP, SR | Identifies Path to AGN5         |



0935

The objective of the LFA protection with a backup node SID feature is to reduce the label stack pushed by AGN2 for BGP label unicast inter-area prefixes. When link AGN2-AGN1 fails, packets are direct away from the failure and forwarded toward ABR2, which acts as the backup for ABR1 (and vice-versa when ABR2 is the primary exit ABR for the BGP label unicast inter-area prefix). This requires that ABR2 advertise a special label for the loopback of ABR1 that will attract packets normally destined for ABR1. These packets will be forwarded by ABR2 to ABR1 via the inter-ABR link.

As a result, AGN2 will push the label advertised by ABR2 to back up ABR1 on top of the BGP label for the remote PE and the service label. This keeps the label stack the same size for the LFA next hop to be the same size as that of the primary next hop. It is also the same size as the remote LFA next hop for the local prefix within the ring.

### 3.3.1 Configuring LFA Using Backup Node SID

LFA using a backup node SID is enabled by configuring a backup node SID at an ABR/ASBR that acts as a backup to the primary exit ABR/ASBR of inter-area/inter-as routes learned as BGP labeled routes.

**CLI Syntax:**

```
config>router>ospf>segment-routing$
 backup-node-sid ip-prefix/prefix-length index
 0..4294967295
 backup-node-sid ip-prefix/prefix-length label
 1..4294967295
```

The user can enter either a label or an index for the backup node SID.




**Note:** This feature only allows the configuration of a single backup node SID per IGP instance and per ABR/ASBR. In other words, only a pair of ABR/ASBR nodes can back up each other in an IGP domain. Each time the user invokes the above command within the same IGP instance, it overrides any previous configuration of the backup node SID. The same ABR/ASBR can, however, participate in multiple IGP instances and provide a backup support within each instance.

### 3.3.2 Detailed Operation of LFA Protection Using Backup Node SID

As shown in [Figure 11](#), LFA for seamless MPLS supports environments where the boundary routers are either:

- ABR nodes that connect with iBGP multiple domains, each using a different area of the same IGP instance
- ASBR nodes that connect domains running different IGP instances and use iBGP within a domain and eBGP to the other domains

**Figure 11 Backup ABR Node SID**



The following steps describe the configuration and behavior of LFA Protection using Backup Node SID:

1. The user configures node SID 100 in ABR1 for its loopback prefix 1.1.1.1/32. This is the regular node SID. ABR1 advertises the prefix SID sub-TLV for this node SID in the IGP and installs the ILM using a unique label.
2. Each router receiving the prefix sub-TLV for node SID 100 resolves it as described in [Segment Routing in Shortest Path Forwarding](#). Changes to the programming of the backup NHLFE of node SID 100 based on receiving the backup node SID for prefix 1.1.1.1/32 are defined in [Duplicate SID Handling](#).
3. The user configures a backup node SID 200 in ABR2 for the loopback 1.1.1.1/32 of ABR1. The SID value must be different from that assigned by ABR1 for the same prefix. ABR2 installs the ILM, which performs a swap operation from the label of SID 200 to that of SID 100. The ILM must point to a direct link and next hop to reach 1.1.1.1/32 of ABR1 as its primary next hop. The IGP examines all adjacencies established in the same area as that of prefix 1.1.1.1/32 and determines which ones have ABR1 as a direct neighbor and with the best cost.

If more than one adjacency has the best cost, the IGP selects the one with the lowest interface index. If there is no adjacency to reach ABR2, the prefix SID for the backup node is flushed and is not resolved. This is to prevent any other non-direct path being used to reach ABR1. As a result, any received traffic on the ILM of SID 200 traffic will be blackholed.

4. If resolved, ABR2 advertises the prefix SID sub-TLV for this backup node SID 200 and indicates in the SR Algorithm field that a modified SPF algorithm, referred to as “Backup-constrained-SPF”, is required to resolve this node SID.
5. Each router receiving the prefix sub-TLV for the backup node SID 200 performs the following steps.

The following resolution steps do not require a CLI command to be enabled.

- i. The router determines which router is being backed up. This is achieved by checking the router ID owner of the prefix sub-TLV that was advertised with the same prefix but without the backup flag and which is used as the best route for the prefix. In this case, it should be ABR1. Then the router runs a modified SPF by removing node ABR1 from the topology to resolve the backup node SID 200. The primary next hop should point to the path to ABR2 in the counter clockwise direction of the ring.  
The router will not compute an LFA or a remote LFA for node SID 200 because the main SPF used a modified topology.
  - ii. The router installs the ILM and primary NHLFE for the backup node SID.  
Only a swap label operation is configured by all routers for the backup node SID. There is no push operation, and no tunnel for the backup node SID is added into the TTM.
  - iii. The router programs the backup node SID as the LFA backup for the SR tunnel to node SID of 1.1.1.1/32 of ABR1. In other words, each router overrides the remote LFA backup for prefix 1.1.1.1/32, which is normally PQ node AGN5.
  - iv. If the router is adjacent to ABR1, for example AGN1, it also programs the backup node SID as the LFA backup for the protection of any adjacency SID to ABR1.
6. When node AGN2 resolves a BGP label route for an inter-area prefix for which the primary ABR exit router is ABR1, it will use the backup node SID of ABR1 as the remote LFA backup instead of the SID to the PQ node (AGN5 in this example) to save on the pushed label stack.

AGN2 continues to resolve the prefix SID for any remote PE prefix that is summarized into the local area of AGN2 as usual. AGN2 programs a primary next hop and a remote LFA next hop. Remote LFA will use AGN5 as the PQ node and will push two labels, as it would for an intra-area prefix SID. There is no need to use the backup node SID for this prefix SID and force its backup path to go to ABR1. The backup path may exit from ABR2 if the cost from ABR2 to the destination prefix is shorter.

7. If the user excludes a link from LFA in the IGP instance (**config>router>ospf>area>interface>loopfree-alternate-exclude** or **config>router>isis>interface>loopfree-alternate-exclude**), a backup node SID that resolves to that interface will not be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.
8. If the OSPF neighbor of a router is put into overload or if the metric of an OSPF interface to that neighbor is set to LSInfinity (0xFFFF), a backup node SID that resolves to that neighbor will not be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.
9. If the IS-IS neighbor of a router is put into overload or if the metric of an IS-IS interface to that neighbor is set to overload max-metric (0xfffffe), a backup node SID that resolves to that neighbor will be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.



**Note:** Other routers in the network will not forward transit traffic to the router in overload.

10. If the IS-IS interface to a neighbor is set to maximum link metric (0xfffff), a backup node SID that resolves to that neighbor will not be used as a remote LFA backup in the same way as regular LFA or PQ remote LFA next hop behavior.
11. LFA policy is supported for IP next hops only. It is not supported with tunnel next hops such as IGP shortcuts or remote LFA tunnels. A backup node SID is also a tunnel next hop and therefore a user-configured LFA policy will not be applied to check constraints such as admin-groups and SRLG against the outgoing interface of the selected backup node SID.

### 3.3.3 Duplicate SID Handling

When the IGP issues or receives an LSA/LSP containing a prefix SID sub-TLV for a node SID or a backup node SID with a SID value that is a duplicate of an existing SID or backup node SID, the resolution in [Table 10](#) is followed.

**Table 10** Handling of Duplicate SIDs

| Old LSA/LSP     | New LSA/LSP     |                       |          |                |
|-----------------|-----------------|-----------------------|----------|----------------|
|                 | Backup Node SID | Local Backup Node SID | Node SID | Local Node SID |
| Backup Node SID | Old             | New                   | New      | New            |



**Table 10** Handling of Duplicate SIDs (Continued)

|                       | New LSA/LSP     |                       |                        |                        |
|-----------------------|-----------------|-----------------------|------------------------|------------------------|
| Old LSA/LSP           | Backup Node SID | Local Backup Node SID | Node SID               | Local Node SID         |
| Local Backup Node SID | Old             | Equal                 | New                    | New                    |
| Node SID              | Old             | Old                   | Equal/Old <sup>1</sup> | Equal/New <sup>2</sup> |
| Local Node SID        | Old             | Old                   | Equal/Old <sup>1</sup> | Equal/Old <sup>1</sup> |

**Notes:**

1. Equal/Old means the following.
  - If the prefix is duplicate, it is equal and no change is needed. Keep the old LSA/LSP.
  - If the prefix is not duplicate, still keep the old LSA/LSP.
2. Equal/New means the following.
  - If the prefix is duplicate, it is equal and no change is needed. Keep the old LSA/LSP.
  - If the prefix is not duplicate, pick a new prefix and use the new LSA/LSP.

### 3.3.4 OSPF Control Plane Extensions

All routers supporting OSPF control plane extensions must advertise support of the new algorithm “Backup-constrained-SPF” of value 2 in the SR-Algorithm TLV, which is advertised in the Router Information Opaque LSA. This is in addition to the default supported algorithm “IGP-metric-based-SPF” of value 0. The following shows the encoding of the prefix SID sub-TLV to indicate a node SID of type backup and to indicate the modified SPF algorithm in the SR Algorithm field. The values used in the Flags field and in the Algorithm field are SR OS proprietary.

The new Algorithm (0x2) field and values are used by this feature.




Table 11 lists OSPF control plane extension flag values.

**Table 11 OSPF Control Plane Extension Fields**

| Field  | Value         |
|--------|---------------|
| Type   | 2             |
| Length | variable      |
| Flags  | 1 octet field |

The following flags are defined; the “B” flag is new:

```

 0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+
| NP | M | E | V | L | B |
+---+---+---+---+---+---+---+

```

[Table 12](#) describes OSPF control plane extension flags.

**Table 12 OSPF Control Plane Extension Flags**

| Flag    | Description                                                                                                                                                                                             |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NP-Flag | No-PHP flag<br>If set, the penultimate hop must not pop the prefix SID before delivering the packet to the node that advertised the prefix SID                                                          |
| M-Flag  | Mapping Server Flag<br>If set, the SID is advertised from the Segment Routing Mapping Server functionality as described in I-D.filsfils-spring-segment-routing-ldp-interop                              |
| E-Flag  | Explicit-Null Flag<br>If set, any upstream neighbor of the prefix SID originator must replace the prefix SID with a prefix SID having an Explicit-NULL value (0 for IPv4) before forwarding the packet. |
| V-Flag  | Value/Index Flag<br>If set, the prefix SID carries an absolute value. If not set, the prefix SID carries an index.                                                                                      |
| L-Flag  | Local/Global Flag<br>If set, the value/index carried by the prefix SID has local significance. If not set, then the value/index carried by this sub-TLV has global significance.                        |
| B-Flag  | This flag is used by the Protection using backup node SID feature. If set, the SID is a backup SID for the prefix. This value is SR OS proprietary.                                                     |

**Table 12 OSPF Control Plane Extension Flags (Continued)**

| Flag            | Description                                                                                                                                                                                                                                                                         |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Other bits      | Reserved<br>These must be zero when sent and are ignored when received                                                                                                                                                                                                              |
| MT-ID           | Multi-Topology ID, as defined in RFC 4915                                                                                                                                                                                                                                           |
| Algorithm       | One octet identifying the algorithm the prefix SID is associated with. A value of (0x2) indicates the modified SPF algorithm, which removes from the topology the node that is backed up by the backup node SID. This value is SR OS proprietary.                                   |
| SID/Index/Label | Based on the V and L flags, it contains either: <ul style="list-style-type: none"> <li>• a 32-bit index defining the offset in the SID/Label space advertised by this router</li> <li>• a 24-bit label where the 20 rightmost bits are used for encoding the label value</li> </ul> |

## 3.4 Segment Routing in Shortest Path Forwarding

OSPF can be configured for segment routing in shortest path forwarding using the same procedures as those used to configure IS-IS. See [Segment Routing in Shortest Path Forwarding](#) in the IS-IS chapter for more information.

This section provides details of segment routing that are specific to the OSPFv2 and OSPFv3 control plane protocols.

### 3.4.1 OSPF Control Protocol Changes

New TLV/sub-TLVs are defined in *draft-ietf-ospf-segment-routing-extensions-04* and are required for the implementation of segment routing in OSPF. Specifically:

- the prefix SID sub-TLV part of the OSPFv2 Extended Prefix TLV
- the prefix SID sub-TLV part of the OSPFv2 Extended Prefix Range TLV
- the adjacency SID sub-TLV part of the OSPFv2 Extended Link TLV
- SID/Label Range Capability TLV
- SR-Algorithm Capability TLV

This section describes the behaviors and limitations of the OSPF support of segment routing TLV and sub-TLVs.

SR OS originates a single prefix SID sub-TLV per OSPFv2 Extended Prefix TLV and processes the first one only if multiple prefix SID sub-TLVs are received within the same OSPFv2 Extended Prefix TLV.

SR OS encodes the 32-bit index in the prefix SID sub-TLV. The 24-bit label or variable IPv6 SID is not supported.

SR OS originates a prefix SID sub-TLV with the following encoding of the flags.

- The NP-Flag is always set, meaning that the label for the prefix SID is pushed by the PHP router when forwarding to this router. SR OS PHP routers properly processes a received prefix SID with the NP-flag set to zero and uses implicit-null for the outgoing label towards the router which advertised it.
- The M-Flag is always unset because SR OS does not support originating a mapping server prefix-SID sub-TLV.

- The E-flag is always set to zero. SR OS PHP routers properly processes a received prefix SID with the E-flag set to 1, and when the NP-flag is also set to 1 it pushes explicit-null for the outgoing label towards the router which advertised it.
- The V-flag is always set to 0 to indicate an index value for the SID.
- The L-flag is always set to 0 to indicate that the SID index value is not locally significant.
- The algorithm field is always set to zero to indicate Shortest Path First (SPF) algorithm based on link metric and is not checked on a received prefix SID sub-TLV.

SR OS resolves a prefix SID received within an Extended Prefix TLV based on the following route preference:

- SID received via an intra-area route in a prefix SID sub-TLV part of Extended Prefix TLV
- SID received via an inter-area route in a prefix SID sub-TLV part of Extended Prefix TLV

SR OS originates an adjacency SID sub-TLV with the following encoding of the flags.

- The F-flag is unset to indicate the Adjacency SID refers to an adjacency with outgoing IPv4 encapsulation.
- The B-flag is set to zero and is not processed on receipt.
- The V-flag is always set.
- The L-flag is always set.
- The S-flag is not supported.
- The weight octet is not supported and is set to all zeros.

An adjacency SID is assigned to next hops over both the primary and secondary interfaces.

SR OS can originate the OSPFv2 Extended Prefix Range TLV as part of the Mapping Server feature and can process it properly if received. The following rules and limitations should be considered.

- Only the prefix SID sub-TLV within the TLV is processed and the ILMs installed if the prefixes are resolved.
- The range and address prefix fields are processed. Each prefix is resolved separately.
- If the same prefix is advertised with both a prefix SID sub-TLV in a IP reachability TLV and a mapping server Prefix-SID sub-TLV, the resolution follows the following route preference:

- the SID received via an intra-area route in a prefix SID sub-TLV part of Extended Prefix TLV
- the SID received via an inter-area route in a prefix SID sub-TLV part of Extended Prefix TLV
- the SID received via an intra-area route in a prefix SID sub-TLV part of a OSPFv2 Extended Range Prefix TLV
- the SID received via an inter-area route in a prefix SID sub-TLV part of a OSPFv2 Extended Range Prefix TLV
- No leaking of the entire TLV is performed between areas. Also, an ABR will not propagate the prefix-SID sub-TLV from the Extended Prefix Range TLV (received from a mapping server) into an Extended Prefix TLV if the latter is propagated between areas.
- The mapping server which advertised the OSPFv2 Extended Prefix Range TLV does not need to be in the shortest path for the FEC prefix.
- If the same FEC prefix is advertised in multiple OSPFv2 Extended Prefix Range TLVs by different routers, the SID in the TLV of the first router which is reachable is used. If that router becomes unreachable, the next reachable one is used.
- No check is performed to determine whether or not the contents of the OSPFv2 Extended Prefix Range TLVs received from different mapping servers are consistent.
- Any other sub-TLV, for example, the ERO metric and unnumbered interface ID ERO, is ignored but the user can get a dump of the octets of the received but not-supported sub-TLVs using the existing IGP **show** command.

SR OS supports propagation on ABR of external prefix LSA into other areas with routeType set to 3 as per *draft-ietf-ospf-segment-routing-extensions-04*.

SR OS supports propagation on ABR of external prefix LSA with route type 7 from NSSA area into other areas with route type set to 5 as per *draft-ietf-ospf-segment-routing-extensions-04*. SR OS does not support propagating of the prefix SID sub-TLV between OSPF instances.

When the user configures an OSPF import policy, the outcome of the policy applies to prefixes resolved in RTM and the corresponding tunnels in TTM. So, a prefix removed by the policy will not appear as both a route in RTM and as an SR tunnel in TTM.

---

## 3.4.2 Entropy Label for OSPF Segment Routing

The router supports the MPLS entropy label, as specified in RFC 6790, on OSPF segment-routed tunnels. LSR nodes in a network can load-balance labeled packets in a more granular way than by hashing on the standard label stack. Refer to the *MPLS Guide* for more information.

## 3.4.3 IPv6 Segment Routing using MPLS Encapsulation in OSPFv3

This feature implements support for SR IPv6 tunnels in OSPFv3 instances 0 to 31. The user can configure a node SID for the primary IPv6 global address of a loopback interface, which then gets advertised in OSPFv3. OSPFv3 automatically assigns and advertises an adjacency SID for each adjacency with an IPv6 neighbor. After the node SID is resolved, it is used to install an IPv6 SR-OSPF3 tunnel in the TTMv6 for use by the routes and services.

### 3.4.3.1 OSPFv3 Extensions

The OSPFv3 extensions consist of support for the following TLVs:

- a prefix SID that is a sub-TLV of the OSPFv3 prefix TLV. The OSPFv3 prefix TLV is a new top-level TLV of the extended prefix LSA introduced in *draft-ietf-ospf-ospfv3-lsa-extend*. The OSPFv3 instance can operate in either LSA sparse mode or extended LSA mode.

The **config>router>extended-lsa only** command advertises the prefix SID sub-TLV in the extended LSA format in both cases.

- an adjacency SID that is a sub-TLV of the OSPFv3 router-link TLV. The OSPFv3 router-link TLV is a new top-level TLV in the extended router LSA introduced in *draft-ietf-ospf-ospfv3-lsa-extend*. The OSPFv3 instance can operate in either LSA sparse mode or extended LSA mode. The **config>router>extended-lsa only** command advertises the adjacency SID sub-TLV in the extended LSA format in both cases.
- the SR-Algorithm TLV and the SID/Label Range TLV. Both of these TLVs are part of the TLV-based OSPFv3 Router Information opaque LSA defined in RFC 7770.

## 3.5 Segment Routing Mapping Server

The mapping server feature allows for the configuration and advertisement, in OSPF, of the node SID index for prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server and using a prefix-SID sub-TLV within an OSPF Extended Prefix Range TLV.

The user configures the SR mapping database in OSPF using the following CLI command:

```
CLI Syntax: configure
 router
 [no] ospf
 segment-routing
 no segment-routing
 mapping-server
 sid-map node-sid {index value
 <0..4294967295> [range value
 <1..65535>]} prefix {{ip-
 address/mask} | {netmask}}
 [scope {area area-id | as}]
 no sid-map node-sid index value
```

The user enters the node SID index, for one prefix or a range of prefixes, by specifying the first index value and, optionally, a range value. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. If the user enters the first prefix with a mask lower than 32, the OSPF Extended Prefix Range TLV is advertised but a router which receives it will not resolve SID and instead originates a trap.

The **no** form of the **sid-map** command deletes the range of node SIDs beginning with the specified index value. The **no** form of the **mapping-server** command deletes all node SID entries in the OSPF instance.

The user specifies the mapping server's own flooding scope for the generated OSPF Extended Prefix Range TLV using the **scope** option. There is no default value. If the scope is a specific area, then the TLV is flooded only in that area.

An ABR that propagates an intra-area OSPF Extended Prefix Range TLV flooded by the mapping server in that area into other areas sets the inter-area flag (IA-flag). The ABR also propagates the TLV if received with the inter-area flag set from other ABR nodes but only from the backbone to leaf areas and not vice-versa. However, if the exact same TLV was advertised as an intra-area TLV in a leaf area, the ABR will not flood the inter-area TLV into that leaf area.





**Note:** SR OS does not leak the OSPF Extended Prefix Range TLV between OSPF instances.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix, or range of prefixes, a prefix-SID sub-TLV within a OSPF Extended Prefix Range TLV in that instance. The flooding scope of the TLV from the mapping server is determined as previously explained. No further check of the reachability of that prefix in the mapping server route table is performed and no check if the SID index is duplicate with some existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

### 3.5.1 Segment Routing Mapping Server Prefix SID Resolution

The rules for IP prefix resolution, prefix SID resolution, and SR tunnel programming are as follows:

- IP Prefix Resolution
  - a. SPF calculates the next hops, up to **max-ecmp**, to reach a destination node
  - b. Each prefix inherits the next hops of one or more destination nodes advertising it
  - c. A prefix advertised by multiple nodes, all reachable with same cost, will inherit up to max-ecmp next hops from the advertising nodes
  - d. Next hop selection, up to **max-ecmp** value, is deterministic and is based on sorting the next hops by:
    - i. lowest next hop router-id
    - ii. lowest interface index, for parallel links to same router-id
 Each next hop keeps a reference to the destination nodes of whom it was inherited
- Prefix SID Resolution
  - a. For a given prefix, IGP selects the SID value among multiple advertised values respecting the following preference order:
    - i. local intra-area SID owned by this router
    - ii. prefix SID sub-TLV advertised within a OSPF Extended Prefix TLV

If multiple SIDs exist in this step, select the SID corresponding to the destination router or ABR with the lowest OSPF Router-ID which is reachable via the first next hop of the prefix

iii. OSPF Extended Prefix Range TLV from mapping server

If multiple SIDs exist in this step, select the following, using the preference rules in *draft-ietf-spring-conflict-resolution-05* when applied to the SRMS entries of the conflicting SIDs. The order of these rules is as follows:

- smallest range
- smallest starting address
- smallest algorithm
- smallest starting SID

- b. The selected SID is used with all ECMP next hops from step (I) towards all destination nodes or ABR nodes which advertised the prefix.
- c. If duplicate prefix SIDs exist for different prefixes after above steps, the first SID which is processed is programmed for its corresponding prefix.

Subsequent SIDs causes a duplicate SID trap and are not programmed. The corresponding prefixes themselves are still resolved normally using IP next-next-hops.

• SR Tunnel Programming

- a. If the prefix SID is resolved from a prefix SID sub-TLV advertised within a OSPF Extended Prefix TLV, the SR ILM label is swapped to a SR NHLFE label like in regular SR tunnel resolution.
- b. If the prefix SID is resolved from a mapping server advertisement:
- i. The SR ILM is stitched to an LDP FEC of the same prefix if one exists. The LDP FEC could have been resolved using a static route, a route within a IS-IS instance, or a route within a OSPF instance. The latter two can be the same or different as the IGP instance which advertised the mapping server prefix SID sub-TLV.
  - ii. Otherwise the SR ILM label is swapped to a SR NHLFE label. This is only possible if a route is exported from another IGP instance into the local IGP instance without propagating the prefix SID sub-TLV with the route. Otherwise, when the prefix SID is propagated with the route, the resolution follows step ([SR Tunnel Programming](#), step a).

---

## 3.6 SPF LSA Filtering

The SR OS OSPF implementation supports a configuration option to filter outgoing OSPF LSAs on selected OSPFv2 or OSPFv3 interfaces. This feature should be used with some caution because it goes against the principle that all OSPF routers in an area should have a synchronized Link State Database (LSDB), but it can be a useful resource saving in certain hub and spoke topologies where learning routes through OSPF is only needed in one direction (for example, from spoke to hub).

Three filtering options are available (configurable per interface):

- Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next hop.
- Flood a minimum set of self-generated LSAs out the interface (e.g. router-LSA, intra-area-prefix-LSA, and link-LSA and network-LSA corresponding to the connected interface); suppress all non-self-originated LSAs. This option is suitable if the neighbor is simply connected and has a statically configured default route with a loopback or system interface address as next hop
- Flood a minimum set of self-generated LSAs (for example, router-LSA, intra-area-prefix-LSA, and link-LSA and network-LSA corresponding to the connected interface) and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route.

---

## 3.7 FIB Prioritization

The RIB processing of specific routes can be prioritized through the use of the `rib-priority` command. This command allows specific routes to be prioritized through the protocol processing so that updates are propagated to the FIB as quickly as possible.

Configuring the **`rib-priority`** command either within the global OSPF or OSPFv3 routing context or under a specific OSPF/OSPFv3 interface context enables this feature. Under the global OSPF context, a prefix list can be specified that identifies which route prefixes should be considered high priority. If the **`rib-priority high`** command is configured under an OSPF interface context then all routes learned through that interface is considered high priority.

The routes that have been designated as high priority will be the first routes processed and then passed to the FIB update process so that the forwarding engine can be updated. All known high priority routes should be processed before the OSPF routing protocol moves on to other standard priority routes. This feature will have the most impact when there are a large number of routes being learned through the OSPF routing protocols.

---

## 3.8 Extended LSA Support in OSPFv3

This feature adds support for the extended LSA format in OSPFv3 as per draft-ietf-ospf-ospfv3-lsa-extend, *OSPFv3 LSA Extendibility*.

Prior to this feature, the SR OS used the fixed-format LSA to carry the prefix and link information as per RFC 5340, *OSPF for IPv6*. The fixed-format is not extensible and hence the need to use the TLV format of the extended LSA.

With this feature, the default mode of operation for OSPFv3 is referred to as **sparse** mode, meaning that the router will always advertise the fixed-format for existing LSAs and will add the TLV-based extended LSA only when it needs to advertise new sub-TLVs. This mode of operation is very similar to the way OSPFv2 advertises the Segment Routing information. It sends the prefix in the original fixed-format prefix LSA and then follows with the extended prefix TLV which is sent in an extended prefix opaque LSA containing the prefix SID sub-TLV.

The **extended-lsa** only value enables the full extended LSA mode and this will cause all existing and new LSAs to use the extended LSA format.

An OSPFv3 area inherits the instance level configuration but can also be configured independently to the sparse or full extended LSA mode.

The OSPFv3 instance must first be shut down before the user can change the mode of operation since the protocol must flush all LSAs and re-establish all adjacencies

---

## 3.9 Support of Multiple Instances of Router Information LSA in OSPFv2 and OSPFv3

This feature adds the support of multiple instances of the Router Information LSA as per RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*.

The original method of advertising router capabilities used options fields in LSAs and hello packets. However, this method is not extensible due to the limited size of the options field. Thus, RFC 4970, *Extensions to OSPF for Advertising Optional Router Capabilities*, defined the Router Information LSA which can carry multiple router capability TLVs. It also defined a single TLV called the Router Information Capabilities TLV to carry all previously defined capabilities in the options field in LSAs and hello packets. The SR OS supports RFC 4970.

RFC 7770 deprecated RFC 4970 by adding the ability to send multiple instances of the Router Information LSA to circumvent the maximum LSA size of 64 kbytes.


There is no CLI to enable the support of multiple instances of the Router Information LSA. The existing Router Information Capabilities TLVs is carried as the first TLV (Opaque ID 0) of the first instance (instance ID 0) of the Router Information LSA. The existing router information TLVs, such as the OSPFv2 SR-Algorithm TLV and the SID/Label Range TLV, are sent in the first instance of the Router Information LSA.

If a router information TLV is received in multiple instances of the Router Information LSA, the default behavior is to process the one in the lowest instance ID and ignore the other ones.

## 3.10 OSPF Configuration Process Overview

Figure 12 displays the process to provision basic OSPF parameters.

**Figure 12** OSPF Configuration and Implementation Flow



## 3.11 Configuration Notes

This section describes OSPF configuration caveats.

### 3.11.1 General

- Before OSPF can be configured, the router ID must be configured.
- The basic OSPF configuration includes at least one area and an associated interface.
- All default and command parameters can be modified.

#### 3.11.1.1 OSPF Defaults

The following list summarizes the OSPF configuration defaults:

- By default, a router has no configured areas.
- An OSPF instance is created in the administratively enabled state.



---

## 3.12 Configuring OSPF with CLI

This section provides information to configure Open Shortest Path First (OSPF) using the command line interface.

### 3.12.1 OSPF Configuration Guidelines

Configuration planning is essential to organize routers, backbone, non-backbone, stub, NSSA areas, and transit links. OSPF provides essential defaults for basic protocol operability. You can configure or modify commands and parameters. OSPF is not enabled by default.

The minimal OSPF parameters which should be configured to deploy OSPF are:

- Router ID

Each router running OSPF must be configured with a unique router ID. The router ID is used by both OSPF and BGP routing protocols in the routing table manager.

When configuring a new router ID, protocols will not automatically be restarted with the new router ID. Shut down and restart the protocol to initialize the new router ID.

- OSPF instance

OSPF instances must be defined when configuring multiple instances and/or the instance being configured is not the base instance.

- An area

At least one OSPF area must be created. An interface must be assigned to each OSPF area.

- Interfaces

An interface is the connection between a router and one of its attached networks. An interface has state information associated with it, which is obtained from the underlying lower level protocols and the routing protocol itself. An interface to a network has associated with it a single IP address and mask (unless the network is an unnumbered point-to-point network). An interface is sometimes also referred to as a link.

### 3.12.2 Basic OSPF Configurations

This section provides information to configure OSPF and OSPF3 as well as configuration examples of common configuration tasks.

The minimal OSPF parameters that need to be configured are:

- A router ID - If a *router-id* is not configured in the **config>router** context, the router's system interface IP address is used.
- One or more areas.
- Interfaces (interface "system").

The following is an example of a basic OSPF configuration:

```
ALA-A>config>router>ospf# info

 area 0.0.0.0
 interface "system"
 exit
 exit
 area 0.0.0.20
 nssa
 exit
 interface "to-104"
 priority 10
 exit
 exit
 area 0.0.1.1
 exit

ALA-A>config>router>ospf#
```

The following is an example of a basic OPSF3 configuration:

```
A:ALA-48>config>router>ospf3# info

 asbr
 overload
 timers
 lsa-arrival 50000
 exit
 export "OSPF-Export"
 area 0.0.0.0
 interface "system"
 exit
 exit
 area 0.0.0.20
 nssa
 exit
 interface "SR1-2"
 exit
 exit

```

```
area 0.0.0.25
 stub
 default-metric 5000
 exit
exit
```

### 3.12.2.1 Configuring the Router ID

The router ID uniquely identifies the router within an AS. In OSPF, routing information is exchanged between autonomous systems, groups of networks that share routing information. It can be set to be the same as the loopback (system interface) address. Subscriber services also use this address as far-end router identifiers when service distribution paths (SDPs) are created. The router ID is used by both OSPF and BGP routing protocols. A router ID can be derived by:

- Defining the value in the **config>router** router-id context.
- Defining the system interface in the **config>router>interface** *ip-int-name* context (used if the router ID is not specified in the **config>router** *router-id* context).
- Inheriting the last four bytes of the MAC address.
- On the BGP protocol level, a BGP router ID can be defined in the **config>router>bgp** *router-id* context and is only used within BGP.
- Defining a router ID when creating an OSPF instance **config>router>ospf** [*instance-id*] [*router-id*]

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the shutdown and no shutdown commands for each protocol that uses the router ID or restart the entire router.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

The following displays a router ID configuration example:

```
A:ALA-B>config>router# info
#-----
IP Configuration
#-----
 interface "system"
 address 10.10.10.104/32
 exit
 interface "to-103"
```

```
 address 10.0.0.104/24
 port 1/1/1
 exit
 autonomous-system 100
 router-id 10.10.10.104
...
#-----
A:ALA-B>config>router#
```

### 3.12.3 Configuring OSPF Components

Use the CLI syntax displayed in the following subsections to configure OSPF components.

#### 3.12.3.1 Configuring OSPF Parameters

The following displays a basic OSPF configuration example:

```
A:ALA-49>config>router>ospf# info

 asbr
 overload
 overload-on-boot timeout 60
 traffic-engineering
 export "OSPF-Export"
 exit

A:ALA-49>config>router>ospf# ex
```

#### 3.12.3.2 Configuring OSPF3 Parameters

Use the following CLI syntax to configure OSPF3 parameters:

**CLI Syntax:**    config>router# ospf3

```
 asbr
 export policy-name [policy-name...(up to 5 max)]
 external-db-overflow limit seconds
 external-preference preference
 overload [timeout seconds]
 overload-include-stub
 overload-on-boot [timeout seconds]
 preference preference
 reference-bandwidth bandwidth-in-kbps
 router-id ip-address
```

```
no shutdown
timers
 lsa-arrival lsa-arrival-time
 lsa-generate max-lsa-wait [lsa-initial-wait
 lsa-initial-wait [lsa-second-wait lsa-
 second-wait]]
 spf-wait max-spf-wait [spf-initial-wait spf-
 initial-wait [spf-second-wait spf-second-
 wait]]
```

The following displays an OSPF3 configuration example:

```
A:ALA-48>config>router>ospf3# info

 asbr
 overload
 timers
 lsa-arrival 50000
 exit
 export "OSPF-Export"

A:ALA-48>config>router>ospf3#
```

OSPF also supports the concept of multi-instance OSPFv2 and OSPFv3 which allows separate instances of the OSPF protocols to run independently within SR OSs.

Separate instances are created by adding a different instance ID as the optional parameter to the **config>router>ospf** and **config>router>ospf3** commands. When this is done a separate OSPF instance is created which maintains separate link state databases for each instance.

### 3.12.3.3 Configuring an OSPF or OSPF3 Area

An OSPF area consists of routers configured with the same area ID. To include a router in a specific area, the common area ID must be assigned and an interface identified.

If your network consists of multiple areas you must also configure a backbone area (0.0.0.0) on at least one router. The backbone is comprised of the area border routers and other routers not included in other areas. The backbone distributes routing information between areas. The backbone is considered to be a participating area within the autonomous system. To maintain backbone connectivity, there must be at least one interface in the backbone area or have a virtual link configured to another router in the backbone area.

The minimal configuration must include an area ID and an interface. Modifying other command parameters are optional.

Use the following CLI syntax to configure an OSPF area:

**CLI Syntax:**

```
ospf ospf-instance
area area-id
 area-range ip-prefix/mask [advertise | not-
 advertise]
 blackhole-aggregate
```

Use the following CLI syntax to configure an OSPF3 area:

**CLI Syntax:**

```
ospf ospf-instance
ospf3
area area-id
 area-range ip-prefix/mask [advertise | not-
 advertise]
 blackhole-aggregate
```

The following displays an OSPF area configuration example:

```
A:ALA-A>config>router>ospf# info

 area 0.0.0.0
 exit
 area 0.0.0.20
 exit

ALA-A>config>router>ospf#A:
```

### 3.12.3.4 Configuring a Stub Area

Configure stub areas to control external advertisements flooding and to minimize the size of the topological databases on an area's routers. A stub area cannot also be configured as an NSSA.

By default, summary route advertisements are sent into stub areas. The **no** form of the summary command disables sending summary route advertisements and only the default route is advertised by the ABR. This example retains the default so the command is not entered.

If this area is configured as a transit area for a virtual link, then existing virtual links of a non-stub or NSSA area are removed when its designation is changed to NSSA or stub.

Stub areas for OSPF3 are configured the same as OSPF stub areas.

Use the following CLI syntax to configure virtual links:

**CLI Syntax:**

```
ospf
 area area-id
 stub
 default-metric metric
 summaries
```

Use the following CLI syntax to configure virtual links for OSPF3:

**CLI Syntax:**

```
ospf
ospf3
 area area-id
 stub
 default-metric metric
 summaries
```

The following displays a stub configuration example:

```
ALA-A>config>router>ospf>area># info

...
 area 0.0.0.0
 exit
 area 0.0.0.20
 stub
 exit
 exit
...

ALA-A>config>router>ospf#
```

The following displays a stub configuration example:

```
ALA-A>config>router>ospf>area># info

...
 area 0.0.0.0
 exit
 area 0.0.0.20
 stub
 exit
 exit
...

ALA-A>config>router>ospf#
```

The following displays a stub configuration example for the OSPF3:

```
A:ALA-48>config>router>ospf3>area# info

 stub
 default-metric 5000
```

```

 exit

A:ALA-48>config>router>ospf3>area#

```

### 3.12.3.5 Configuring a Not-So-Stubby Area

You must explicitly configure an area to be a Not-So-Stubby Area (NSSA) area. NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes it learns throughout its area and by an area border router to the entire OSPF domain. An area cannot be both a stub area and an NSSA.

If this area is configured as a transit area for a virtual link, then existing virtual links of a non-stub or NSSA area are removed when its designation is changed to NSSA or stub.

Use the following CLI syntax to configure stub areas:

**CLI Syntax:**

```

ospf ospf-instance
 area area-id
 nssa
 area-range ip-prefix/mask [advertise|not-
 advertise]
 originate-default-route [type-7]
 redistribute-external
 summaries

```

Use the following CLI syntax to configure stub areas for the OSPF3:

**CLI Syntax:**

```

ospf ospf-instance
ospf3
 area area-id
 nssa
 area-range ip-prefix/mask [advertise|not-
 advertise]
 originate-default-route [type-7]
 redistribute-external
 summaries

```

The following displays an NSSA configuration example:

```

A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60

```



```

traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
exit
area 0.0.0.20
stub
exit
exit
area 0.0.0.25
nssa
exit
exit

```

-----  
A:ALA-49>config>router>ospf#

The following displays a OSPF3 NSSA configuration example:

A:ALA-48>config>router>ospf3# info

```

asbr
overload
timers
 lsa-arrival 50000
exit
export "OSPF-Export"
area 0.0.0.0
exit
area 0.0.0.20
stub
exit
exit
area 0.0.0.25
nssa
exit
exit
area 4.3.2.1
exit

```

-----  
A:ALA-48>config>router>ospf3#

### 3.12.3.6 Configuring a Virtual Link

The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone then the area border routers must be connected via a virtual link. The two area border routers will form a point-to-point-like adjacency across the transit area. A virtual link can only be configured while in the area 0.0.0.0 context.

The **router-id** parameter specified in the **virtual-link** command must be associated with the virtual neighbor, that is, enter the virtual neighbor's router ID, not the local router ID. The transit area cannot be a stub area or an NSSA.

Use the following CLI syntax to configure stub areas:

**CLI Syntax:**

```
ospf ospf-instance
 area area-id
 virtual-link router-id transit-area area-id
 authentication-key [authentication-key|hash-key] [hash]
 authentication-type [password|message-digest]
 dead-interval seconds
 hello-interval seconds
 message-digest-key key-id md5 [key|hash-key]
 [hash|hash2]
 retransmit-interval seconds
 transit-delay
 no shutdown
```

The following displays a virtual link configuration example:

```
A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
exit
area 0.0.0.20
 stub
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 1.2.3.4
 exit

A:ALA-49>config>router>ospf#
```

The following displays an OSPF3 virtual link configuration example:

```
A:ALA-48>config>router>ospf3# info
```

```

asbr
overload
timers
 lsa-arrival 50000
exit
export "OSPF-Export"
area 0.0.0.0
 virtual-link 4.3.2.1 transit-area 4.3.2.1
 exit
exit
area 0.0.0.20
 stub
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 4.3.2.1
exit

A:ALA-48>config>router>ospf3#
```

### 3.12.3.7 Configuring an Interface

In OSPF, an interface can be configured to act as a connection between a router and one of its attached networks. An interface includes state information that was obtained from underlying lower level protocols and from the routing protocol itself. An interface to a network is associated with a single IP address and mask (unless the network is an unnumbered point-to-point network). If the address is merely changed, then the OSPF configuration is preserved.

The passive command enables the passive property to and from the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol. By default, only interface addresses that are configured for OSPF are advertised as OSPF interfaces. The passive parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol. When enabled, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

An interface can be part of more than one area, as specified in RFC 5185. To do this, add the keyword **secondary** when creating the interface.

Use the following CLI syntax to configure an OSPF interface:

**CLI Syntax:**

```
ospf ospf-instance
 area area-id
 interface ip-int-name
 advertise-subnet
```

```

authentication-key [authentication-key|hash-
 key] [hash|hash2]
authentication-type [password|message-digest]
bfd-enable
dead-interval seconds
hello-interval seconds
interface-type {broadcast|point-to-point |non-
 broadcast}
message-digest-key key-id md5 [key|hash-
 key] [hash|hash2]
metric metric
mtu bytes
passive
priority number
retransmit-interval seconds
no shutdown
transit-delay seconds

```

The following displays an interface configuration example:

```
A:ALA-49>config>router>ospf# info
```

```

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 1.2.3.4
exit
area 4.3.2.1
 interface "SR1-3"
 exit
exit
area 4.3.2.1
 interface "SR1-3" secondary
 exit

```

```

exit

A:ALA-49>config>router>ospf# area 0.0.0.20

```

The following displays an interface configuration:

```

A:ALA-48>config>router>ospf3# info

asbr
overload
timers
 lsa-arrival 50000
exit
export "OSPF-Export"
area 0.0.0.0
 virtual-link 4.3.2.1 transit-area 4.3.2.1
 exit
 interface "system"
 exit
exit
area 0.0.0.20
 stub
 exit
 interface "SR1-2"
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 4.3.2.1
exit

A:ALA-48>config>router>ospf3#

```

### 3.12.3.8 Configuring Authentication

#### 3.12.3.8.1 Overview

The use of protocol authentication is recommended to protect against malicious attack on the communications between routing protocol neighbors. These attacks could aim to either disrupt communications or to inject incorrect routing information into the systems routing table. The use of authentication keys can help to protect the routing protocols from these types of attacks.

Authentication must be explicitly configured and can be done so through two separate mechanisms. First is configuration of an explicit authentication key and algorithm through the use of the authentication and authentication-type commands. The second method is through the use of the authentication keychain mechanism. Both mechanisms are described in the following sections.

### 3.12.3.8.2 Configuring Authentication Keys and Algorithms

The following authentication commands can be configured on the interface level or the virtual link level:

- **authentication-key** — Configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.
- **authentication-type** — Enables authentication and specifies the type of authentication to be used on the OSPF interface, either password or message digest.
- **message-digest-key** — Use this command when message-digest keyword is selected in the authentication-type command. The Message Digest 5 (MD5) hashing algorithm is used for authentication. MD5 is used to verify data integrity by creating a 128-bit message digest from the data input. It is unique to that specific data.

An special checksum is included in transmitted packets and are used by the far-end router to verify the packet by using an authentication key (a password). Routers on both ends must use the same MD5 key.

MD5 can be configured on each interface and each virtual link. If MD5 is enabled on an interface, then that interface accepts routing updates only if the MD5 authentication is accepted. Updates that are not authenticated are rejected. A router accepts only OSPF packets sent with the same key-id value defined for the interface.

When the hash parameter is not used, non-encrypted characters can be entered. Once configured using the **message-digest-key** command, then all keys specified in the command are stored in encrypted format in the configuration file using the **hash** keyword. When using the **hash** keyword the password must be entered in encrypted form. Hashing cannot be reversed. Issue the **no message-digest-key key-id** command, then re-enter the command *without* the **hash** parameter to configure an unhashed key.

The following CLI commands are displayed to illustrate the key authentication features. These command parameters can be defined at the same time interfaces and virtual-links are being configured. See [Configuring an Interface](#) and [Configuring a Virtual Link](#).

Use the following CLI syntax to configure authentication:

**CLI Syntax:**

```
ospf ospf-instance
 area area-id
 interface ip-int-name
 authentication-key [authentication-key|hash-
 key] [hash]
 authentication-type [password|message-digest]
 message-digest-key key-id md5 key [hash]
 virtual-link router-id transit-area area-id
 authentication-key [authentication-key|hash-
 key] [hash]
 authentication-type [password|message-digest]
 message-digest-key key-id md5 key [hash]
```

The following displays authentication configuration examples:

```
A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 0.0.0.40
 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
exit
area 1.2.3.4
exit

A:ALA-49>config>router>ospf#
A:ALA-49>config>router>ospf# info
```

```

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 10.0.0.1 transit-area 0.0.0.1
 authentication-type message-digest
 message-digest-key 2 md5 "Mi6BQAFi3MI" hash
 exit
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.1
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 0.0.0.40
 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
exit
area 1.2.3.4
exit

A:ALA-49>config>router>ospf#

```

### 3.12.3.8.3 Configuring Authentication using Keychains

The use of authentication mechanism is recommended to protect against malicious attack on the communications between routing protocol neighbors. These attacks could aim to either disrupt communications or to inject incorrect routing information into the systems routing table. The use of authentication keys can help to protect the routing protocols from these types of attacks. In addition, the use of authentication keychains provides the ability to configure authentication keys and make changes to them without affecting the state of the routing protocol adjacencies.



To configure the use of an authentication keychain within OSPF, use the following steps:

1. Configure an authentication keychain within the `config>system>security` context. The configured keychain must include at least one valid key entry, using a valid authentication algorithm for the OSPF protocol.
2. Associate the configured authentication keychain within OSPF. Authentication keychains can be used to specify the authentication key and algorithm on a per interface basis within the configuration for the OSPF protocol.

For a key entry to be valid, it must include a valid key, the current system clock value must be within the begin and end time of the key entry, and the algorithm specified in the key entry must be supported by the OSPF protocol.

The OSPF protocol supports the following algorithms:

- clear text password
- MD5
- HMAC-SHA-1-96
- HMAC-SHA-1
- HMAC-SHA-256

Keychain Error handling:

- If a keychain exists but there are no active key entries with an authentication type that is valid for the associated protocol, then inbound protocol packets will not be authenticated and discarded and no outbound protocol packets will be sent.
- If keychain exists, but the last key entry has expired, a log entry will be raised indicating that all keychain entries have expired. The OSPF protocol requires that the protocol continue to authenticate inbound and outbound traffic using the last valid authentication key.

### 3.12.3.9 Assigning a Designated Router

A designated router is elected according to the priority number advertised by the routers. When a router starts up, it checks for a current designated router. If a designated router is present, then the router accepts that designated router, regardless of its own priority designation. When a router fails, then new designated and backup routers are elected according to their priority numbers.

The **priority** command is only used if the interface is a broadcast type. The designated router is responsible for flooding network link advertisements on a broadcast network to describe the routers attached to the network. A router uses hello packets to advertise its priority. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be a designated router or a backup designated router. At least one router on each logical IP network or subnet must be eligible to be the designated router. By default, routers have a priority value of 1.

Use the following CLI syntax to configure the designated router:

**CLI Syntax:**

```
ospf ospf-instance
 area area-id
 interface ip-int-name
 priority number
```

The following displays a priority designation example:

```
A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 10.0.0.1 transit-area 0.0.0.1
 authentication-type message-digest
 message-digest-key 2 md5 "Mi6BQAFi3MI" hash
 exit
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.1
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
 interface "if2"
 priority 100
 exit
exit
area 0.0.0.40
```

```

 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
 exit
 area 1.2.3.4
 exit

A:ALA-49>config>router>ospf#

```

### 3.12.3.10 Configuring Route Summaries

Area border routers send summary (type 3) advertisements into a stub area or NSSA to describe the routes to other areas. This command is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or NSSA.

By default, summary route advertisements are sent into the stub area or NSSA. The no form of the summaries command disables sending summary route advertisements and, in stub areas, the default route is advertised by the area border router.

The following CLI commands are displayed to illustrate route summary features. These command parameters can be defined at the same time stub areas and NSSAs are being configured. See [Configuring a Stub Area](#) and [Configuring a Not-So-Stubby Area](#).

Use the following CLI syntax to configure a route summary:

```

CLI Syntax: ospf ospf-instance
 area area-id
 stub
 summaries
 nssa
 summaries

```

The following displays a stub route summary configuration example:

```

A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 10.0.0.1 transit-area 0.0.0.1
 authentication-type message-digest

```

```

 message-digest-key 2 md5 "Mi6BQAFi3MI" hash
 exit
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.1
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
 interface "if2"
 priority 100
 exit
exit
area 0.0.0.40
 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
exit
area 1.2.3.4
exit

```

```

A:ALA-49>config>router>ospf#

```

The following displays a stub route summary configuration example:

```

A:ALA-48>config>router>ospf3# info

```

```

asbr
overload
timers
 lsa-arrival 50000
exit
export "OSPF-Export"
area 0.0.0.0
 virtual-link 4.3.2.1 transit-area 4.3.2.1
 exit
 interface "system"
 exit
exit
area 0.0.0.20
 stub
 exit
 interface "SR1-2"
 exit
exit
area 0.0.0.25

```

```
 nssa
 exit
 exit
 area 4.3.2.1
 exit

A:ALA-48>config>router>ospf3#
```

### 3.12.3.11 Configuring Route Preferences

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference value is used to decide which route is installed in the forwarding table if several protocols calculate routes to the same destination. The route with the lowest preference value is selected

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in [Table 13](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

**Table 13** Route Preference Defaults by Route Type

| Route Type             | Preference | Configurable     |
|------------------------|------------|------------------|
| Direct attached        | 0          | No               |
| Static routes          | 5          | Yes              |
| OSPF internal          | 10         | Yes <sup>1</sup> |
| IS-IS level 1 internal | 15         | Yes              |
| IS-IS level 2 internal | 18         | Yes              |
| OSPF external          | 150        | Yes              |
| IS-IS level 1 external | 160        | Yes              |
| IS-IS level 2 external | 165        | Yes              |
| BGP                    | 170        | Yes              |

Note:

1. Preference for OSPF internal routes is configured with the preference command.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

The following CLI commands are displayed to illustrate route preference features. The command parameters can be defined at the same time you are configuring OSPF. See [Configuring OSPF Components](#).

Use the following CLI syntax to configure a route preference:

**CLI Syntax:**    *ospf ospf-instance*  
                          *preference preference*  
                          *external-preference preference*

Use the following CLI syntax to configure a route preference for the OSPF3:

**CLI Syntax:**    *ospf ospf-instance*  
                  *ospf3*  
                          *preference preference*  
                          *external-preference preference*

The following displays a route preference configuration example:

```
A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
preference 9
external-preference 140
export "OSPF-Export"
exit
area 0.0.0.0
 virtual-link 10.0.0.1 transit-area 0.0.0.1
 authentication-type message-digest
 message-digest-key 2 md5 "Mi6BQAFi3MI" hash
 exit
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
exit
area 0.0.0.1
exit
area 0.0.0.20
 stub
 exit
 interface "to-103"
 exit
exit
area 0.0.0.25
 nssa
 exit
 interface "if2"
 priority 100
```

```

 exit
 exit
 area 0.0.0.40
 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
 exit
 area 1.2.3.4
 exit

```

-----

The following displays a route preference configuration example:

A:ALA-48>config>router>ospf3# info

```

asbr
overload
timers
 lsa-arrival 50000
exit
preference 9
external-preference 140
export "OSPF-Export"
area 0.0.0.0
 virtual-link 4.3.2.1 transit-area 4.3.2.1
 exit
 interface "system"
 exit
exit
area 0.0.0.20
 stub
 exit
 interface "SR1-2"
 exit
exit
area 0.0.0.25
 nssa
 exit
exit
area 4.3.2.1
exit

```

-----

A:ALA-48>config>router>ospf3#

## 3.13 OSPF Configuration Management Tasks

This section discusses OSPF configuration management tasks.

### 3.13.1 Modifying a Router ID

Since the router ID is defined in the `config>router` context, not in the OSPF configuration context, the protocol instance is not aware of the change. Re-examine the plan detailing the router ID. Changing the router ID on a device could cause configuration inconsistencies if associated values are not also modified.

After you have changed a router ID, manually shut down and restart the protocol using the `shutdown` and `no shutdown` commands in order for the changes to be incorporated.

Use the following CLI syntax to change a router ID number:

**CLI Syntax:** `config>router# router-id router-id`

The following displays a NSSA router ID modification example:

```
A:ALA-49>config>router# info

IP Configuration

 interface "system"
 address 10.10.10.104/32
 exit
 interface "to-103"
 address 10.0.0.103/24
 port 1/1/1
 exit
 autonomous-system 100
 router-id 10.10.10.104

A:ALA-49>config>router#

ALA-48>config>router# info

IP Configuration

 interface "system"
 address 10.10.10.103/32
 exit
 interface "to-104"
 address 10.0.0.104/24
 port 1/1/1
 exit
```



```

autonomous-system 100
router-id 10.10.10.103

ALA-48>config>router#

```

### 3.13.2 Deleting a Router ID

You can modify a router ID, but you cannot delete the parameter. When the **no router *router-id*** command is issued, the router ID reverts to the default value, the system interface address (which is also the loopback address). If a system interface address is not configured, then the last 32 bits of the chassis MAC address is used as the router ID.

### 3.13.3 Modifying OSPF Parameters

You can change or remove existing OSPF parameters in the CLI or NMS. The changes are applied immediately.

The following example displays an OSPF modification in which an interface is removed and another interface added.

**Example:**

```

config>router# ospf 1
config>router>ospf# area 0.0.0.20
config>router>ospf>area# no interface "to-103"
config>router>ospf>area# interface "to-HQ"
config>router>ospf>area>if$ priority 50
config>router>ospf>area>if# exit
config>router>ospf>area# exit

```

The following example displays the OSPF configuration with the modifications entered in the previous example:

```

A:ALA-49>config>router>ospf# info

asbr
overload
overload-on-boot timeout 60
traffic-engineering
preference 9
external-preference 140
export "OSPF-Export"
exit
area 0.0.0.0
virtual-link 10.0.0.1 transit-area 0.0.0.1
authentication-type message-digest
message-digest-key 2 md5 "Mi6BQAFi3MI" hash

```

```
 exit
 virtual-link 1.2.3.4 transit-area 1.2.3.4
 hello-interval 9
 dead-interval 40
 exit
 interface "system"
 exit
 exit
 area 0.0.0.1
 exit
 area 0.0.0.20
 stub
 exit
 interface "to-HQ"
 priority 50
 exit
 exit
 area 0.0.0.25
 nssa
 exit
 interface "if2"
 priority 100
 exit
 exit
 area 0.0.0.40
 interface "test1"
 authentication-type password
 authentication-key "3WErEDozxyQ" hash
 exit
 exit
 area 1.2.3.4
 exit

A:ALA-49>config>router>ospf#
```

## 3.14 OSPF Configuration Command Reference

### 3.14.1 Command Hierarchies

- [Configuration Commands](#)

#### 3.14.1.1 Configuration Commands

```

config
 — router
 — [no] ospf [ospf-instance] [router-id]
 — advertise-router-capability {link | area | as}
 — no advertise-router-capability
 — advertise-tunnel-link
 — no advertise-tunnel-link
 — [no] area area-id
 — [no] advertise-router-capability
 — area-range ip-prefix/mask [advertise | not-advertise]
 — no area-range ip-prefix/mask
 — [no] blackhole-aggregate
 — [no] database-export-exclude
 — [no] export policy-name [policy-name]
 — [no] import policy-name [policy-name]
 — [no] interface ip-int-name [secondary]
 — [no] adjacency-set id
 — adjacency-sid label value
 — no adjacency-sid
 — [no] advertise-router-capability
 — [no] advertise-subnet
 — auth-keychain name
 — authentication-key [authentication-key | hash-key] [hash |
 hash2]
 — no authentication-key
 — authentication-type {password | message-digest}
 — no authentication-type
 — bfd-enable [remain-down-on-failure]
 — no bfd-enable
 — dead-interval seconds
 — no dead-interval
 — hello-interval seconds
 — no hello-interval
 — interface-type {broadcast | point-to-point | non-broadcast}
 — no interface-type
 — lfa-policy-map route-nh-template template-name
 — no lfa-policy-map
 — [no] loopfree-alternate-exclude

```

- 
- **lsa-filter-out** [all | except-own-rtrlsa | except-own-rtrlsa-and-defaults]
  - **no lsa-filter-out**
  - **message-digest-key** *key-id* **md5** [*key* | *hash-key* | *hash2-key*] [*hash* | *hash2*]
  - **no message-digest-key** *key-id*
  - **metric** *metric*
  - **no metric**
  - **mtu** *bytes*
  - **no mtu**
  - [no] **neighbor** *ipv4-address*
  - **node-sid** **index** *index-value* [clear-n-flag]
  - **node-sid** **label** *label-value* [clear-n-flag]
  - **no node-sid**
  - [no] **passive**
  - **poll-interval** *seconds*
  - **no poll-interval**
  - **priority** *number*
  - **no priority**
  - **retransmit-interval** *seconds*
  - **no retransmit-interval**
  - **rib-priority** {high}
  - **no rib-priority**
  - [no] **shutdown**
  - [no] **sid-protection**
  - **transit-delay** *seconds*
  - **no transit-delay**
  - [no] **loopfree-alternate-exclude**
  - [no] **nssa**
    - **area-range** *ip-prefix/mask* [advertise | not-advertise]
    - **no area-range** *ip-prefix/mask*
    - **originate-default-route** [type-7] [no-adjacency-check]
    - **no originate-default-route**
    - [no] **redistribute-external**
    - [no] **summaries**
  - [no] **stub**
    - **default-metric** *metric*
    - **no default-metric**
    - [no] **summaries**
  - [no] **virtual-link** *router-id* **transit-area** *area-id*
    - **auth-keychain** *name*
    - **authentication-key** [*authentication-key* | *hash-key*] [*hash* | *hash2*]
    - **no authentication-key**
    - **authentication-type** {password | message-digest}
    - **no authentication-type**
    - **dead-interval** *seconds*
    - **no dead-interval**
    - **hello-interval** *seconds*
    - **no hello-interval**
    - **message-digest-key** *key-id* **md5** [*key* | *hash-key* | *hash2-key*] [*hash* | *hash2*]
    - **no message-digest-key** *key-id*
    - **retransmit-interval** *seconds*

---

```

— no retransmit-interval
— [no] shutdown
— transit-delay seconds
— no transit-delay
— asbr [trace-path domain-id]
— no asbr
— [no] compatible-rfc1583
— database-export [identifier id] [bgp-ls-identifier bgp-ls-id]
— [no] disable-ldp-sync
— entropy-label
 — [no] override-tunnel-elc
— export policy-name [policy-name]
— no export
— export-limit number [log percentage]
— no export-limit
— external-db-overflow limit seconds
— no external-db-overflow
— external-preference preference
— no external-preference
— [no] graceful-restart
 — [no] helper-disable
 — [no] strict-lsa-checking
— igp-shortcut
 — [no] shutdown
 — tunnel-next-hop
 — family {ipv4 | srv4}
 — resolution {any | disabled | filter | match-family-ip}
 — resolution-filter
 — [no] rsvp
 — [no] sr-te
— import policy-name [policy-name]
— no import
— [no] ldp-over-rsvp
— [no] loopfree-alternates
 — exclude
 — prefix-policy prefix-policy prefix-policy [prefix-policy]
 — no prefix-policy
 — remote-lfa [max-pq-cost value]
 — no remote-lfa
 — node-protect [max-pq-nodes value]
 — no node-protect
 — ti-lfa [max-sr-frr-labels value]
 — no ti-lfa
 — [no] node-protect
— [no] multicast-import
— overload [timeout seconds]
— no overload
— [no] overload-include-ext-2
— [no] overload-include-stub
— overload-on-boot [timeout seconds]
— no overload-on-boot
— preference preference
— no preference
— reference-bandwidth bandwidth-in-kbps

```

- 
- **reference-bandwidth** [tbps *Tera-bps*] [gbps *Giga-bps*] [mbps *Mega-bps*] [kbits *Kilo-bps*]
  - **no reference-bandwidth**
  - **rib-priority** high *prefix-list-name*
  - **no rib-priority**
  - **router-id** *ip-address*
  - **no router-id**
  - **rtr-adv-lsa-limit** limit [log-only] [threshold *percent*] [overload-timeout {seconds | forever}]
  - **no rtr-adv-lsa-limit**
  - [no] **segment-routing**
    - **adj-sid-hold** seconds
    - **no adj-sid-hold**
    - [no] **adjacency-set** id
      - **parallel** [no-advertise]
      - **no parallel**
      - **sid** label value
      - **no sid**
    - **backup-node-sid** ip-prefix/prefix-length index value
    - **backup-node-sid** ip-prefix/prefix-length label value
    - **no backup-node-sid**
    - **entropy-label** {force-disable | enable}
    - **no entropy-label**
    - **export-tunnel-table** ldp
    - **no export-tunnel-table**
    - [no] **mapping-server**
      - [no] **shutdown**
      - **sid-map** node-sid index index-value [range range-value] prefix ip-address/mask [netmask]
      - **sid-map** node-sid index index-value [range range-value] prefix ip-address/mask [netmask] scope {area area-id | as}
      - [no] **sid-map** node-sid index index-value
    - **prefix-sid-range** {global | start-label label-value max-index index-value}
    - **no prefix-sid-range**
    - [no] **shutdown**
    - **srlb** reserved-label-block-name
    - **no srlb**
    - **tunnel-mtu** bytes
    - **no tunnel-mtu**
    - **tunnel-table-pref** preference
    - **no tunnel-table-pref**
  - [no] **shutdown**
  - **timers**
    - **incremental-spf-wait** inc-spf-wait
    - **no incremental-spf-wait**
    - **lsa-accumulate** lsa-accum-time
    - **no lsa-accumulate**
    - **lsa-arrival** lsa-arrival-time
    - **no lsa-arrival**
    - [no] **lsa-generate** max-lsa-wait [lsa-initial-wait lsa-initial-wait [lsa-second-wait lsa-second-wait]]
    - **no lsa-generate**

---

```

— [no] redistribute-delay redist-wait
— [no] spf-wait max-spf-wait [spf-initial-wait initial-wait [spf-second-
wait second-wait]]
— [no] traffic-engineering
— [no] unicast-import-disable
— ospf3 [ospf-instance] [router-id]
— no ospf3 [ospf-instance]
— advertise-router-capability {link | area | as}
— no advertise-router-capability
— [no] area area-id
— [no] advertise-router-capability
— area-range ipv4-prefix/mask | ipv6-prefix/prefix-length [advertise | not-
advertise]
— no area-range ipv4-prefix/mask | ipv6-prefix/prefix-length
— [no] blackhole-aggregate
— export policy-name [policy-name]
— no export
— extended-lsa {inherit | only}
— no extended-lsa
— import policy-name [policy-name]
— no import
— interface ip-int-name [secondary]
— no interface ip-int-name
— [no] advertise-router-capability
— authentication bidirectional sa-name
— authentication inbound sa-name outbound sa-name
— no authentication
— bfd-enable [remain-down-on-failure]
— no bfd-enable
— dead-interval seconds
— no dead-interval
— hello-interval seconds
— no hello-interval
— interface-type {broadcast | point-to-point | non-broadcast}
— no interface-type
— lfa-policy-map route-nh-template template-name
— no lfa-policy-map
— [no] loopfree-alternate-exclude
— lsa-filter-out [all | except-own-rtrlsa | except-own-rtrlsa-and-
defaults]
— [no] lsa-filter-out
— metric metric
— no metric
— mtu bytes
— no mtu
— [no] neighbor ipv6-address
— node-sid index index-value [clear-n-flag]
— node-sid label label-value [clear-n-flag]
— no node-sid
— [no] passive
— poll-interval seconds
— no poll-interval
— priority number
— no priority

```

---

```

— retransmit-interval seconds
— no retransmit-interval
— rib-priority high
— no rib-priority
— [no] shutdown
— transit-delay seconds
— no transit-delay
— key-rollover-interval seconds
— [no] loopfree-alternate-exclude
— [no] nssa
— area-range ipv4-prefix/mask | ipv6-prefix/prefix-length
 [advertise | not-advertise]
— no area-range ipv4-prefix/mask | ipv6-prefix/prefix-length
— originate-default-route [type-nssa] [no-adjacency-check]
— no originate-default-route
— [no] redistribute-external
— [no] summaries
— [no] stub
— default-metric metric
— no default-metric
— [no] summaries
— [no] virtual-link router-id transit-area area-id
— authentication bidirectional sa-name
— authentication inbound sa-name outbound sa-name
— no authentication
— dead-interval seconds
— no dead-interval
— hello-interval seconds
— no hello-interval
— retransmit-interval seconds
— no retransmit-interval
— [no] shutdown
— transit-delay seconds
— no transit-delay
— [no] asbr
— [no] disable-ldp-sync
— export policy-name [policy-name]
— no export
— export-limit number [log percentage]
— no export-limit
— extended-lsa {sparse | only}
— no extended-lsa
— external-db-overflow limit seconds
— no external-db-overflow
— external-preference preference
— no external-preference
— [no] graceful-restart
— [no] helper-disable
— [no] strict-lsa-checking
— igp-shortcut
— [no] shutdown
— tunnel-next-hop
— family ipv6
— resolution {any | disabled | filter}

```



---

```

 — resolution-filter
 — [no] rsvp
 — [no] sr-te
— import policy-name [policy-name]
— no import
— [no] loopfree-alternates
 — exclude
 — prefix-policy prefix-policy prefix-policy [prefix-policy]
 — no prefix-policy
 — remote-lfa [max-pq-cost value]
 — no remote-lfa
 — node-protect [max-pq-nodes value]
 — no node-protect
 — ti-lfa [max-sr-frr-labels value]
 — no ti-lfa
 — [no] node-protect
— [no] multicast-import
— overload [timeout seconds]
— no overload
— [no] overload-include-ext-2
— [no] overload-include-stub
— overload-on-boot [timeout seconds]
— no overload-on-boot
— preference preference
— no preference
— reference-bandwidth bandwidth-in-kbps
— reference-bandwidth [tbps Tera-bps] [gbps Giga-bps] [mbps Mega-bps]
 [kbps Kilo-bps]
— no reference-bandwidth
— rib-priority {high} prefix-list-name
— no rib-priority {high}
— router-id ip-address
— no router-id
— rtr-adv-lsa-limit limit [log-only] [threshold percent] [overload-timeout
 {seconds | forever}]
— no rtr-adv-lsa-limit
— [no] segment-routing
 — adj-sid-hold seconds
 — no adj-sid-hold
 — prefix-sid-range {global | start-label label-value max-index index-
 value}
 — no prefix-sid-range
 — [no] shutdown
 — tunnel-mtu bytes
 — no tunnel-mtu
 — tunnel-table-pref preference
 — no tunnel-table-pref
— [no] shutdown
— timers
 — incremental-spf-wait inc-spf-wait
 — no incremental-spf-wait
 — lsa-accumulate lsa-accumulate
 — no lsa-accumulate
 — lsa-arrival lsa-arrival-time

```

- **no** **lsa-arrival**
- **lsa-generate** *max-lsa-wait* [**lsa-initial-wait** *lsa-initial-wait* [**lsa-second-wait** *lsa-second-wait*]]
- **no** **lsa-generate**
- **redistribute-delay** *redistribute-delay*
- **no** **redistribute-delay**
- **spf-wait** *max-spf-wait* [**spf-initial-wait** *initial-wait* [**spf-second-wait** *second-wait*]]
- **[no]** **spf-wait**
- **[no]** **unicast-import-disable**

## 3.14.2 Command Descriptions

### 3.14.2.1 Generic Commands

#### shutdown

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no]</b> <b>shutdown</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf>area>interface<br>config>router>ospf>segment-routing<br>config>router>ospf>segment-routing>mapping-server<br>config>router>ospf>area>virtual-link<br>config>router>ospf>igp-shortcut<br>config>router>ospf>segment-routing<br>config>router>ospf3<br>config>router>ospf3>area>interface<br>config>router>ospf3>area>virtual-link                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>The <b>shutdown</b> command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the <b>no shutdown</b> command.</p> <p>The <b>shutdown</b> command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>Unlike other commands and parameters where the default state is not indicated in the configuration file, <b>shutdown</b> and <b>no shutdown</b> are always indicated in system generated configuration files.</p> <p>The <b>no</b> form of the command puts an entity into the administratively enabled state.</p> |

---

**Special Cases**     **OSPF Interface** — When an IP interface is configured as an OSPF interface, OSPF on the interface is in the **no shutdown** state by default.

### 3.14.2.2    OSPF Global Commands

#### ospf

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ospf</b> <i>ospf-instance</i> [ <i>router-id</i> ]<br>[ <b>no</b> ] <b>ospf</b> <i>ospf-instance</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command creates an OSPF routing instance and then enters the associated context to configure the associated protocol parameters.</p> <p>Additionally, the router ID can be specified as another parameter of the OSPF command. This parameter is required for all non-base OSPF instances.</p> <p>The default value for the base instance is inherited from the configuration in the <b>config&gt;router</b> context. When that is not configured, the following apply:</p> <ol style="list-style-type: none"> <li>1. the system uses the system interface address (which is also the loopback address)</li> <li>2. if a system interface address is not configured, it uses the last 32 bits of the chassis MAC address</li> </ol> <p>This is a required command when configuring multiple instances and the instance being configured is not the base instance. When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To prevent this from happening, all routers in a domain should be configured with the same domain ID. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.</p> <p>The default value for non-base instances is 0.0.0.0 and is invalid; in this case, the instance of OSPF will not start. When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.</p> <p>Issue the shutdown and no shutdown commands for the instance for the new router ID to be used, or reboot the entire router.</p> <p>OSPF instances are <b>shutdown</b> when created, so that all parameters can be configured prior to the instance being enabled.</p> <p>The <b>no</b> form of the command to reverts to the default value.</p> |
| <b>Default</b>     | no ospf                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Parameters**     *ospf-instance* — Specifies a unique integer that identifies a specific instance of a version of the OSPF protocol running in the router instance specified by the router ID.

**Values**        1 to 31

*router-id* — Specifies the OSPF router ID to be used with the associated OSPF instance. This IP address must be given a dot decimal notation format.

## ospf3

**Syntax**        **ospf3** [*ospf-instance*] [*router-id*]  
                  [**no**] **ospf3** *instance-id*

**Context**        config>router

**Description**   This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.

OSPFv3 instances are **shutdown** when created, so that all parameters can be configured prior to the instance being enabled.

The **no** form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.

**Parameters**     *ospf-instance* — Specifies the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family. For IPv6 unicast, the instance id must be between 0 and 31. For IPv4 unicast the instance id must be between 64-95.

**Values**        0 to 31: IPV6 unicast

**Values**        64 to 95: IPV4 unicast

*router-id* — Specifies the OSPF router ID to be used with the associated OSPF instance. This IP address must be given a dot decimal notation format.

## advertise-router-capability

**Syntax**        **advertise-router-capability** {link | area | as}  
                  **no advertise-router-capability**

**Context**        config>router>ospf  
                  config>router>ospf3

**Description**     This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled

- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The parameters (**link**, **area** and **as**) control the scope of the capability advertisements.

The **no** form of this command disables this capability.

|                   |                                                                                                                                                                                                                                                                      |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no advertise-router-capability                                                                                                                                                                                                                                       |
| <b>Parameters</b> | <p><b>link</b> — capabilities are only advertised over local links and not flooded beyond</p> <p><b>area</b> — capabilities are only advertised within the area of origin</p> <p><b>as</b> — capabilities are advertised throughout the entire autonomous system</p> |

## advertise-tunnel-link

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>advertise-tunnel-link</b><br><b>no advertise-tunnel-link</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>ospf                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command enables the forwarding adjacency feature. With this feature, IS-IS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs per RFC 3906.</p> <p>The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both <b>igp-shortcut</b> and <b>advertise-tunnel-link</b> options are enabled for a given IGP instance, then the <b>advertise-tunnel-link</b> will win.</p> <p>When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router-id of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA/LSP advertising the corresponding P2P unnumbered link is installed in the local routing database, then the node performs an SPF using it like any other link LSA/LSP. The link bi-directional check requires that a link, regular link or tunnel link, exists in the reverse direction for the tunnel to be used in SPF.</p> <p>That the <b>igp-shortcut</b> option under the LSP name governs the use of the LSP with both the <b>igp-shortcut</b> and the <b>advertise-tunnel-link</b> options in IGP. In other words, the user can exclude a specific RSVP LSP from being used as a forwarding adjacency by entering the command <b>config&gt;router&gt;mpls&gt;lsp&gt;no igp-shortcut</b>.</p> <p>The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.</p> |

The **no** form of this command disables forwarding adjacency and hence disables the advertisement of RSVP LSP into IGP.

**Default** no advertise-tunnel-link

asbr

**Syntax** [no] asbr [trace-path *domain-id*]  
**no asbr**  
**[no] asbr**


**Context** config>router>ospf  
config>router>ospf3

**Description** This command configures the router as an Autonomous System Boundary Router (ASBR) if the router is to be used to export routes from the Routing Table Manager (RTM) into this instance of OSPF. Once a router is configured as an ASBR, the export policies into this OSPF domain take effect. If no policies are configured, no external routes are redistributed into the OSPF domain.

The **no** form of the command removes the ASBR status and withdraws the routes redistributed from the Routing Table Manager into this instance of OSPF from the link state database.

When configuring multiple instances of OSPF, there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To prevent this from happening, all routers in a domain should be configured with the same domain ID. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.



Domain IDs are incompatible with any other use of normal tags. The domain ID should be configured with a value between 1 and 31 by each router in a given OSPF domain (OSPF Instance).

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding (1-31) bit is set in the AS-external LSA.

As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

|                   |                                             |
|-------------------|---------------------------------------------|
| <b>Default</b>    | no asbr                                     |
| <b>Parameters</b> | <i>domain-id</i> — Specifies the domain ID. |
| <b>Values</b>     | 1 to 31                                     |
| <b>Default</b>    | 0                                           |

## compatible-rfc1583

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] compatible-rfc1583</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ospf                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.</p> <p>RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.</p> <p>Although it would be favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.</p> <p>The <b>no</b> form of the command enables the post-RFC1583 method of summary and external route calculation.</p> |
| <b>Default</b>     | compatible-rfc1583                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## database-export

|                    |                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database-export [identifier <i>id</i>] [bgp-ls-identifier <i>bgp-ls-id</i>]<br/>no database-export</b>                          |
| <b>Context</b>     | config>router>ospf                                                                                                                 |
| <b>Description</b> | This command enables the population of the extended TE Database (TE-DB) with the link-state information from a given IGP instance. |

The extended TE-DB is used as a central point for importing all link-state information, link, node, and prefix, from IGP instances on the router or the vSROS controller of the NSP and to exporting it to BGP-LS on the router and to Java-VM proxy on the vSROS controller. This information includes the IGP, TE, and the SR information, prefix SID sub-TLV, adjacency SID sub-TLV, and router SR capability TLV.

The **no** form of the command disables database exportation.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no database-export                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b> | <p><b>identifier</b> — This parameter is used to uniquely identify the IGP instance in the BGP-LS NLRI when a router has interfaces participating in multiple IGP instances. This parameter defaults to the IGP instance ID assigned by SR OS. However, given that the concept of instance ID defined in IS-IS (RFC 6822) is unique within a routing domain while the one specified for OSPF is local subnet significant (RFC 6549), the user can remove any overlap by configuring the new <b>identifier</b> value to be unique within a given IGP domain when this router sends the IGP link-state information using BGP-LS.</p> <p><i>id</i> — Specifies an entry ID to export.</p> <p><b>Values</b> 0 to 18446744073709551615</p> <p><b>bgp-ls-identifier</b> — This parameter is used, along with the Autonomous System Number (ASN) to correlate the BGP-LS NLRI advertisements of multiples BGP-LS speakers of the same IGP domain. If an NRC-P network domain has multiple IGP domains, BGP-LS speakers within each IGP domain must be configured with the same unique {bgp-ls-identifier, asn} tuple.</p> <p>The BGP-LS identifier is optional and is only sent in a BGP-LS NLRI if configured in the IGP instance of an IGP domain.</p> <p>Note that if this IGP instance participates in traffic engineering with RSVP-TE or SR-TE, the <b>traffic-engineering</b> option is not strictly required because enabling the extended TE-DB populates this information automatically. It is, however, recommended to enable it to make the configuration consistent with other routers in the network that do not require the enabling of the extended TE-DB.</p> <p><i>bgp-ls-id</i> — Specifies a BGP LS ID to export.</p> <p><b>Values</b> 0 to 4294967295</p> |

## disable-ldp-sync

|                |                                           |
|----------------|-------------------------------------------|
| <b>Syntax</b>  | <b>[no] disable-ldp-sync</b>              |
| <b>Context</b> | config>router>ospf<br>config>router>ospf3 |



**Description** This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It will then disable IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the `ldp-sync-timer` is configured.

**Default** no disable-ldp-sync

## entropy-label

**Syntax** **entropy-label**

**Context** config>router>isis  
config>router>ospf

**Description** This command enables the context for the configuration of entropy label capabilities for the routing protocol.

## override-tunnel-elc

**Syntax** [**no**] **override-tunnel-elc**

**Context** config>router>isis>entropy-label  
config>router>ospf>entropy-label

**Description** This command configures the ability to override any received entropy label capability advertisement. When enabled, the system assumes that all nodes for an IGP domain are capable of receiving and processing the entropy label on segment routed tunnels. This command can only be configured if **entropy-label** is enabled via the **config>router>isis>segment-routing>entropy-label** or **config>router>ospf>segment-routing>entropy-label** command.

The **no** version of this command disables the override. The system assumes entropy label capability for other nodes in the IGP instance if capability advertisements are received.

**Default** no override-tunnel-elc

## export

**Syntax** **export** *policy-name* [*policy-name*]

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>no export</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command associates export route policies to determine which routes are exported from the route table to OSPF. Export policies are only in effect if OSPF is configured as an ASBR.</p> <p>If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.</p> <p>If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.</p> <p>The <b>no</b> form of the command removes all policies from the configuration.</p> |
| <b>Default</b>     | no export                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>policy-name</i> — The export route policy name, up to a maximum of 5. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p>The specified names must already be defined.</p>                                                                                                                                                                                                                                                                                                                                          |

## export-limit

|                    |                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-limit</b> <i>number</i> [ <i>log percentage</i> ]<br><b>no export-limit</b>                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the maximum number of routes (prefixes) that can be exported into OSPF from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.</p> <p>The <b>no</b> form of the command removes the parameters from the configuration.</p>                           |
| <b>Default</b>     | no export-limit                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into OSPF from the route table</p> <p><b>Values</b> 1 to 4294967295</p> <p><i>percentage</i> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p><b>Values</b> 1 to 100</p> |

## extended-lsa

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>extended-lsa</b> { <b>sparse</b>   <b>only</b> }<br><b>no extended-lsa</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the use of extended LSA format in OSPFv3 as per <i>draft-ietf-ospf-ospfv3-lsa-extend</i>.</p> <p>Prior to this feature, SR OS used the fixed format LSA to carry the prefix and link information as per RFC 5340, <i>OSPF for IPv6</i>. The fixed format is not extensible and the TLV format of the extended LSA must be used.</p> <p>With this feature, the default mode of operation for OSPFv3 is referred to as <b>sparse</b> mode, meaning that the router will always advertise the fixed format for existing LSAs and will add the TLV-based extended LSA only when it needs to advertise new sub-TLVs. This mode of operation is similar to the way OSPFv2 advertises the segment routing information. It sends the prefix in the original fixed-format prefix LSA and then follows with the extended prefix TLV which is sent in an extended prefix opaque LSA containing the prefix SID sub-TLV.</p> <p>The <b>extended-lsa only</b> value enables the full extended LSA mode. This causes all existing and new LSAs to use the extended LSA format.</p> <p>The OSPFv3 instance must first be shut down before the user can change the mode of operation since the protocol must flush all LSAs and re-establish all adjacencies.</p> <p>The <b>no</b> form at the OSPFv3 instance level reverts the OSPFv3 instance into the default sparse mode of operation.</p> |
| <b>Default</b>     | extended-lsa sparse                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <p><b>sparse</b> — Enables the sparse mode of operation in an OSPFv3 instance</p> <p><b>only</b> — Enables the full extended LSA mode of operation in an OSPFv3 instance</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## external-db-overflow

|                    |                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>external-db-overflow</b> <i>limit interval</i><br><b>no external-db-overflow</b>                                                                                                             |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                       |
| <b>Description</b> | This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded. |

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of the command disables limiting the number of non-default AS-external-LSA entries.

**Default** no external-db-overflow

**Parameters** *limit* — The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

**Values** -1 to 2147483647

**Note:** Setting a value of -1 is equivalent to no **external-db-overflow**.



*interval* — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

**Values** 0 to 2147483647

external-preference

**Syntax** external-preference *preference*  
no external-preference

**Context** config>router>ospf  
config>router>ospf3

**Description** This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the [Table 14](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

The **no** form of the command reverts to the default value.

**Default** external-preference 150

**Parameters** *preference* — The preference for external routes expressed as a decimal integer. Defaults for different route types are listed in [Table 14](#).

**Table 14** Route Preference Defaults by Route Type

| Route Type             | Preference | Configurable     |
|------------------------|------------|------------------|
| Direct attached        | 0          | No               |
| Static routes          | 5          | Yes              |
| OSPF internal          | 10         | Yes <sup>1</sup> |
| IS-IS level 1 internal | 15         | Yes              |
| IS-IS level 2 internal | 18         | Yes              |
| RIP                    | 100        | Yes              |
| OSPF external          | 150        | Yes              |
| IS-IS level 1 external | 160        | Yes              |
| IS-IS level 2 external | 165        | Yes              |
| BGP                    | 170        | Yes              |

**Note:**

1. Preference for OSPF internal routes is configured with the **preference** command.

**Values** 1 to 255

## graceful-restart

**Syntax** **[no] graceful-restart**

**Context** config>router>ospf  
config>router>ospf3

**Description** This command enables OSPF graceful restart (GR) to minimize service disruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve OSPF forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router comes back up within the grace period, the restarting router resumes normal OSPF operation. If the grace period expires, then the restarting router is presumed inactive and the OSPF topology is recalculated to route traffic around the failure.

The **no** form of the command disables graceful restart and removes the graceful restart configuration from the OSPF instance.

**Default** no graceful-restart

## helper-disable

**Syntax** [no] helper-disable

**Context** config>router>ospf>graceful-restart  
config>router>ospf3>graceful-restart

**Description** This command disables helper support for OSPF graceful restart (GR).

When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart), a restarting router, or both. The router only supports helper mode. It will not act as a restarting router because the high availability feature set already preserves OSPF forwarding information so that this functionality is not needed. This command is a historical command and should not be disabled. Configuring **helper-disable** has the effect of disabling graceful restart because the router only supports helper mode.

The **no helper-disable** command enables helper support and is the default when **graceful-restart** is enabled.

**Default** no helper-disable

## strict-lsa-checking

**Syntax** [no] strict-lsa-checking

**Context** config>router>ospf>graceful-restart  
config>router>ospf3>graceful-restart

**Description** This command indicates whether an OSPF restart helper should terminate graceful restart when there is a change to an LSA that would be flooded to the restarting router during the restart process.

The default OSPF behavior is to terminate a graceful restart if an LSA changes, which causes the OSPF neighbor to go down.

The **no** form of the command disables strict LSA checking.

**Default** strict-lsa-checking

## igp-shortcut

**Syntax** **igp-shortcut**

**Context** config>router>ospf  
config>router>ospf3

**Description** This command enables the use of an RSVP-TE or SR-TE shortcut for resolving IGP routes by OSPF or IS-IS routing protocols.

This command instructs IGP to include RSVP LSPs and SR-TE LSPs originating on this node and terminating on the router ID of a remote node as direct links with a metric equal to the metric provided by MPLS.

During the IP reach calculation to determine the reachability of nodes and prefixes, LSPs are overlaid and the LSP metric is used to determine the subset of paths that are equal lowest cost to reach a node or a prefix. If the user enabled the **relative-metric** option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix that is resolved to the LSP.

When a prefix is resolved to a tunnel next hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP-TE or SR-TE LSP, as well as the explicit-null IPv6 label at the bottom of the stack in the case of an IPv6 prefix. Any network event causing one or more IGP shortcuts to go down will trigger a full SPF computation, which may result in installing a new route over an updated set of tunnel next-hops and IP next-hops.

When **igp-shortcut** is enabled at the IGP instance level, all RSVP-TE and SR-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router ID of a remote node. LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IGP. The user can, however, exclude a specific RSVP-TE or SR-TE LSP from being used as a shortcut for resolving IGP routes by entering the **config>router>mpls>lsp>no igp-shortcut** command.

The SPF in IGP only uses RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If two or more options are enabled in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

The SPF in IGP uses SR-TE LSPs as IGP shortcuts only.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of tunnel and IP next-hops to program in the data path.

- Where a destination is a tunnel-endpoint (including external prefixes with tunnel-endpoint as the next hop), the tunnel with lowest tunnel-index is selected (the IP next hop is never used in this case).
- Where a destination is not a tunnel-endpoint:
  - LSPs with metric higher than underlying IGP cost between the endpoint of the LSP are excluded
  - Tunnel next-hops are preferred over IP next-hops
  - Within tunnel next-hops:
    - i. The lowest endpoint-to-destination cost is selected
    - ii. If the endpoint-to-destination costs are the same, the lowest endpoint node router ID is selected
    - iii. If the router IDs are the same, the lowest tunnel index is selected
  - Within IP next-hops:
    - i. The lowest downstream router ID is selected
    - ii. If the downstream router IDs are the same, the lowest interface-index is selected



**Note:** Although ECMP is not performed across both the IP and tunnel next-hops, the tunnel endpoint may lie in one of the shortest IGP paths for that prefix. In that case, the tunnel next hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

When both RSVP-TE and SR-TE IGP shortcuts are available, the IP reach calculation, in the unicast routing table, will first follow the above ECMP tunnel and IP next hop selection rules when resolving a prefix over IGP shortcuts. After the set of ECMP tunnel and IP next-hops have been selected, the preference of tunnel type is then applied based on the user setting of the resolution of the family of the prefix. If the user enabled resolution of the prefix family to both RSVP-TE and SR-TE tunnel types, the TTM tunnel preference value is used to select one type for the prefix. In other words, the RSVP-TE LSP type is preferred to an SR-TE LSP type on a per-prefix basis.

The ingress IOM sprays the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next hop when both the **igp-shortcut** and the **multicast-import** options are enabled in IGP. The unicast RTM can still make use of the tunnel next hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

## tunnel-next-hop

**Syntax**    **tunnel-next-hop**



---

|                    |                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf>igp-shortcut<br>config>router>ospf3>igp-shortcut                                                                                                                                                                   |
| <b>Description</b> | This command enables the context to configure the resolution of IGP IPv4 prefix families, IGP IPv6 prefix families, SR-ISIS IPv4 tunnel families, SR-ISIS IPv6 tunnel families, and SR-OSPF IPv4 tunnel families using IGP shortcuts. |

The **resolution** node is introduced to provide flexibility in the selection of the tunnel types for each of the IP prefix and SR tunnel families.

The IPv4 **family** option causes the IS-IS or OSPF SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv4 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv4 prefixes owned by this node or for which this node is the IPv4 next hop.

The IPv6 **family** option causes the IS-IS or OSPFv3 SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv6 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv6 prefixes owned by this node or for which this node is the IPv6 next hop. The resolution of IPv6 prefixes is supported in OSPFv3 and in both IS-IS MT=0 and MT=2.

The IS-IS and OSPFv3 IPv6 routes resolved to IPv4 IGP shortcuts are used to:

- forward packets of IS-IS or OSPFv3 prefixes matching these routes
- forward CPM-originated IPv6 packets
- resolve the BGP next hop of BGP IPv6 prefixes
- resolve the indirect next hop of static IPv6 routes

In the data path, a packet for an IPv6 prefix has a label stack that consists of the IPv6 Explicit-Null label value of 2 at the bottom of the label stack followed by the label stack of the IPv4 RSVP-TE LSP.

There is no default behavior for IPv4 prefixes to automatically resolve to RSVP-TE or SR-TE LSPs used as IGP shortcuts by only enabling the **igp-shortcut** context. Instead, the user must enable the **ipv4 family** or **ipv6 family** and set the resolution to the value of **rsvp-te** to select the RSVP-TE tunnel type, or to the value of **sr-te** to select the SR-TE tunnel type.

Setting the **resolution** to the **any** value means that IGP selects the tunnels used as IGP shortcuts according to the TTM preference for the tunnel type. The RSVP-TE LSP type is of higher priority than the SR-TE LSP type.

An IP prefix of family=ipv4 or family= ipv6 always resolves to a single type of tunnel **rsvp-te** or **sr-te**. **Rsvp-te** type is preferred if both types are allowed by the prefix family resolution and both types exist in the set of tunnel next-hops of the prefix. The feature does not support mixing tunnel types per prefix.

If **resolution** for the IPv4 or IPv6 family is set to **disabled**, the corresponding prefixes are resolved to IP next-hops in the multicast routing table.

The **srv4 family** enables the resolution of SR-OSPF IPv4 tunnels and SR-ISIS IPv4 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-OSPF or an SR-ISIS IPv4 tunnel.

The **srv6 family** enables the resolution of SR-ISIS IPv6 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-ISIS IPv6 tunnel.

One or more RSVP-TE LSPs can be selected if **resolution=match-family-ip** and the corresponding IPv4 or IPv6 prefix resolves to RSVP-TE LSPs.



**Note:** An SR tunnel cannot resolve to SR-TE IGP shortcuts.

If **resolution** for the SRv4 or SRv6 tunnel family is set to **disabled**, the corresponding tunnels are resolved to IP next-hops in the multicast routing table.

To enable or disable IGP shortcuts in the IGP instance, the user must perform a **shutdown** or **no shutdown** in the **igp-shortcut** context.

family

|                    |                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>family {ipv4   srv4}</b>                                                                                                            |
| <b>Context</b>     | config>router>ospf>igp-shortcut>tunnel-next-hop                                                                                        |
| <b>Description</b> | This command enables the context to configure the resolution of the IGP IPv4 prefix family or SR-OSPF IPv4 tunnel using IGP shortcuts. |
| <b>Parameters</b>  | <b>ipv4</b> — Selects the IPv4 address family.<br><b>srv4</b> — Selects the SR-OSPF IPv4 tunnel family.                                |

family

|                    |                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>family ipv6</b>                                                                                              |
| <b>Context</b>     | config>router>ospf3>igp-shortcut>tunnel-next-hop                                                                |
| <b>Description</b> | This command enables the context to configure the resolution of the IGP IPv6 prefix family using IGP shortcuts. |
| <b>Parameters</b>  | <b>ipv6</b> — Selects the IPv6 address family.                                                                  |

## resolution

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resolution {any   disabled   filter   match-family-ip}</b>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>ospf>igp-shortcut>tunnel-next-hop>family                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command configures resolution mode in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><b>any</b> — Enables the binding to any supported tunnel type following TTM preference.</p> <p><b>disabled</b> — Disables the resolution of the IP prefix or SR tunnel family using IGP shortcuts.</p> <p><b>filter</b> — Enables the binding to the subset of tunnel types configured under <b>resolution-filter</b>.</p> <p><b>match-family-ip</b> — Enables the resolution of the SR tunnel family to match that of the corresponding IP prefix family.</p> |

## resolution

|                    |                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resolution {any   disabled   filter}</b>                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>ospf3>igp-shortcut>tunnel-next-hop>family                                                                                                                                                                                                                                                     |
| <b>Description</b> | This command configures resolution mode in the resolution of the IPv6 prefix using IGP shortcuts.                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><b>any</b> — Enables the binding to any supported tunnel type following TTM preference.</p> <p><b>disabled</b> — Disables the resolution of the IPv6 prefix using IGP shortcuts.</p> <p><b>filter</b> — Enables the binding to the subset of tunnel types configured under <b>resolution-filter</b>.</p> |

## resolution-filter

|                    |                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resolution-filter</b>                                                                                                                                              |
| <b>Context</b>     | config>router>ospf>igp-shortcut>tunnel-next-hop>family<br>config>router>ospf3>igp-shortcut>tunnel-next-hop>family                                                     |
| <b>Description</b> | This command enables the context to configure the subset of tunnel types that can be used in the resolution of the IP prefix or SR tunnel family using IGP shortcuts. |

## rsvp

|               |                  |
|---------------|------------------|
| <b>Syntax</b> | <b>[no] rsvp</b> |
|---------------|------------------|

**Context** config>router>ospf>igp-shortcut>tunnel-next-hop>family  
config>router>ospf3>igp-shortcut>tunnel-next-hop>family

**Description** This command selects the RSVP-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

## sr-te

**Syntax** [no] sr-te

**Context** config>router>ospf>igp-shortcut>tunnel-next-hop>family  
config>router>ospf3>igp-shortcut>tunnel-next-hop>family

**Description** This command selects the SR-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.

## import

**Syntax** import *policy-name* [*policy-name*]  
no import

**Context** config>router>ospf  
config>router>ospf3

**Description** This command applies one or more (up to 5) route policies as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy, it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions. The **no** form of the command removes all policies from the configuration.

**Default** no import

**Parameters** *policy-name* — The export route policy name, up to a maximum of 5. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.  
The specified names must already be defined.

## ldp-over-rsvp

**Syntax** [no] ldp-over-rsvp

**Context** config>router>ospf

---

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Description</b> | This command allows LDP-over-RSVP processing in this OSPF instance. |
| <b>Default</b>     | no ldp-over-rsvp                                                    |

## loopfree-alternates

|                    |                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>loopfree-alternates [remote-lfa]</b><br><b>loopfree-alternates remote-lfa [max-pq-cost value]</b><br><b>no loopfree-alternates</b> |
| <b>Context</b>     | config>router>ospf                                                                                                                    |
| <b>Description</b> | This command enables loop-free alternate (LFA) computation by SPF under the OSPFv2 routing protocol instance.                         |

When this command is enabled, it instructs the IGP SPF to attempt to precalculate both a primary next hop and an LFA backup next hop for every learned prefix. When found, the LFA next hop is populated into the routing table along with the primary next hop for the prefix.

The user enables the remote LFA next hop calculation by the IGP LFA SPF by using the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the additional remote LFA computation that follows the regular LFA next hop calculation when the latter calculation results in no protection for one or more prefixes that are resolved to a particular interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing or tearing down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node that puts the packets back into the shortest path without looping them back to the node that forwarded them over the repair tunnel. The remote LFA node is referred to as a PQ node. A repair tunnel can, in theory, be an RSVP-TE LSP, an LDP-in-LDP tunnel, or a segment routing (SR) tunnel. In this command, **remote-lfa** is restricted to using an SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix calculation like the regular LFA algorithm. The remote LFA algorithm provides protection for all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for all the destinations.

The **no** form of this command disables the LFA computation by the IGP SPF.

|                   |                                                                                                                                                                                                                                  |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no loopfree-alternates                                                                                                                                                                                                           |
| <b>Parameters</b> | <i>value</i> — specifies the integer used to limit the search of candidate P- and Q-nodes in the remote LFA by setting the maximum IGP cost from the router performing the remote LFA calculation to the candidate P- or Q-node. |
| <b>Values</b>     | 0 to 4294967295                                                                                                                                                                                                                  |

---

## loopfree-alternates

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] loopfree-alternates</b>                                                                               |
| <b>Context</b>     | config>router>ospf3                                                                                           |
| <b>Description</b> | This command enables Loop-Free Alternate (LFA) computation by SPF under the OSPFv3 routing protocol instance. |

When this command is enabled, it instructs the IGP SPF to attempt to precalculate both a primary next hop and an LFA next hop for every learned prefix. When found, the LFA next hop is populated into the routing table along with the primary next hop for the prefix.

The user enables the remote LFA next hop calculation by the IGP LFA SPF by appending the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next hop calculation when the latter resulted in no protection for one or more prefixes which are resolved to a particular interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing or tearing down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node that puts the packets back into the shortest path without looping them back to the node that forwarded them over the repair tunnel. The remote LFA node is referred to as a PQ node. A repair tunnel can, in theory, be an RSVP-TE LSP, an LDP-in-LDP tunnel, or a segment routing (SR) tunnel. In this command, **remote-lfa** is restricted to using an SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix calculation like the regular LFA algorithm. The remote LFA algorithm provides protection for all destination prefixes that share the protected link by using the neighbor on the other side of the protected link as a proxy for all the destinations.

The Topology-Independent LFA (TI-LFA) further improves the protection coverage of a network topology by computing and automatically instantiating a repair tunnel to a Q node which is not in shortest path from the computing node. The repair tunnel uses shortest path to the P node and a source routed path from the P node to the Q node.

In addition, the TI-LFA algorithm selects the backup path which matches the post-convergence path. This helps the capacity planning in the network since traffic will always flow on the same path when transitioning to the FRR next hop and then onto the new primary next hop.

At a high level, the TI-LFA protection algorithm is searching for a candidate P-Q set separated with a number of hops such that the label stack size does not exceed the value of **ti-lfa max-sr-frr-labels**, on each of the post-convergence paths to each destination node or prefix D.

When the **ti-lfa** option is enabled in OSPF, it provides TI-LFA node-protect or link-protect backup path for a SR-OSPF IPV4 tunnel (node SID and adjacency SID), and for a IPv4 SR-TE LSP.

The **max-sr-frr-labels** parameter is used to limit the search for the TI-LFA backup next hop:

- a. 0 — The IGP LFA SPF restricts the search to TI-LFA backup next hop which does not require a repair tunnel, meaning that P node and Q node are the same and match a neighbor. This is also the case when both P and Q node match the advertising router for a prefix.
- b. 1 to 3 — The IGP LFA SPF will widen the search to include a repair tunnel to a P node which itself is connected to the Q nodes with a 0-to-2 hops for a total of maximum of three labels: one node SID to P node and two adjacency SIDs from P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed and meaning that up to three adjacency SIDs can separate the P and Q nodes.
- c. 2 (default) — Corresponds to a repair tunnel to a non-adjacent P which is adjacent to the Q node. If the P node is a neighbor of the computing node, then the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

The TI-LFA repair tunnel can have a maximum of three labels pushed in addition to the label of the destination node or prefix. The user can set a lower maximum value for the additional FRR labels by configuring the CLI option **max-sr-frr-labels** *labels*. The default value is 2.

When the **node-protect** command is enabled, the router will prefer a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The parameter **max-pq-nodes** in Remote LFA controls the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

The **no** form of this command disables the LFA computation by the IGP SPF.

**Default** no loopfree-alternates

## exclude

**Syntax** **exclude**

**Context** config>router>ospf>loopfree-alternates  
config>router>ospf3>loopfree-alternates

**Description** This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.



**Note:** Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

**Default** no exclude

## prefix-policy

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] prefix-policy</b> <i>prefix-policy</i> [ <i>prefix-policy</i> ]                                                                                                                                             |
| <b>Context</b>     | config>router>ospf>loopfree-alternates>exclude<br>config>router>ospf3>loopfree-alternates>exclude                                                                                                                   |
| <b>Description</b> | This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this OSPF or OSPF3 instance.<br><br>The <b>no</b> form of this command deletes the exclude prefix policy. |
| <b>Default</b>     | no prefix-policy                                                                                                                                                                                                    |
| <b>Parameters</b>  | <b>prefix-policy</b> <i>prefix-policy</i> — Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.           |

## remote-lfa

|                    |                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>remote-lfa</b> [ <b>max-pq-cost</b> <i>value</i> ]<br><b>no remote-lfa</b>                                                                                                                                                                                         |
| <b>Context</b>     | config>router>ospf>loopfree-alternates<br>config>router>ospf3>loopfree-alternates                                                                                                                                                                                     |
| <b>Description</b> | This command enables the use of the Remote LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance.<br><br>The <b>no</b> form of this command disables the use of the Remote LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance. |



---

|                   |                                                                                                                                                                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no remote-lfa                                                                                                                                                                                                                                     |
| <b>Parameters</b> | <b>max-pq-cost</b> <i>value</i> — Specifies the integer used to limit the search of candidate P and Q nodes in the remote LFA by setting the maximum IGP cost from the router performing the remote LFA calculation to the candidate P or Q node. |
| <b>Values</b>     | 0 to 4294967295                                                                                                                                                                                                                                   |
| <b>Default</b>    | 4261412864                                                                                                                                                                                                                                        |

## node-protect

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>node-protect</b> [ <b>max-pq-nodes</b> <i>value</i> ]<br><b>no node-protect</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>ospf>loopfree-alternates>remote-lfa<br>config>router>ospf3>loopfree-alternates>remote-lfa                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.<br><br>The <b>no</b> form of this command disables node-protect.                                                                                                                                                                                                                                                          |
| <b>Default</b>     | no node-protect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <b>max-pq-nodes</b> <i>value</i> — Specifies the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the <b>max-pq-cost</b> parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel |
| <b>Values</b>      | 1 to 32                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | 16                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## ti-lfa

|                    |                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ti-lfa</b> [ <b>max-sr-frr-labels</b> <i>value</i> ]<br><b>no ti-lfa</b>                                                       |
| <b>Context</b>     | config>router>ospf>loopfree-alternates<br>config>router>ospf3>loopfree-alternates                                                 |
| <b>Description</b> | This command enables the use of the Topology-Independent LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance. |

The **no** form of this command disables the use of the Topology-Independent LFA algorithm in the LFA SPF calculation in this OSPF or OSPF3 instance.

|                   |                                                                                                                                                                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no ti-lfa                                                                                                                                                                                                                                         |
| <b>Parameters</b> | <b>max-sr-frr-labels</b> [ <i>value</i> ] — Specifies the maximum number of labels that the TI-LFA backup next hop can use. The TI-LFA algorithm uses this value to limit the search for the Q node from the P node on the post-convergence path. |
| <b>Values</b>     | 0 to 3                                                                                                                                                                                                                                            |
| <b>Default</b>    | 2                                                                                                                                                                                                                                                 |

## node-protect

|                    |                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] node-protect</b>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>ospf>loopfree-alternates>ti-lfa<br>config>router>ospf3>loopfree-alternates>ti-lfa                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.<br><br>The <b>no</b> form of this command disables node-protect. |
| <b>Default</b>     | no node-protect                                                                                                                                                                                                                                                                                                                                                              |

## multicast-import

|                    |                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] multicast-import</b>                                                                                                                                                                           |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                              |
| <b>Description</b> | This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF.<br><br>The <b>no</b> form of the command disables the submission of routes into the multicast RTM. |
| <b>Default</b>     | no multicast-import                                                                                                                                                                                    |

## overload

|               |                                                         |
|---------------|---------------------------------------------------------|
| <b>Syntax</b> | <b>overload</b> [timeout seconds]<br><b>no overload</b> |
|---------------|---------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continues to reach the router.</p> <p>To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a <b>no overload</b> command is executed.</p> <p>If the <b>overload</b> command is encountered during the execution of an <a href="#">overload-on-boot</a> command then this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the <b>overload-on-boot</b> command is saved after the <b>overload</b> command. <b>However</b>, when <b>overload-on-boot</b> is configured under OSPF with no timeout value configured, the router will remain in overload state indefinitely after a reboot.</p> <p>Use the <b>no</b> form of this command to return to the default. When the <b>no overload</b> command is executed, the overload state is terminated regardless of the reason the protocol entered overload state.</p> |
| <b>Default</b>     | no overload                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>timeout seconds</i> — Specifies the number of seconds to reset overloading.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Values</b>      | 1 to 1800 in the <b>config&gt;router&gt;ospf</b> context<br>60 to 1800 in the <b>config&gt;router&gt;ospf3</b> context                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## overload-include-ext-2

|                    |                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] overload-include-ext-2</b>                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command is used to control if external type-2 routes should be re-advertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes will be advertised with the maximum metric. |
| <b>Default</b>     | no overload-include-ext-2                                                                                                                                                                                                                                                                                 |

## overload-include-stub

|                |                                   |
|----------------|-----------------------------------|
| <b>Syntax</b>  | <b>[no] overload-include-stub</b> |
| <b>Context</b> | config>router>ospf                |

---

```
config>router>ospf3
```

**Description** This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.

**Default** no overload-include-stub

## overload-on-boot

**Syntax** **overload-on-boot** [*timeout seconds*]  
**no overload**

**Context** config>router>ospf  
config>router>ospf3

**Description** When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- the timeout timer expires
- a manual override of the current overload state is entered with the **no overload** command

The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

The default timeout value is 60 seconds, which means after 60 seconds overload status the SR will recover (change back to non-overload status). However, when **overload-on-boot** is configured under OSPF with no timeout value the router will remain in overload state indefinitely after a reboot.

**Default** no overload-on-boot

**Parameters** *timeout seconds* — Specifies the number of seconds to reset overloading.

**Values** 1 to 1800 in the **config>router>ospf** context  
60 to 1800 in the **config>router>ospf3** context

## preference

**Syntax** **preference** *preference*  
**no preference**

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the preference for OSPF internal routes.</p> <p>A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference is used to decide which route will be used.</p> <p>Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in <a href="#">Table 15</a>. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.</p> <p>If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the <b>ecmp</b> in the <b>config&gt;router</b> context.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | preference 10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>preference</i> — The preference for internal routes expressed as a decimal integer.<br>Defaults for different route types are listed in <a href="#">Table 15</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 15** Route Preference Defaults by Route Type

| Route Type             | Preference | Configurable     |
|------------------------|------------|------------------|
| Direct attached        | 0          | No               |
| Static routes          | 5          | Yes              |
| OSPF internal          | 10         | Yes <sup>1</sup> |
| IS-IS level 1 internal | 15         | Yes              |
| IS-IS level 2 internal | 18         | Yes              |
| RIP                    | 100        | Yes              |
| OSPF external          | 150        | Yes              |
| IS-IS level 1 external | 160        | Yes              |
| IS-IS level 2 external | 165        | Yes              |
| BGP                    | 170        | Yes              |

Note:

1. Preference for OSPF internal routes is configured with the **preference** command.

**Values** 1 to 255

# reference-bandwidth

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>reference-bandwidth</b> <i>bandwidth-in-kbps</i><br><b>reference-bandwidth</b> [ <b>tbps</b> <i>Tera-bps</i> ] [ <b>gbps</b> <i>Giga-bps</i> ] [ <b>mbps</b> <i>Mega-bps</i> ] [ <b>kbps</b> <i>Kilo-bps</i> ]<br><b>no reference-bandwidth</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the reference bandwidth in kilobits per second (kb/s) that provides the reference for the default costing of interfaces based on their underlying link speed.</p> <p>The default interface cost is calculated as follows:</p> $\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$ <p>The default <i>reference-bandwidth</i> is 100,000,000 kb/s or 100 Gb/s, so the default auto-cost metrics for various link speeds are as follows:</p> <ul style="list-style-type: none"> <li>• 10 Mb/s link default cost of 10000</li> <li>• 100 Mb/s link default cost of 1000</li> <li>• 1 Gb/s link default cost of 100</li> <li>• 10 Gb/s link default cost of 10</li> </ul> <div>  <p><b>Note:</b> The default reference-bandwidth must be manually configured to a higher value if interface speeds are greater than 100 Gb/s, and metrics based on link speed are used. When the default reference-bandwidth is used, a metric of 1 is set on all interface speeds ≥ 100 Gb/s. For example, 100 GE, 100 GE LAG, 400 GE, and 400 GE LAG interfaces will all have a metric of 1.</p> </div> <p>The <b>reference-bandwidth</b> command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the <b>metric</b> <i>metric</i> command in the <b>config&gt;router&gt;ospf&gt;area&gt;interface</b> <i>ip-int-name</i> context.</p> <p>The <b>no</b> form of the command reverts the reference-bandwidth to the default value.</p> |
| <b>Default</b>     | reference-bandwidth 100000000                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>bandwidth-in-kbps</i> — The reference bandwidth in kilobits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4000000000</p> <p><i>tbps</i> <i>Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**gbps** *Giga-bps* — The reference bandwidth in gigabits per second expressed as a decimal integer.

**Values** 1 to 999

**mbps** *Mega-bps* — The reference bandwidth in megabits per second expressed as a decimal integer.

**Values** 1 to 999

**kbps** *Kilo-bps* — reference bandwidth in kilobits per second expressed as a decimal integer.

**Values** 1 to 999

## rib-priority

|                    |                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rib-priority</b> { <b>high</b> } <i>prefix-list-name</i><br><b>no rib-priority</b> { <b>high</b> }                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command enabled RIB prioritization for the OSPF protocol and specifies the prefix list that will be used to select the specific routes that should be processed through the OSPF route calculation process at a higher priority.<br><br>The <b>no</b> form of <b>rib-priority</b> command disables RIB prioritization at the associated level. |
| <b>Default</b>     | no rib-priority high                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <i>prefix-list-name</i> — specifies the prefix list, up to a maximum of 32 characters, which is used to select the routes that are processed at a higher priority through the route calculation process.                                                                                                                                            |

## router-id

|                    |                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>router-id</b> <i>ip-address</i><br><b>no router-id</b>                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the router ID for the OSPF instance. This command configures the router ID for the OSPF instance.<br><br>When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the <b>config&gt;router</b> context. |

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **config>router** context is not configured, the following applies:

- the system uses the system interface address (which is also the loopback address)
- if a system interface address is not configured, it uses the last 32 bits of the chassis MAC address

This is a **required** command when configuring multiple instances and the instance being configured is not the base instance.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

The **no** form of the command to reverts to the default value.

## rtr-adv-lsa-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rtr-adv-lsa-limit</b> <i>limit</i> [ <b>log-only</b> ] [ <b>threshold percent</b> ]<br><b>rtr-adv-lsa-limit</b> <i>limit</i> [ <b>log-only</b> ] [ <b>threshold percent</b> ] [ <b>overload-timeout</b> { <i>seconds</i>   <b>forever</b> }]<br><b>no rtr-adv-lsa-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command configures the maximum number of LSAs OSPF can learn from another router, in order to protect the system from a router that accidentally advertises a large number of LSAs. When the number of advertised LSAs reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, OSPF goes into overload.</p> <p>The overload-timeout option allows the administrator to control how long OSPF is in overload as a result of the advertised LSA limit being reached. At the end of this duration of time the system automatically attempts to restart OSPF. One possible value for the overload-timeout is forever, which means OSPF is never restarted automatically and this corresponds to the default behavior when the overload-timeout option is not configured.</p> <p>The <b>no</b> form of the command removes the <b>rtr-adv-lsa-limit</b>.</p> |
| <b>Default</b>     | no rtr-adv-lsa-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><b>log-only</b> — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, overload is not set.</p> <p><b>percent</b> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p><b>Values</b> 0 to 100</p> <p><b>limit</b> — The number of LSAs that can be learned expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4294967295</p> <p><b>second</b> — Specifies duration in minutes before restarting OSPF.</p> <p><b>Values</b> Values 1 to 1800</p> <p><b>forever</b> — Specifies that OSPF is restarted only after the <b>clear router ospf   ospf3 overload rtr-adv-lsa-limit</b> command is executed.</p> |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## segment-routing

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] segment-routing</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command enables the context to configure segment routing parameters within an IGP instance.</p> <p>Segment routing adds to IS-IS, OSPF, or OSPF3 routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface or next hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as a segment ID (SID).</p> <p>When segment routing is used together with the MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels.</p> <p>Segment routing using MPLS labels can be used in both shortest path routing applications and traffic engineering applications. This feature implements the shortest path forwarding application.</p> <p>After segment routing is successfully enabled in the IS-IS, OSPF, or OSPF3 instance, the router will perform the following operations:</p> <ul style="list-style-type: none"><li>• Advertise the Segment Routing Capability sub-TLV to routers in all areas or levels of the IGP instance. However, only neighbors with which the IGP instance established an adjacency will interpret the SID and label range information and use it for calculating the label to swap to or push for a particular resolved prefix SID.</li></ul> |

- Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node SID flag) set. The segment routing module then programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
- Automatically assign and advertise an adjacency SID label for each formed adjacency over a network IP interface in the new adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, in effect with a swap to an implicit null label operation, for each advertised adjacency SID.
- Resolve received prefixes, and if a prefix SID sub-TLV exists, the segment routing module programs the ILM with a swap operation and programs an LSP ID to NHLFE (LTN) with a push operation, both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in an IGP instance, the main SPF and LFA SPF are computed normally and the primary next hop and LFA backup next hop for a received prefix are added to the RTM without the label information advertised in the prefix SID sub-TLV.

## adjacency-set

|                    |                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] adjacency-set <i>id</i></b>                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>ospf>segment-routing<br>config>router>isis>segment-routing                                                                                                                                                                                              |
| <b>Description</b> | This command creates an adjacency set. An adjacency set consists of one or more adjacency SIDs originating on this node. The constituent adjacencies may terminate on different nodes.<br><br>The <b>no</b> form of this command removes the specified adjacency set. |
| <b>Parameters</b>  | <i>id</i> — Specifies an unsigned integer representing the identifier of the adjacency set.                                                                                                                                                                           |
| <b>Values</b>      | 1 to 4294967295                                                                                                                                                                                                                                                       |

## parallel

|                    |                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>parallel [no-advertise]<br/>no parallel</b>                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>ospf>segment-routing>adjacency-set<br>config>router>isis>segment-routing>adjacency-set                                                                                                                                                                                                                                                                             |
| <b>Description</b> | This command indicates that all members of the adjacency set must terminate on the same neighboring node. The system raises a trap if a user attempts to add an adjacency terminating on a neighboring node that differs from the existing members of the adjacency set. In addition, the system stops advertising the adjacency set in IS-IS or OSPF and locally deprograms it. |

By default, parallel adjacency sets are advertised in the IGP. The **no-advertise** option prevents an adjacency set from being advertised in the IGP. It is only allowed in CLI and SNMP if the **parallel** command is configured.

The **no** form of this command indicates that the adjacency set can include adjacencies to different next hop nodes.

**Default** parallel

sid

**Syntax** sid label *value*

**Context** config>router>ospf>segment-routing>adjacency-set  
config>router>isis>segment-routing>adjacency-set

**Description** This command allows a static SID value to be assigned to an adjacency set in IS-IS or OSPF segment routing.

The **label** option specifies the value is assigned to an MPLS label.

The **no** form of this command removes the adjacency SID.

**Parameters** label *value* — Specifies the value of adjacency SID label.

**Values** 18432 to 524287 | 1048575 (FP4 only)

adj-sid-hold

**Syntax** adj-sid-hold *seconds*  
no adj-sid-hold

**Context** config>router>ospf>segment-routing  
config>router>ospf3>segment-routing

**Description** This command configures a timer to hold the ILM or LTM of an adjacency SID following a failure of the adjacency.

When an adjacency to a neighbor fails, the IGP will withdraw the advertisement of the link TLV information as well as its adjacency SID sub-TLV. However, the LTN or ILM record of the adjacency SID must be kept in the data path to maintain forwarding using the LFA or remote LFA backup for sufficient length of time to allow the ingress LER and other routers that use this adjacency SID to activate a new path after the IGP converges.

If the adjacency is restored before the timer expires, the timer is aborted as soon as the new ILM or LTN records are updated with the new primary and backup NHLFE information.

The **no** form of the command removes the adjacency SID hold time.

|                   |                                                                    |
|-------------------|--------------------------------------------------------------------|
| <b>Default</b>    | adj-sid-hold 15                                                    |
| <b>Parameters</b> | <i>seconds</i> — Specifies the adjacency SID hold time, in seconds |
| <b>Values</b>     | 1 to 300                                                           |

## backup-node-sid

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>backup-node-sid</b> <i>ip-prefix/prefix-length index index</i><br><b>backup-node-sid</b> <i>ip-prefix/prefix-length label label</i><br><b>no backup-node-sid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>ospf>segment-routing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command enables LFA Protection using segment routing backup node SID.</p> <p>The objective of this feature is to reduce the label stack pushed in a LFA tunnel next hop of inter-area and inter-domain prefixes. This is applicable in MPLS deployments across multiple IGP areas or domains such in seamless MPLS design.</p> <p>The user enables the feature by configuring a backup node SID at an ABR/ASBR that is acting as a backup to the primary exit ABR/ASBR of inter-area or inter-as routes learned as BGP labeled routes. The user can enter either a label or an index for the backup node SID.</p> <p>When a node in a IGP domain resolves a BGP label route for an inter-area or inter-domain prefix via the primary ABR exit router, it will use the backup node SID of this router, which is advertised by the backup ABR/ASBR, as the LFA backup instead of the SID to the remote LFA PQ node to save on the pushed label stack</p> <p>This feature only allows the configuration of a single backup node SID per IGP instance and per ABR/ASBR. In other words, only a pair of ABR/ASBR nodes can back up each other in an IGP domain. Each time the user invokes the above command within the same IGP instance, it will override any previous configuration of the backup node SID. The same ABR/ASBR can, however, participate in multiple IGP instances and provide backup support within each instance.</p> |
| <b>Default</b>     | no backup-node-sid                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>ip-prefix/prefix-length</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.<br><b>Values</b> ip-prefix/mask: <ul style="list-style-type: none"> <li>• ip-prefix a.b.c.d (host bits must be 0)</li> </ul> ipv6-prefix: <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

---

- d: [0 to 255]D

prefix-length: 0 to 128

*index* — specifies the index for this backup node SID

**Values** 0 to 4294967295

*label* — specifies the SID value for this backup node SID.

**Values** 1 to 4294967295

## entropy-label

|                    |                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>entropy-label {force-disable   enable}</b><br><b>no entropy-label</b>                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>isis>segment-routing<br>config>router>ospf>segment-routing                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command instructs the system to ignore any received IGP advertisements of entropy label capability relating to remote nodes in the network. It also prevents a user from configuring <b>override-tunnel-eltc</b> for the IGP instance.</p> <p>The <b>no</b> version of this command enables the processing of any received IGP advertisements of entropy label capability.</p> |
| <b>Default</b>     | entropy-label enable                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><b>force-disable</b> — Forces the system to ignore any received entropy label capability signaled in the IGP.</p> <p><b>enable</b> — Enables the system to process any received entropy label capability signaled in the IGP.</p>                                                                                                                                                   |

## export-tunnel-table

|                    |                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-tunnel-table ldp</b><br><b>no export-tunnel-table</b>                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ospf>segment-routing                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command enables exporting, to an IGP instance, the LDP tunnels for the purpose of stitching a SR tunnel to a LDP FEC for the same destination IPv4 /32 prefix.</p> <p>In the SR-to-LDP data path direction, the SR mapping server provides a global policy for the prefixes corresponding to the LDP FECs the SR stitches to.</p> |

When this command is enabled in the segment-routing context of an IGP instance, IGP listens to LDP tunnel entries in the TTM. Whenever a LDP tunnel destination matches a prefix for which IGP received a prefix-SID sub-TLV from a mapping server, it instructs the SR module to program the SR ILM and to stitch it to the LDP tunnel endpoint. The LDP FEC can be resolved via a static route, a IS-IS instance, or an OSPF instance.

When an SR tunnel is stitched to a LDP FEC, packets forwarded will benefit from the protection of the LFA backup next hop of the LDP FEC.

When resolving a node SID, IGP will prefer resolution of prefix SID received in a IP Reach TLV over a prefix SID received via the mapping server. In other words the swapping of the SR ILM to a SR NHLFE is preferred over stitching it to a LDP tunnel endpoint.

It is recommended to enable the **bfd-enable** option on the interfaces in both LDP and IGP instance contexts to speed up the failure detection and the activation of the LFA/remote-LFA backup next hop in either direction of the stitching

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

The **no** form of this command disables the exporting of LDP tunnels to the IGP instance.

## mapping-server

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mapping-server</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ospf>segment-routing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the context for the Segment Routing mapping server feature in an OSPF instance.</p> <p>The mapping server feature allows the configuration and advertisement in OSPF of the node SID index for OSPF prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server and using a prefix-SID sub-TLV within the Extended Prefix Range TLV in OSPF.</p> <p>The <b>no</b> form of this command deletes all node SID entries in the OSPF instance.</p> |

## sid-map

|                |                                                                                                                                                                                                                                                                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>sid-map node-sid index <i>index-value</i> [range <i>range-value</i>] prefix <i>ip-address/mask</i> [<i>netmask</i>]</b><br><b>sid-map node-sid index <i>index-value</i> [range <i>range-value</i>] prefix <i>ip-address/mask</i> [<i>netmask</i>]</b><br><b>scope {<i>area area-id</i>   as}</b><br><b>no sid-map node-sid index <i>index-value</i></b> |
| <b>Context</b> | config>router>ospf>segment-routing>mapping-server                                                                                                                                                                                                                                                                                                          |

**Description** This command configures the Segment Routing mapping server database in OSPF.

The user enters the node SID index for one or a range of prefixes by specifying the first index value and optionally a range value. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. If the user enters the first prefix with a mask lower than 32, the OSPF Extended Prefix Range TLV is advertised but a router which receives it will not resolve SID and instead originates a trap.

The user specifies the mapping server own flooding scope for the generated OSPF Extended Prefix Range TLV using the scope option. There is no default value. If the scope is a specific area, then the TLV is flooded only in that area.

An ABR that propagates an intra-area OSPF Extended Prefix Range TLV flooded by the mapping server in that area into other areas, sets the inter-area flag (IA-flag). The ABR also propagates the TLV if received with the inter-area flag set from other ABR nodes but only from the backbone to leaf areas and not vice-versa. However, if the exact same TLV is advertised as an intra-area TLV in a leaf area, the ABR will not flood the inter-area TLV into that leaf area.



**Note:** SR OS does not leak this TLV between OSPF instances.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix, or range of prefixes, a prefix-SID sub-TLV within a OSPF Extended Prefix Range TLV in that instance. The flooding scope of the TLV from the mapping server is determined as previously explained. No further check of the reachability of that prefix in the mapping server route table is performed and no check if the SID index is duplicate with some existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

The **no** form of the command deletes the range of node SIDs beginning with the specified index value.

**Default** no prefix-sid-range

**Parameters** **index** *index-value* — Specifies the index

**Values** 0 to 4294967295

**range** *range-value* — Specifies the range

**Values** 1 to 65535

**prefix** *ip-address/mask* — The IP address in dotted decimal notation

**Values** ip-address/mask:

• ip-address a.b.c.d (host bits must be 0)

mask: 0 to 132

*netmask* — Specifies the netmask.

**Values** netmask — a.b.c.d (network bits all 1 and host bits all 0)

**area** *area-id* — configures the mapping server own flooding scope for the generated OSPF Extended Prefix Range TLV

**Values** ip-address | 0 to 4294967295

## prefix-sid-range

|                    |                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefix-sid-range global</b><br><b>prefix-sid-range start-label label-value max-index index-value</b><br><b>no prefix-sid-range</b> |
| <b>Context</b>     | config>router>ospf>segment-routing<br>config>router>ospf3>segment-routing                                                             |
| <b>Description</b> | This command configures the prefix SID index range and offset label value for an IGP instance.                                        |

The key parameter is the configuration of the prefix SID index range and the offset label value that this IGP instance will use. Because each prefix SID represents a network global IP address, the SID index for a prefix must be unique network-wide. Therefore, all routers in the network are expected to configure and advertise the same prefix SID index range for an IGP instance. However, the label value used by each router to represent this prefix, that is, the label programmed in the ILM, can be local to that router by the use of an offset label, referred to as a start label:

*Local Label (Prefix SID) = start-label + {SID index}*

The label operation in the network is very similar to LDP when operating in independent label distribution mode (RFC 5036, *LDP Specification*), with the difference being that the label value used to forward a packet to each downstream router is computed by the upstream router based on the advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the **global** mode of operation, the user configures the global value and this IGP instance will assume the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. Once one IGP instance selected the global option for the prefix SID range, all IGP instances on the system will be restricted to do the same. The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. Once the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change will be failed if an already allocated SID index/label goes out of range.



In per-instance mode, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index) must be within the SRGB or the configuration will fail. The 7750 SR checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce no overlapping of these ranges. The user must shut down the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. A range change will fail if an already allocated SID index/label goes out of range. The user can change the SRGB without shutting down the segment routing context as long as it does not reduce the current per-IGP instance SID index/label range defined with the **prefix-sid-range** command. Otherwise, shut down the segment routing context of the IGP instance, and disable and re-enable the **prefix-sid-range** command.

|                   |                                                                                                                                                                                                                                                      |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no prefix-sid-range                                                                                                                                                                                                                                  |
| <b>Parameters</b> | <i>label-value</i> — Specifies the label offset for the SR label range of this IGP instance<br><b>Values</b> 0 to 524287<br><i>index-value</i> — Specifies the maximum value of the prefix SID index range for this IGP<br><b>Values</b> 1 to 524287 |

## srlb

|                    |                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>srlb</b> <i>reserved-label-block-name</i><br><b>no srlb</b>                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>ospf>segment-routing<br>config>router>isis>segment-routing                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command specifies the reserved label block to use for the Segment Routing Local Block (SRLB) for the specified IS-IS or OSPF instance. The named reserved label block must already have been configured under <b>configure&gt;router&gt;mpls&gt;mpls-labels</b> .<br><br>The <b>no</b> form of the command removes an SRLB. |
| <b>Parameters</b>  | <i>reserved-label-block-name</i> — Specifies the name of the reserved label block, up to 64 characters.                                                                                                                                                                                                                          |

## tunnel-mtu

|                    |                                                                             |
|--------------------|-----------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-mtu</b> <i>bytes</i><br><b>no tunnel-mtu</b>                      |
| <b>Context</b>     | config>router>ospf>segment-routing<br>config>router>ospf3>segment-routing   |
| <b>Description</b> | This command configures the MTU of all SR tunnels within each IGP instance. |

The MTU of a SR tunnel populated into the TTM is determined as the same as an IGP tunnel; for example, for an LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add, at most, one more label to the tunnel for a total of two labels. There is no default value for this command. If the user does not configure an SR tunnel MTU, the MTU will be determined by IGP as follows:

The MTU of the SR tunnel in bytes is then determined as follows:

$$SR\_Tunnel\_MTU = MIN \{ Cfg\_SR\_MTU, IGP\_Tunnel\_MTU - (1 + frr\_overhead) \times 4 \}$$

Where:

- *Cfg\_SR\_MTU* is the MTU configured by the user for all SR tunnels within an IGP instance using the **tunnel-mtu** command. If no value is configured by the user, the SR tunnel MTU is determined by the IGP interface calculation explained in the next bullet point.
- *IGP\_Tunnel\_MTU* is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel
- *frr-overhead* is set to 1 if the **segment-routing** and **remote-lfa** options are enabled in the IGMP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated whenever any of the above parameters used in its calculation changes. This includes if the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.

|                   |                                                       |
|-------------------|-------------------------------------------------------|
| <b>Default</b>    | no tunnel-mtu                                         |
| <b>Parameters</b> | <i>bytes</i> — Specifies the size of the MTU in bytes |
| <b>Values</b>     | 512 to 9198                                           |

## tunnel-table-pref

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-table-pref</b> <i>preference</i><br><b>no tunnel-table-pref</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ospf>segment-routing<br>config>router>ospf3>segment-routing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command configures the TTM preference of shortest path SR tunnels created by the IGP instance. This is used for BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the tunnel binding commands are configured to the <b>any</b> value, which parses the TTM for tunnels in the protocol preference order. The user can choose to either accept the global TTM preference or explicitly list the tunnel types they want to use. If the user lists the tunnel types explicitly, the TTM preference is still used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected type fails. A reversion to a more preferred tunnel type is performed as soon as one is available. |

The segment routing module adds to the TTM an SR tunnel entry for each resolved remote node SID prefix and programs the data path having the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the value of the default preference for the various tunnel types. This includes the preference of SR tunnels based on shortest path (referred to as SR-ISIS and SR-OSPF).



**Note:** The preference of an SR-TE LSP is not configurable and is the second most preferred tunnel type after RSVP-TE. The preference is the same whether if the SR-TE LSP was resolved in IS-IS or OSPF.

The global default TTM preference for the tunnel types is as follows:

- ROUTE\_PREF\_RSVP 7
- ROUTE\_PREF\_SR\_TE 8
- ROUTE\_PREF\_LDP 9
- ROUTE\_PREF\_OSPF\_TTM 10
- ROUTE\_PREF\_ISIS\_TTM 11
- ROUTE\_PREF\_BGP\_TTM 12
- ROUTE\_PREF\_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of a SR tunnel in this case will be based on the lowest IGP instance ID. Similarly, IPv6 SR-ISIS and SR-OSPF3 tunnels are programmed into TTMv6 with the same default preference value as IPv4 SR-ISIS and IPv4 SR-OSPF respectively.

|                   |                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | tunnel-table-pref 10                                                                                                       |
| <b>Parameters</b> | <i>preference</i> — Specifies the integer value to represent the preference of IS-IS, OSPF, or OSPF3 SR tunnels in the TTM |
| <b>Values</b>     | 1 to 255                                                                                                                   |

timers

|                |                                           |
|----------------|-------------------------------------------|
| <b>Syntax</b>  | <b>timers</b>                             |
| <b>Context</b> | config>router>ospf<br>config>router>ospf3 |

**Description** This command enables the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network re-convergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase re-convergence time.

## incremental-spf-wait

**Syntax** **incremental-spf-wait** *inc-spf-wait*  
**no incremental-spf-wait**

**Context** config>router>ospf>timers  
config>router>ospf3>timers

**Description** This command sets the delay before an incremental SPF calculation is performed when LSA types 3, 4, 5, or 7 are received. This allows multiple updates to be processed in the same SPF calculation. Type 1 or type 2 LSAs are considered a topology change and will always trigger a full SPF calculation.

The **no incremental-spf-wait** form of the command resets the timer value back to the default value.



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default** incremental-spf-wait 1000

**Parameters** *inc-spf-wait* — Specifies the OSPF incremental SPF calculation delay in milliseconds.

**Values** 0 to 1000

## lsa-accumulate

**Syntax** **lsa-accumulate** *lsa-accumulate*  
**no lsa-accumulate**

**Context** config>router>ospf>timers  
config>router>ospf3>timers

**Description** This command sets the internal OSPF delay to allow for the accumulation of multiple LSA so OSPF messages can be sent as efficiently as possible. The **lsa-accumulate** timer applies to all LSAs except Type 1 and Type 2 LSAs, which are sent immediately. LSAs are accumulated and then sent when:

- its size reaches the MTU size of the interface
- a new LSA is received on the interface
- the **lsa-accumulate** timer expires

Shorting this delay can speed up the advertisement of LSAs to OSPF neighbors but may increase the number of OSPF messages sent.



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default** lsa-accumulate 1000

**Parameters** *lsa-accumulate* — Specifies the LSA accumulation delay in milliseconds.

**Values** 0 to 1000

## lsa-arrival

**Syntax** **lsa-arrival** *lsa-arrival-time*  
**no lsa-arrival**

**Context** config>router>ospf>timers  
config>router>ospf3>timers

**Description** This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.

It is recommended that the neighbors configured **lsa-generate** *lsa-second-wait* interval is equal or greater than the **lsa-arrival** timer configured here.

Use the **no** form of this command to return to the default.



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default** lsa-arrival 1000

**Parameters**    *lsa-arrival-time* — Specifies the timer in milliseconds.

**Values**        0 to 600000

lsa-generate

**Syntax**        **lsa-generate** *max-lsa-wait* [**lsa-initial-wait** *lsa-initial-wait* [**lsa-second-wait** *lsa-second-wait*]]

**no lsa-generate**

**Context**        config>router>ospf>timers  
config>router>ospf3>timers

**Description**    This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

Use the **no** form of this command to return to the default.



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default**        lsa-generate 5000

**Parameters**    *max-lsa-wait* — Specifies the maximum interval in milliseconds between two consecutive occurrences of an LSA being generated.

**Values**        10 to 600000

**Default**        5000

*lsa-initial-wait* — Specifies the first waiting period between link-state advertisements (LSA) originate(s), in milliseconds. When the LSA exceeds the **lsa-initial-wait** timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If, within the specified **lsa-initial-wait** period and another topology change occurs, then the **lsa-initial-wait** timer applies.

**Values**        10 to 600000

**Default**        5000

*lsa-second-wait* — Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time). This assumes that each failure occurs within the relevant wait period.

**Values** 10 to 600000  
**Default** 5000

redistribute-delay

|                    |                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>redistribute-delay</b> <i>redistribute-delay</i><br><b>no redistribute-delay</b>                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ospf>timers<br>config>router>ospf3>timers                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command sets the internal OSPF hold down timer for external routes being redistributed into OSPF.</p> <p>Shorting this delay can speed up the advertisement of external routes into OSPF but can result in additional OSPF messages if that source route is not yet stable.</p> <p>The <b>no</b> form of the command resets the timer value back to the default value.</p> |



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

|                   |                                                                                                                                                                         |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | redistribute-delay 1000                                                                                                                                                 |
| <b>Parameters</b> | <i>redistribute-delay</i> — Specifies the OSPF redistribution hold down time in milliseconds for external routes being advertised into OSPF.<br><b>Values</b> 0 to 1000 |

spf-wait

|                |                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>spf-wait</b> <i>max-spf-wait</i> [ <b>spf-initial-wait</b> <i>spf-initial-wait</i> [ <b>spf-second-wait</b> <i>spf-second-wait</i> ]]<br><b>no spf-wait</b> |
| <b>Context</b> | config>router>ospf>timers<br>config>router>ospf3>timers                                                                                                        |

**Description** This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default** `spf-wait 10000`

**Parameters** *max-spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

**Values** 10 to 120000

**Default** 10000

*spf-initial-wait* — Specifies the initial SPF calculation delay in milliseconds after a topology change.

**Values** 10 to 100000

**Default** 1000

*spf-second-wait* — Specifies the hold time in milliseconds between the first and second SPF calculation.

**Values** 10 to 100000

**Default** 1000

traffic-engineering

**Syntax** `[no] traffic-engineering`

**Context** `config>router>ospf`

**Description** This command enables the advertisement of the traffic engineering information for the router and its links.



Traffic engineering enables the router to perform route calculations constrained by nodes or links. The traffic engineering of this router are limited to calculations based on link and nodal constraints.

The **no** form of the command disables the advertisement of the traffic engineering information.

**Default** no traffic-engineering

## unicast-import-disable

**Syntax** [no] unicast-import-disable

**Context** config>router>ospf  
config>router>ospf3

**Description** This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured.

**Default** no unicast-import-disable

### 3.14.2.3 OSPF Area Commands

## area

**Syntax** [no] area *area-id*

**Context** config>router>ospf  
config>router>ospf3

**Description** This command creates the context to configure an OSPF or OSPF3 area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, and address-ranges and so on, that are currently assigned to this area.

**Default** no area

---

|                   |                                                                                                        |
|-------------------|--------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer. |
| <b>Values</b>     | 0.0.0.0 to 255.255.255.255 (dotted decimal), 0 to 4294967295 (decimal integer)                         |

## adjacency-set

|                    |                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] adjacency-set <i>id</i></b>                                                                                                                                                                                      |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>isis>interface                                                                                                                                                        |
| <b>Description</b> | This command associates an interface with an adjacency set. The adjacency set must have been defined under the IS-IS or OSPF segment-routing context.<br><br>The <b>no</b> form of this command removes the association. |
| <b>Parameters</b>  | <i>id</i> — Specifies an unsigned integer representing the identifier of the adjacency set.                                                                                                                              |
| <b>Values</b>      | 1 to 4294967295                                                                                                                                                                                                          |

## adjacency-sid

|                    |                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adjacency-sid label <i>value</i></b><br><b>no adjacency-sid</b>                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ospf>area>interface                                                                                                                                                                                                                         |
| <b>Description</b> | This command allows a static value to be assigned to an adjacency SID in OSPF segment routing.<br><br>The <b>label</b> option specifies that the value is assigned to an MPLS label.<br><br>The <b>no</b> form of this command removes the adjacency SID. |
| <b>Parameters</b>  | <b>label <i>value</i></b> — Specifies the value of adjacency SID label.                                                                                                                                                                                   |
| <b>Values</b>      | 18432 to 52428   1048575 (FP4 only)                                                                                                                                                                                                                       |

## advertise-router-capability

|                |                                                                                                                                |
|----------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] advertise-router-capability</b>                                                                                        |
| <b>Context</b> | config>router>ospf>area<br>config>router>ospf>area>interface<br>config>router>ospf3>area<br>config>router>ospf3>area>interface |

**Description** This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF stub router support: yes
- OSPF traffic engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF experimental TE: no

The **no** form of this command disables this capability.

**Default** advertise-router-capability

## area-range

**Syntax** **area-range** *ip-prefix/mask* [**advertise** | **not-advertise**]  
**no area-range** *ip-prefix/mask*

**Context** config>router>ospf>area  
config>router>ospf>area>nssa

**Description** This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of the command deletes the range (non) advertisement.

**Default** no area-range

**Special Cases** **NSSA Context** — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

**Area Context** — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

**Parameters** *ip-prefix* — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

**Values** ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0)

*mask* — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

**Values** 0 to 32 (mask length), 0.0.0.0 to 255.255.255.255 (dotted decimal)

**advertise** | **not-advertise** — Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

**Default** advertise

## area-range

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>area-range</b> <i>ipv4-prefix/mask</i>   <i>ipv6-prefix/prefix-length</i> [ <b>advertise</b>   <b>not-advertise</b> ]<br><b>no area-range</b> <i>ipv4-prefix/mask</i>   <i>ipv6-prefix/prefix-length</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>       | config>router>ospf3>area<br>config>router>ospf3>area>nssa                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>   | <p>This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.</p> <p>ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.</p> <p>The <b>no</b> form of the command deletes the range (non) advertisement.</p> |
| <b>Default</b>       | no area-range                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Special Cases</b> | <p><b>NSSA Context</b> — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.</p> <p><b>Area Context</b> — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>    | <p><i>ip-prefix/prefix-length</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.</p> <p><b>Values</b> ip-prefix/mask:</p> <ul style="list-style-type: none"> <li>• ip-prefix a.b.c.d (host bits must be 0)</li> </ul> <p>ipv6-prefix:</p> <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> </ul>                                                                                                                                                                                                                                                               |

• d: [0 to 255]D

prefix-length: 0 to 128

**advertise** | **not-advertise** — Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.

**Default** advertise

## blackhole-aggregate

**Syntax** [no] **blackhole-aggregate**

**Context** config>router>ospf>area  
config>router>ospf3>area

**Description** This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

When performing area aggregation, addresses may be included in the range for which no actual route exists, which can cause routing loops. To avoid this problem, configure the **blackhole-aggregate** option.

The **no** form of this command removes this option.

**Default** blackhole-aggregate

## database-export-exclude

**Syntax** [no] **database-export-exclude**

**Context** config>router>ospf>area

**Description** This command allows the user to prune the IGP link-state information of a specific OSPF level or OSPF area from being exported into the extended TE-DB.

The **no** form of the command returns to the default behavior inherited from the database-export command at the IS-IS or OSPF instance level.

**Default** no database-export-exclude

## export

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

---

|                    |                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf>area<br>config>router>ospf3>area                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures ABR export policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being exported into an area.</p> <p>This command cannot be used in OSPF area 0.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>                      |
| <b>Default</b>     | no export                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><i>policy-name</i> — Specifies up to five export route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p>The specified names must already be defined.</p> |

## extended-lsa

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>extended-lsa {inherit   only}</b><br><b>no extended-lsa</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>ospf3>area                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures the use of extended LSA format in a OSPFv3 area as per <i>draft-ietf-ospf-ospfv3-lsa-extend</i>.</p> <p>By default, the area inherits the instance level configuration. The latter defaults to <b>sparse</b> mode of operation. The <b>extended-lsa only</b> value enables the full extended LSA mode and this will cause all existing and new LSAs to use the extended LSA format.</p> <p>The OSPFv3 instance must first be shut down before the user can change the mode of operation since the protocol must flush all LSAs and re-establish all adjacencies.</p> <p>The <b>no</b> form at the area level returns the area into the default mode of inheriting the mode from the OSPFv3 instance level.</p> |
| <b>Default</b>     | extended-lsa inherit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><b>inherit</b> — Configures the area to inherit the mode of operation enabled at the OSPFv3 instance level.</p> <p><b>only</b> — Enables the full extended LSA mode of operation in an OSPFv3 area.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## import

|                    |                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] import</b> <i>policy-name</i> [ <i>policy-name</i> ]                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ospf>area<br>config>router>ospf3>area                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command configures ABR import policies to filter OSPFv2 Type 3 Summary-LSAs or OSPFv3 Inter-Area-Prefix-LSA between areas, in order to only permit the specified routes from being imported into an area.<br><br>This command cannot be used in OSPF area 0.<br><br>The <b>no</b> form of the command reverts to the default value.                      |
| <b>Default</b>     | no import                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>policy-name</i> — Specifies up to five import route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.<br><br>The specified names must already be defined. |

## key-rollover-interval

|                    |                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>key-rollover-interval</b> <i>seconds</i>                                                                            |
| <b>Context</b>     | config>router>ospf3>area                                                                                               |
| <b>Description</b> | This command configures the key rollover interval.                                                                     |
| <b>Default</b>     | key-rollover-interval 10                                                                                               |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the time, in seconds, after which a key rollover will start.<br><br><b>Values</b> 10 to 300 |

## lfa-policy-map

|                    |                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lfa-policy-map route-nh-template</b> <i>template-name</i><br><b>no lfa-policy-map</b>                  |
| <b>Context</b>     | config>router>ospf>area>interface                                                                         |
| <b>Description</b> | This command specifies the next hop template to be applied to those prefixes which primary next-hops use. |
| <b>Parameters</b>  | <i>template-name</i> — Specifies a template name up to 32 characters.                                     |

## loopfree-alternate-exclude

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>loopfree-alternate-exclude prefix-policy</b> <i>prefix-policy</i> [ <i>prefix-policy</i> ...(up to 5 max)]<br><b>no loopfree-alternate-exclude</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ospf<br>config>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command excludes from the LFA SPF calculation those prefixes that match a prefix entry or a tag entry in a prefix policy. If a prefix is excluded from LFA, it is not included in LFA calculations regardless of its priority. The prefix tag will, however, be used in the main SPF.</p> <p>The implementation also allows the user to exclude a specific interface in IS-IS or OSPF, a or all interfaces in an OSPF area or IS-IS level from the LFA SPF.</p> <p> <b>Note:</b> Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.</p> <p>The default action of the <b>loopfree-alternate-exclude</b> command, when not explicitly specified by the user in the prefix policy, is “reject”. Therefore, regardless of whether the user explicitly added the statement “default-action reject” to the prefix policy, a prefix that does not match any entry in the policy will be accepted into LFA SPF.</p> <p>The <b>no</b> form deletes the exclude prefix policy.</p> |
| <b>Default</b>     | no loopfree-alternate-exclude                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><i>prefix-policy</i> — specifies the name of the prefix policy, up to a maximum of five policies. The specified policy name must already be defined. Prefix policies are created with the command <b>config&gt;router&gt;policy-options&gt;prefix-list</b> (for information on prefix lists, refer to the 7705 SAR Router Configuration Guide, “Route Policies”).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## loopfree-alternate-exclude

|                    |                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] loopfree-alternate-exclude</b>                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>ospf>area<br>config>router>ospf>area>interface<br>config>router>ospf3>area<br>config>router>ospf3>area>interface                                                                                                                           |
| <b>Description</b> | <p>This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.</p> |



When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command re-instates the default value for this command.

**Default** no loopfree-alternate-exclude

## nssa

**Syntax** [no] nssa

**Context** config>router>ospf>area  
config>router>ospf3>area

**Description** This command creates the context to configure an OSPF or OSPF3 Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF or OSPF3 domain.

Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of the command removes the NSSA designation and configuration context from the area.

**Default** no nssa

## originate-default-route

**Syntax** originate-default-route [type-7] [no-adjacency-check]  
originate-default-route [type-nssa] [no-adjacency-check]  
no originate-default-route

**Context** config>router>ospf>area>nssa  
config>router>ospf3>area>nssa

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR).</p> <p>When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.</p> <p>The <b>no</b> form of the command disables origination of a default route.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>     | no originate-default-route                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><b>type-7</b> — Specifies a type 7 LSA should be used for the default route in the <b>config&gt;router&gt;ospf&gt;area&gt;nssa</b> context.</p> <p>Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.</p> <p>To revert to a type 3 LSA, enter <b>originate-default-route</b> without the <b>type-7</b> parameter.</p> <p><b>Default</b> Type 3 LSA default route.</p> <p><b>type-nssa</b> — Specifies an NSSA-LSA type should be used for the default route in the <b>config&gt;router&gt;ospf3&gt;area&gt;nssa</b> context.</p> <p><b>no-adjacency-check</b> — Specifies whether or not adjacency checks are performed before originating a default route. If this parameter is configured, then no area 0 adjacency is required for the ABR to advertise the default route.</p> <p><b>Default</b> Adjacency checks are performed, and an area 0 adjacency is required for the ABR to advertise the default route.</p> |

## redistribute-external

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] redistribute-external</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>ospf>area>nssa<br>config>router>ospf3>area>nssa                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.</p> <p>NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF or OSPF3 areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF or OSPF3 domain.</p> <p>The <b>no</b> form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.</p> |
| <b>Default</b>     | redistribute-external                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## summaries

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] summaries</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>ospf>area>nssa<br>config>router>ospf>area>stub<br>config>router>ospf3>area>nssa<br>config>router>ospf3>area>stub                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR).</p> <p>This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or NSSA area. (Default: summary)</p> <p>By default, summary route advertisements are sent into the stub area or NSSA.</p> <p>The <b>no</b> form of the command disables sending summary route advertisements and, for stub areas; only the default route is advertised by the ABR.</p> |
| <b>Default</b>     | summaries                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## stub

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] stub</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>ospf>area<br>config>router>ospf3>area                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command enables access to the context to configure an OSPF or OSPF3 stub area and adds/removes the stub designation from the area.</p> <p>External routing information is not flooded into stub areas. All routers in the stub area must be configured with the <b>stub</b> command. An OSPF or OSPF3 area cannot be both an NSSA and a stub area.</p> <p>Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.</p> <p>By default, an area is not a stub area.</p> <p>The <b>no</b> form of the command removes the stub designation and configuration context from the area.</p> |
| <b>Default</b>     | no stub                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

---

## default-metric

|                    |                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-metric</b> <i>metric</i><br><b>no default-metric</b>                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ospf>area>stub<br>config>router>ospf3>area>stub                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command configures the metric used by the area border router (ABR) for the default route into a stub area.</p> <p>The default metric should only be configured on an ABR of a stub area.</p> <p>An ABR generates a default route if the area is a <b>stub</b> area.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | default-metric 1                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>metric</i> — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.<br><b>Values</b> 1 to 16777214                                                                                                                                                                                               |

### 3.14.2.4 Interface/Virtual Link Commands

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] interface</b> <i>ip-int-name</i> [ <b>secondary</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ospf>area<br>config>router>ospf3>area                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command creates a context to configure an OSPF interface.</p> <p>By default, interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.</p> <p>The <b>no</b> form of the command deletes the OSPF interface configuration for this interface. The <b>shutdown</b> command in the <b>config&gt;router&gt;ospf&gt;interface</b> context can be used to disable an interface without removing the configuration for the interface.</p> |
| <b>Default</b>     | no interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |


- 
- Parameters**    *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config>router>interface** and **config>service>ies>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
- If the IP interface name does not exist or does not have an IP address configured an error message will be returned.
- If the IP interface exists in a different area it will be moved to this area.
- secondary** — Allows multiple secondary adjacencies to be established over a single IP interface.

## virtual-link

- Syntax**        **[no] virtual-link** *router-id* **transit-area** *area-id*
- Context**        config>router>ospf>area  
                    config>router>ospf3>area
- Description**    This command configures a virtual link to connect area border routers to the backbone via a virtual link.
- The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in the picture below) then the area border routers (routers 1 and 2 in the picture below) must be connected via a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area. (area 0.0.0.1 in the picture below). A virtual link can only be configured while in the area 0.0.0.0 context.
- The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).
- The **no** form of the command deletes the virtual link.
- By default, no virtual link is defined.
- Default**        no virtual-link
- Parameters**    *router-id* — The router ID of the virtual neighbor in IP address dotted decimal notation.
- area-id* — The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in [Figure 13](#)) then the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (see Area 0.0.0.4).

**Figure 13** OSPF Areas



## advertise-subnet

|                    |                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <code>[no] advertise-subnet</code>                                                                                                                                             |
| <b>Context</b>     | <code>config&gt;router&gt;ospf&gt;area&gt;interface</code>                                                                                                                     |
| <b>Description</b> | This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes. |

The **no** form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

**Default** advertise-subnet

## auth-keychain

**Syntax** **auth-keychain**

**Context** config>router>ospf>area>interface  
config>router>ospf>area>virtual-link

**Description** This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

**Default** no auth-keychain

**Parameters** *name* — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

## authentication

**Syntax** **authentication bidirectional** *sa-name*  
**authentication** [**inbound** *sa-name* **outbound** *sa-name*]  
**no authentication**

**Context** config>router>ospf3>area>interface  
config>router>ospf3>area>virtual-link

**Description** This command configures the password used by the OSPF3 interface or virtual-link to send and receive OSPF3 protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication.

By default, no authentication key is configured.

The **no** form of the command removes the authentication.

**Default** no authentication

**Parameters** **bidirectional** *sa-name* — Specifies bidirectional OSPF3 authentication.  
**inbound** *sa-name* — Specifies the inbound security association (SA) name for OSPF3 authentication.  
**outbound** *sa-name* — Specifies the outbound SA name for OSPF3 authentication.

## authentication-key

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ]<br><b>no authentication-key</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf>area>virtual-link                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.</p> <p>All neighboring routers must use the same type of authentication and password for proper protocol communication. If the <b>authentication-type</b> is configured as <b>password</b>, then this key must be configured.</p> <p>By default, no authentication key is configured.</p> <p>The <b>no</b> form of the command removes the authentication key.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>     | no authentication-key                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p><i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the <b>hash2</b> encrypted variable cannot be copied and pasted. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> |

## authentication-type

|                |                                                                                                         |
|----------------|---------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>authentication-type</b> { <b>password</b>   <b>message-digest</b> }<br><b>no authentication-type</b> |
| <b>Context</b> | config>router>ospf>area>interface<br>config>router>ospf>area>virtual-link                               |



---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables authentication and specifies the type of authentication to be used on the OSPF interface.</p> <p>Both simple <b>password</b> and <b>message-digest</b> authentication are supported.</p> <p>By default, authentication is not enabled on an interface.</p> <p>The <b>no</b> form of the command disables authentication on the interface.</p>                                                                      |
| <b>Default</b>     | no authentication-type                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><b>password</b> — This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple <b>password</b> authentication is enabled.</p> <p><b>message-digest</b> — This keyword enables message digest MD5 authentication in accordance with RFC1321. If this option is configured, then at least one message-digest-key must be configured.</p> |

## bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bfd-enable [remain-down-on-failure]</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | <pre>config&gt;router&gt;ospf&gt;area&gt;interface config&gt;router&gt;ospf3&gt;area&gt;interface</pre>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.</p> <p>The <b>no</b> form of this command removes BFD from the associated OSPF protocol adjacency.</p> |
| <b>Default</b>     | no bfd-enable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><b>remain-down-on-failure</b> — If the BFD session does not come back up within 10 seconds, then OSPF will bring down the adjacency and wait on BFD again. This can cause OSPF neighbors to flap, because OSPF will form the adjacency and then bring it down if the BFD session is still down. If this parameter is not configured, the OSPF adjacency will form even if the BFD adjacency does not come back up after a failure.</p>                                                       |

## dead-interval

|                |                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <p><b>dead-interval</b> <i>seconds</i></p> <p><b>no dead-interval</b></p>                                 |
| <b>Context</b> | <pre>config&gt;router&gt;ospf&gt;area&gt;interface config&gt;router&gt;ospf&gt;area&gt;virtual-link</pre> |

---

|                      |                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | <pre>config&gt;router&gt;ospf3&gt;area&gt;interface config&gt;router&gt;ospf3&gt;area&gt;virtual-link</pre>                                                                                                                                                                                                                                                               |
| <b>Description</b>   | <p>This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>       | dead-interval 40                                                                                                                                                                                                                                                                                                                                                          |
| <b>Special Cases</b> | <p><b>OSPF Interface</b> — If the <b>dead-interval</b> configured applies to an interface, then all nodes on the subnet must have the same dead interval.</p> <p><b>Virtual Link</b> — If the <b>dead-interval</b> configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.</p>           |
| <b>Parameters</b>    | <p><i>seconds</i> — The dead interval expressed in seconds.</p> <p><b>Values</b> 1 to 65535</p>                                                                                                                                                                                                                                                                           |

## hello-interval

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <pre>hello-interval seconds no hello-interval</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>       | <pre>config&gt;router&gt;ospf&gt;area&gt;interface config&gt;router&gt;ospf&gt;area&gt;virtual-link config&gt;router&gt;ospf3&gt;area&gt;interface config&gt;router&gt;ospf3&gt;area&gt;virtual-link</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>   | <p>This command configures the interval between OSPF hellos issued on the interface or virtual link.</p> <p>The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.</p> <p>Reducing the interval, in combination with an appropriate reduction in the associated <b>dead-interval</b>, allows for faster detection of link and/or router failures at the cost of higher processing costs.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |
| <b>Default</b>       | hello-interval 10                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Special Cases</b> | <p><b>OSPF Interface</b> — If the <b>hello-interval</b> configured applies to an interface, then all nodes on the subnet must have the same hello interval.</p> <p><b>Virtual Link</b> — If the <b>hello-interval</b> configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.</p>                                                                                                                                                                                                                             |

---

|                   |                                                                                  |
|-------------------|----------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>seconds</i> — The hello interval, in seconds, expressed as a decimal integer. |
| <b>Values</b>     | 1 to 65535                                                                       |

## interface-type

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>interface-type {broadcast   point-to-point  non-broadcast}</b><br><b>no interface-type</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>       | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>   | <p>This command configures the interface type to be one of broadcast, point-to-point, or non-broadcast.</p> <p>Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead of the Ethernet link provided the link is used as point-to-point.</p> <p>If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.</p> <p>The <b>no</b> form of the command returns the setting to the default value.</p> |
| <b>Default</b>       | <p>interface-type point-to-point (if the physical interface is SONET)</p> <p>interface-type broadcast (if the physical interface is Ethernet or unknown)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Special Cases</b> | <b>Virtual-Link</b> — A virtual link is always regarded as a point-to-point interface and not configurable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>    | <p><b>broadcast</b> — Configures the interface as a broadcast network. To significantly improve adjacency forming and network convergence, configure a network as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.</p> <p><b>point-to-point</b> — Configures the interface as a point-to-point link.</p> <p><b>non-broadcast</b> — Configures the interface as a non-broadcast network.</p>                                                                                                                                                                                         |

## lfa-policy-map

|                    |                                                                                          |
|--------------------|------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lfa-policy-map route-nh-template <i>template-name</i></b><br><b>no lfa-policy-map</b> |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                  |
| <b>Description</b> | This command applies a route next hop policy template to an OSPF or IS-IS interface.     |

When a route next hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next hop policy template to an OSPF or IS-IS interface.

|                   |                                                                                 |
|-------------------|---------------------------------------------------------------------------------|
| <b>Default</b>    | no lfa-policy-map                                                               |
| <b>Parameters</b> | <i>template-name</i> — Specifies the name of the template, up to 32 characters. |

## lsa-filter-out

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsa-filter-out</b> [ <b>all</b>   <b>except-own-rtrlsa</b>   <b>except-own-rtrlsa-and-defaults</b> ]<br><b>no lsa-filter-out</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:</p> <ul style="list-style-type: none"> <li>• Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next hop.</li> <li>• Flood the router's own router-LSA out the interface and suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and has a statically configured default route with a loopback or system interface address (contained in the router-LSA) as next hop.</li> <li>• Flood the router's own router-LSA and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route.</li> </ul> <p>The <b>no</b> form of this command disables OSPF LSA filtering (normal operation).</p> |
| <b>Default</b>     | no lsa-filter-out                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## message-digest-key

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>message-digest-key</b> <i>key-id</i> <b>md5</b> [ <i>key</i>   <i>hash-key</i>   <i>hash2-key</i> ] [ <b>hash</b>   <b>hash2</b> ]<br><b>no message-digest-key</b> <i>key-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf>area>virtual-link                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.</p> <p>The <b>no</b> form of the command removes the message digest key identified by the <i>key-id</i>.</p> <p>By default, no message keys are defined.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>key-id</i> — The <i>keyid</i> is expressed as a decimal integer.</p> <p><b>Values</b> 1 to 255</p> <p><i>key</i> — Specifies the MD5 key. The <i>key</i> can be any alphanumeric string up to 16 characters in length.</p> <p><i>hash-key</i>   <i>hash2-key</i> — The MD5 hash or hash2 key. the hash key. The key can be any combination of ASCII characters up to 32 (hash1-key) or 55 (hash2-key) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the <b>hash2</b> encrypted variable cannot be copied and pasted. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> |

## metric

|                    |                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>metric</b> <i>metric</i><br><b>no metric</b>                                                                                                               |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                       |
| <b>Description</b> | This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link. |

The **no** form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

|                   |                                                                                           |
|-------------------|-------------------------------------------------------------------------------------------|
| <b>Default</b>    | no metric                                                                                 |
| <b>Parameters</b> | <i>metric</i> — The metric to be applied to the interface expressed as a decimal integer. |
| <b>Values</b>     | 1 to 65535                                                                                |

## mtu

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mtu</b> <i>bytes</i><br><b>no mtu</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:</p> <ul style="list-style-type: none"> <li>• <b>config&gt;port&gt;ethernet</b></li> <li>• <b>config&gt;port&gt;sonet-sdh&gt;path</b></li> <li>• <b>config&gt;port&gt;tdm&gt;t3-e3</b></li> <li>• <b>config&gt;port&gt;tdm&gt;t1-e1&gt;channel-group</b></li> </ul> <p>If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.</p> <p>To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.</p> <p>The <b>no</b> form of this command reverts to the default derived from the MTU configured in the <b>config&gt;port</b> context.</p> |
| <b>Default</b>     | no mtu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>bytes</i> — The MTU to be used by OSPF for this logical interface in bytes.</p> <p><b>Values</b>      512 to 9198 in the <b>config&gt;router&gt;ospf&gt;area&gt;interface</b> context.<br/>1280 to 9198 in the <b>config&gt;router&gt;ospf3&gt;area&gt;interface</b> context.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## neighbor

|               |                                                                                      |
|---------------|--------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>[no] neighbor</b> <i>ipv4-address</i><br><b>[no] neighbor</b> <i>ipv6-address</i> |
|---------------|--------------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures an OSPF non-broadcast multi-access (NBMA) neighbor. The OSPF interface must be configured as an NBMA interface with the <b>interface-type non-broadcast</b> command. An NBMA network has no broadcast or multicast capabilities, so the router cannot discover its neighbors dynamically. All neighbors must be configured statically with the <b>neighbor</b> command.</p> <p>In addition to configuring the OSPF NBMA neighbor's IP address, the neighbor's MAC address may need to be configured with the <b>config&gt;router&gt;interface&gt;static-arp</b> command for OSPFv2 neighbors using its IPv4 address, and the <b>config&gt;router&gt;interface&gt;ipv6&gt;neighbor</b> command for OSPFv3 neighbors using its IPv6 link-local address.</p> <p>The <b>no</b> form of the command removes the <b>neighbor</b> configuration.</p> |
| <b>Default</b>     | no neighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>ipv4-address</i> — Specifies the OSPFv2 neighbor's IPv4 address.</p> <p><b>Values</b>      <i>ipv4-address</i> — a.b.c.d</p> <p><i>ipv6-address</i> — Specifies the OSPFv3 neighbor's IPv6 link-local address.</p> <p><b>Values</b></p> <p>        <i>ipv6-address</i>:    x:x:x:x:x:x:x [-interface]<br/>                               x:x:x:x:x:x:d.d.d.d [-interface]<br/>                               x: [0..FFFF]H<br/>                               d: [0..255]D<br/>                               interface — 32 characters maximum for link local addresses.</p>                                                                                                                                                                                                                                                                                      |

## node-sid

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>node-sid index</b> <i>index-value</i> [ <b>clear-n-flag</b> ]<br><b>node-sid label</b> <i>label-value</i> [ <b>clear-n-flag</b> ]<br><b>no node-sid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of type loopback. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.</p> <p>The above command should fail if the network interface is not of type loopback or if the interface is defined in an IES or a VPRN context. Assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.</p> |

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, the segment routing module checks that the same index or label value is not assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required because the index and therefore the label ranges, of IGP instances are not allowed to overlap.

The **clear-n-flag** option allows the user to clear the N-flag (node-sid flag) in an OSPF or OSPF3 prefix SID sub-TLV originated for the prefix of a loopback interface on the system. By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID, meaning that it belongs to this node only. However, when the user wants to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, you must clear the N-flag to assure interoperability with third-party implementations, which may perform a strict check on the receive end and drop duplicate prefix SID sub-TLVs when the N-flag is set.

The SR OS implementation is relaxed on the receive end and accepts duplicate prefix SIDs with the N-flag set or clear. SR OS will resolve to the closest owner, or owners if ECMP, of the prefix SID cost-wise.

|                   |                                                          |
|-------------------|----------------------------------------------------------|
| <b>Parameters</b> | <i>index-value</i> — Specifies the node SID index value. |
|                   | <b>Values</b> 0 to 4294967295                            |
|                   | <i>label-value</i> — Specifies the node SID label value. |
|                   | <b>Values</b> 0 to 4294967295                            |
|                   | <b>clear-n-flag</b> — Clears the node SID flag.          |
|                   | <b>Default</b> no clear-n-flag                           |

## passive

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] passive</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.</p> <p>By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The <b>passive</b> parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.</p> <p>While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.</p> <p>Service interfaces defined in <b>config&gt;router&gt;service-prefix</b> are passive. All other interfaces are not passive.</p> <p>The <b>no</b> form of the command removes the passive property from the OSPF interface.</p> |



**Default** no passive

## poll-interval

**Syntax** **poll-interval** *seconds*  
**no poll-interval**

**Context** config>router>ospf>area>interface  
config>router>ospf3>area>interface

**Description** This command configures the poll interval, in seconds. The poll interval is the time between two Hello packets to a dead (non-adjacent) OSPF NBMA neighbor. The default value of the poll interval timer is higher than the hello interval timer to avoid wasting bandwidth on non-broadcast networks, since OSPF messages are unicast to each configured neighbor. The poll interval timer is used only on **non-broadcast** interface types and has no effect if configured on other interface types.

The **no** form of the command removes the **poll-interval** configuration.

**Default** 120

**Parameters** *seconds* — Specifies the poll interval, in seconds.

**Values** 0 to 4294967295

## priority

**Syntax** **priority** *number*  
**no priority**

**Context** config>router>ospf>area>interface  
config>router>ospf3>area>interface

**Description** This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

The **no** form of the command reverts the interface priority to the default value.

**Default** priority 1

**Parameters** *number* — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router or Backup Designated Router on the interface subnet.

**Values** 0 to 255

---

## retransmit-interval

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retransmit-interval</b> <i>seconds</i><br><b>no retransmit-interval</b>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf>area>virtual-link<br>config>router>ospf3>area>interface<br>config>router>ospf3>area>virtual-link                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.</p> <p>The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgment has been received, the LSA will be retransmitted.</p> <p>The <b>no</b> form of this command reverts to the default interval.</p> |
| <b>Default</b>     | retransmit-interval 5                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>seconds</i> — The retransmit interval in seconds expressed as a decimal integer.<br><b>Values</b> 1 to 1800                                                                                                                                                                                                                                                                                                                                                                              |

## rib-priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rib-priority</b> { <b>high</b> }<br><b>no rib-priority</b>                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf3>area>interface                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface will be processed through the OSPF route calculation process at a higher priority.</p> <p>The <b>no</b> form of <b>rib-priority</b> command disables RIB prioritization at the associated level.</p> |
| <b>Default</b>     | no rib-priority                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <b>high</b> — Specifies that the name of the prefix list which contains prefixes get high priority for RIB-download. The high priority prefixes are downloaded first to the RIB. In doing so, the convergence time for these prefixes is better.                                                                                                                                 |

## sid-protection

|               |                            |
|---------------|----------------------------|
| <b>Syntax</b> | <b>[no] sid-protection</b> |
|---------------|----------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ospf>area>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command enables or disables adjacency SID protection by LFA and remote LFA.</p> <p>LFA and remote LFA Fast-Reroute (FRR) protection is enabled for all node SIDs and local adjacency SIDs when the user enables the <b>loopfree-alternate</b> option in IS-IS or OSPF at the LER and LSR. However, may be applications where the user never wants traffic to divert from the strict hop computed by CSPF for an SR-TE LSP. In this case, the user can disable protection for all adjacency SIDs formed over a particular network IP interface using this command.</p> <p>The protection state of an adjacency SID is advertised in the B-FLAG of the IS-IS or OSPF Adjacency SID sub-TLV.</p> |
| <b>Default</b>     | sid-protection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## transit-delay

|                    |                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>transit-delay</b> <i>seconds</i><br><b>no transit-delay</b>                                                                                                                                                                          |
| <b>Context</b>     | config>router>ospf>area>interface<br>config>router>ospf>area>virtual-link<br>config>router>ospf3>area>interface<br>config>router>ospf3>area>virtual-link                                                                                |
| <b>Description</b> | <p>This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link.</p> <p>The <b>no</b> form of this command reverts to the default delay time</p> |
| <b>Default</b>     | transit-delay 1                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>seconds</i> — The transit delay in seconds expressed as a decimal integer.<br><b>Values</b> 1 to 1800                                                                                                                                |



## 3.15 Show, Clear, and Debug Command Reference

### 3.15.1 Command Hierarchies

#### 3.15.1.1 Show Commands

```

show
 — ospf
 — ospf3
 — neighbor [detail]

show
 — router
 — ospf [ospf-instance | all]
 — area [area-id] [detail] [lfa]
 — capabilities [router-id]
 — database [type {router | network | summary | asbr-summary | external |
 nssa | all} [area area-id] [adv-router router-id] [link-state-id] [detail]
 [filtered]
 — interface [ip-int-name | ip-address | area area-id] [detail]
 — interface [ip-int-name | ip-address] database [detail]
 — lfa-coverage
 — neighbor [ip-int-name | ip-address] [detail]
 — neighbor overview
 — neighbor [remote ip-address] [detail]
 — opaque-database [area area-id | as] [adv-router router-id] [ls-id] [detail]
 — prefix-sids [ip-prefix[/prefix-length]] [sid sid] [adv-router router-id]
 — range [area-id]
 — routes [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary]
 [exclude-shortcut]
 — sham-link [interface-name] [detail]
 — sham-link interface-name remote ip-address [detail]
 — sham-link-neighbor [detail]
 — sham-link-neighbor interface-name remote ip-address [detail]
 — spf [lfa]
 — sr-adj-sets [adjacency-set-id] [detail]
 — sr-lfa-coverage [area area-id]
 — statistics
 — status
 — virtual-link [detail]
 — virtual-link database [detail]
 — virtual-neighbor [remote ip-address] [detail]
 — ospf3 [ospf-instance | all]
 — area [area-id] [detail] [lfa]
 — capabilities [router-id]
 — database [type database type] [area area-id] [adv-router router-id] [link-state-
 id] [detail] [filtered]

```

- **interface** [area area-id] [detail]
- **interface** [ip-int-name | ip-address | ipv6-address] [detail]
- **interface** ip-int-name | ip-address | ipv6-address **database** [detail]
- **lfa-coverage**
- **neighbor** [ip-int-name] [router-id] [detail]
- **neighbor overview**
- **prefix-sids** [ip-prefix[/prefix-length]] [sid sid] [adv-router router-id]
- **range** [area-id]
- **routes** [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary]
- **spf** [lfa]
- **sr-lfa-coverage** [area area-id]
- **statistics**
- **status**
- **virtual-link** [detail]
- **virtual-link** [detail] **database**
- **virtual-neighbor** [remote ip-address] [detail]

### 3.15.1.2 Clear Commands

- ```
clear
  — router
    — ospf [ospf-instance]
      — database [purge]
      — export
      — neighbor [ip-int-name | ip-address]
      — overload {rtm | fib | rtr-adv-lsa-limit}
      — statistics
    — ospf3 [ospf-instance]
      — database [purge]
      — export
      — neighbor [ip-int-name | router-id]
      — overload {rtm | fib | rtr-adv-lsa-limit}
      — statistics
```

3.15.1.3 Debug Commands

- ```
debug
 — router
 — ospf [ospf-instance]
 — area [area-id]
 — no area
 — area-range [ip-address]
 — no area-range
 — cspf [ip-address]
 — no cspf
 — [no] graceful-restart
 — interface [ip-int-name | ip-address]
 — no interface
```

---

```

— leak [ip-address]
— no leak
— lsdb [type] [ls-id] [adv-rtr-id] [area area-id]
— no lsdb
— [no] misc
— neighbor [ip-int-name | ip-address]
— no neighbor
— nssa-range [ip-address]
— no nssa-range
— packet [packet-type] [interface-name] [ingress | egress] [detail]
— no packet
— rsvp-shortcut [ip-address]
— no rsvp-shortcut
— rtm [ip-address]
— no rtm
— sham-neighbor [ip-address]
— no sham-neighbor
— spf [type] [dest-addr]
— no spf
— tunnel-endpoint [tunnel-spf ip-address] [tunnel-leak ip-address]
— no tunnel-endpoint
— virtual-neighbor [router-id]
— no virtual-neighbor
— ospf3 [ospf-instance]
— area [area-id]
— no area
— area-range [ipv4-address | ipv6-address]
— no area-range
— [no] graceful-restart
— interface [interface-name]
— no interface
— leak [ipv4-address | ipv6-address]
— no leak
— lsdb [type] [ls-id] [adv-rtr-id] [area area-id]
— no lsdb
— [no] misc
— neighbor [ip-int-name | router-id]
— no neighbor
— nssa-range [ipv4-address | ipv6-address]
— no nssa-range
— packet [packet-type] [interface-name] [ingress | egress | drop] [detail]
— no packet
— rtm [ipv4-address | ipv6-address]
— no rtm
— spf [type] [ipv4-address | ipv6-address]
— no spf
— tunnel-endpoint [tunnel-spf ip-address] [tunnel-leak ip-address]
— no tunnel-endpoint
— virtual-neighbor [router-id]
— no virtual-neighbor

```

### 3.15.1.4 Tools Commands

```
tools
 — dump
 — router
 — ospf
 — sr-adjacencies [ip-int-name | ip-address] [detail]
 — sr-adjacencies [remote ip-address] [detail]
 — sr-database [sid sid] [prefix ip-address] [nh-type {sr | ldp | rsvp}]
 [detail]
 — ospf3
 — sr-adjacencies [ip-int-name | ip-address] [[detail]
 — sr-adjacencies [remote ip-address] [detail]
 — sr-database [sid sid] [prefix ip-address] [nh-type {sr | ldp | rsvp}]
 [detail]
```

## 3.15.2 Command Descriptions

### 3.15.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

#### neighbor

|             |                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------|
| Syntax      | neighbor [detail]                                                                                       |
| Context     | show>ospf>neighbor<br>show>ospf3>neighbor                                                               |
| Description | This command shows all OSPFv2 or OSPFv3 neighbor information.                                           |
| Parameters  | <b>detail</b> — Displays the detailed OSPF and OSPF3 neighbor status for all routers and all instances. |
| Output      | The following output is an example of OSPF and OSPF3 neighbor information.                              |

#### Sample Output

```
*A:Dut-C# show ospf neighbor

=====
Rtr Base OSPFv2 Instance 0 Neighbors
=====
Interface-Name Rtr Id State Pri RetxQ TTL
```



```

Area-Id

DUTC_TO_DUTB.1.0 10.20.1.2 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTE.1.0 10.20.1.5 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTB.2.0 10.20.1.2 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTE.2.0 10.20.1.5 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTA.13.0 10.20.1.1 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTA.14.0 10.20.1.1 Full 1 0 9
 0.0.0.0

No. of Neighbors: 6
=====

Rtr vprn1 OSPFv2 Instance 0 Neighbors
=====
Interface-Name Rtr Id State Pri RetxQ TTL
Area-Id

DUTC_TO_DUTA.7.0 10.20.4.1 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTD.7.0 10.20.4.4 Full 1 0 12
 0.0.0.0
DUTC_TO_DUTA.8.0 10.20.4.1 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTD.8.0 10.20.4.4 Full 1 0 12
 0.0.0.0

No. of Neighbors: 4
=====
*A:Dut-C# show ospf neighbor detail

=====
Rtr Base OSPFv2 Instance 0 Neighbors (detail)
=====
Neighbor Rtr Id : 10.20.1.2 Interface: DUTC_TO_DUTB.1.0

Neighbor IP Addr : 1.0.23.2
Local IF IP Addr : 1.0.23.3
Area Id : 0.0.0.0
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : - E - - - - O --
Events : 4 Last Event Time : 09/20/2018 14:38:42
Up Time : 0d 00:02:34 Time Before Dead : 9 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown (0)
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches : 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never
<snip>

```

```

=====
*A:Dut-C#
*A:Dut-C#
*A:Dut-C# show ospf3 neighbor

=====
Rtr Base OSPFv3 Instance 0 Neighbors
=====
Interface-Name Rtr Id State Pri RetxQ TTL
 Area-Id

DUTC_TO_DUTB.3.0 10.20.2.2 Full 1 0 8
 0.0.0.0
DUTC_TO_DUTE.3.0 10.20.2.5 Full 1 0 8
 0.0.0.0
DUTC_TO_DUTB.4.0 10.20.2.2 Full 1 0 8
 0.0.0.0
DUTC_TO_DUTE.4.0 10.20.2.5 Full 1 0 8
 0.0.0.0
DUTC_TO_DUTA.15.0 10.20.2.1 Full 1 0 8
 0.0.0.0
DUTC_TO_DUTA.16.0 10.20.2.1 Full 1 0 9
 0.0.0.0

No. of Neighbors: 6
=====

=====
Rtr Base OSPFv3 Instance 64 Neighbors
=====
Interface-Name Rtr Id State Pri RetxQ TTL
 Area-Id

DUTC_TO_DUTB.5.0 10.20.3.2 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTE.5.0 10.20.3.5 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTB.6.0 10.20.3.2 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTE.6.0 10.20.3.5 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTA.17.0 10.20.3.1 Full 1 0 9
 0.0.0.0
DUTC_TO_DUTA.18.0 10.20.3.1 Full 1 0 9
 0.0.0.0

No. of Neighbors: 6
=====
<snip>
=====

*A:Dut-C# show ospf3 neighbor detail

=====
Rtr Base OSPFv3 Instance 0 Neighbors (detail)
=====
Neighbor Rtr Id : 10.20.2.2 Interface: DUTC_TO_DUTB.3.0

```

```

Neighbor IP Addr : fe80::200:ff:fe00:2-"DUTC_TO_DUTB.3.0"
Local IF IP Addr : fe80::200:ff:fe00:3-"DUTC_TO_DUTB.3.0"
Area Id : 0.0.0.0
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : ----R--EV6
Events : 4 Last Event Time : 09/20/2018 14:38:43
Up Time : 0d 00:03:01 Time Before Dead : 8 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason : Unknown (0)
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.20.2.5 Interface: DUTC_TO_DUTE.3.0

Neighbor IP Addr : fe80::200:ff:fe00:5-"DUTC_TO_DUTE.3.0"
Local IF IP Addr : fe80::200:ff:fe00:3-"DUTC_TO_DUTE.3.0"
Area Id : 0.0.0.0
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : ----R--EV6
Events : 6 Last Event Time : 09/20/2018 14:38:43
Up Time : 0d 00:03:02 Time Before Dead : 9 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason : Unknown (0)
Bad Nbr States : 0 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.20.2.2 Interface: DUTC_TO_DUTB.4.0

Neighbor IP Addr : fe80::200:ff:fe00:2-"DUTC_TO_DUTB.4.0"
Local IF IP Addr : fe80::200:ff:fe00:3-"DUTC_TO_DUTB.4.0"
Area Id : 0.0.0.0
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : ----R--EV6
Events : 4 Last Event Time : 09/20/2018 14:38:43
Up Time : 0d 00:03:02 Time Before Dead : 9 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason : Unknown (0)
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

<snip>
=====
*A:Dut-C#

```

---

## ospf

|                    |                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ospf</b> [ <i>ospf-instance</i>   <b>all</b> ]                                                                                       |
| <b>Context</b>     | show>router                                                                                                                             |
| <b>Description</b> | This command enables the context to display OSPF information.                                                                           |
| <b>Parameters</b>  | <i>ospf-instance</i> — Shows the specified OSPF instance.<br><b>Values</b> 0 to 31<br><b>all</b> — Shows all configured OSPF instances. |

## ospf3

|                    |                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ospf3</b> [ <i>ospf-instance</i>   <b>all</b> ]                                                                                                                                                                                     |
| <b>Context</b>     | show>router                                                                                                                                                                                                                            |
| <b>Description</b> | This command enables the context to display OSPF3 information.                                                                                                                                                                         |
| <b>Parameters</b>  | <i>ospf-instance</i> — Shows the specified OSPF3 instance.<br><b>Values</b> 0 to 31   64 to 95<br>0 to 31 — IPv6-unicast address-family<br>64 to 95— IPv4-unicast address-family<br><b>all</b> — Shows all configured OSPF3 instances. |

## area

|                    |                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>area</b> [ <i>area-id</i> ] [ <b>detail</b> ] [ <b>lfa</b> ]                                                                                                                                                                                                                                                            |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays configuration information about all areas or the specified area. When detail is specified operational and statistical information will be displayed.                                                                                                                                                 |
| <b>Parameters</b>  | <i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.<br><b>Values</b> ip-address — a.b.c.d<br>area — 0 to 4294967295<br><b>detail</b> — Displays detailed information about the specified area.<br><b>lfa</b> — Displays Loop-Free Alternate (LFA) next-hop information. |
| <b>Output</b>      | OSPF Area Output                                                                                                                                                                                                                                                                                                           |

The following table describes the standard and detailed command output fields for an OSPF area.

**Table 16**      **OSPF Area Output Fields**

| Label             | Description                                                                                                                                                                    |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Area Id           | A 32 bit integer uniquely identifying an area.                                                                                                                                 |
| Type              | NSSA — This area is configured as an NSSA area.<br>Standard — This area is configured as a standard area (not NSSA or Stub).<br>Stub — This area is configured as a stub area. |
| SPF Runs          | The number of times that the intra-area route table has been calculated using this area's link state database.                                                                 |
| LSA Count         | The total number of link-state advertisements in this area's link state database, excluding AS External LSAs.                                                                  |
| LSA Cksum Sum     | The 32-bit unsigned sum of the link-state database advertisements LS checksums contained in this area's link state database. This checksum excludes AS External LSAs (type-5). |
| No. of OSPF Areas | The number of areas configured on the router.                                                                                                                                  |
| Virtual Links     | The number of virtual links configured through this transit area.                                                                                                              |
| Active IFs        | The active number of interfaces configured in this area.                                                                                                                       |
| Area Bdr Rtrs     | The total number of ABRs reachable within this area.                                                                                                                           |
| AS Bdr Rtrs       | The total number of ASBRs reachable within this area.                                                                                                                          |
| Last SPF Run      | The time when the last intra-area SPF was run on this area.                                                                                                                    |
| Router LSAs       | The total number of router LSAs in this area.                                                                                                                                  |
| Network LSAs      | The total number of network LSAs in this area.                                                                                                                                 |
| Summary LSAs      | The summary of LSAs in this area.                                                                                                                                              |
| Asbr-summ LSAs    | The summary of ASBR LSAs in this area.                                                                                                                                         |
| Nssa-ext LSAs     | The total number of NSSA-EXT LSAs in this area.                                                                                                                                |
| Area opaque LSAs  | The total number of opaque LSAs in this area.                                                                                                                                  |
| Total Nbrs        | The total number of neighbors in this area.                                                                                                                                    |
| Total IFs         | The total number of interfaces configured in this area.                                                                                                                        |
| Total LSAs        | The sum of LSAs in this area excluding autonomous system external LSAs.                                                                                                        |

**Table 16 OSPF Area Output Fields (Continued)**

| Label             | Description (Continued)                                                                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Blackhole Range   | False — No blackhole route is installed for aggregates configured in this area.<br>True — A lowest priority blackhole route is installed for aggregates configured in this area. |
| Export Database   | False — The extended TE Database (TE-DB) is not exported in this area.<br>True — The extended TE Database (TE-DB) is exported in this area.                                      |
| Export Policies   | The export policies configured in this area.                                                                                                                                     |
| Export Fltrd LSAs | The total number of LSAs filtered by area export policies.                                                                                                                       |
| Import Policies   | The import policies configured in this area.                                                                                                                                     |
| Import Fltrd LSAs | The total number of LSAs filtered by area import policies.                                                                                                                       |

**Sample Output**

```

A:SetupCLI# show router ospf 0 area detail
=====
Rtr Base OSPFv2 Instance 0 Areas (detail)
=====

Area Id: 0.0.0.0

Area Id : 0.0.0.0 Type : Standard
Key Rollover Int.: 10 LFA : Include
Virtual Links : 0 Total Nbrs : 0
Active IFs : 0 Total IFs : 2
Area Bdr Rtrs : 0 AS Bdr Rtrs : 0
SPF Runs : 0 Last SPF Run : Never
Router LSAs : 0 Network LSAs : 0
Summary LSAs : 0 Asbr-summ LSAs : 0
Nssa ext LSAs : 0 Area opaque LSAs : 1
Total LSAs : 1 LSA Cksum Sum : 0xd6af
Blackhole Range : False Unknown LSAs : 0
Export Database : False
Export Policies : None
Export Fltrd LSAs: 0
Import Policies : None
Import Fltrd LSAs: 0

Area Id: 1.1.1.1

Area Id : 1.1.1.1 Type : Stub
Default Cost : 16777215 Import Summary : Send Summary
Key Rollover Int.: 10 LFA : Exclude
Virtual Links : 0 Total Nbrs : 0
Active IFs : 0 Total IFs : 1
Area Bdr Rtrs : 0 AS Bdr Rtrs : 0

```

```

SPF Runs : 0
Router LSAs : 0
Summary LSAs : 0
Nssa ext LSAs : 0
Total LSAs : 1
Blackhole Range : False
Export database : False
Export Policies : None
Export Fltrd LSAs: 0
Import Policies : None
Import Fltrd LSAs: 0
Last SPF Run : Never
Network LSAs : 0
Asbr-summ LSAs : 0
Area opaque LSAs : 1
LSA Cksum Sum : 0xf493
Unknown LSAs : 0

```

-----  
Area Id: 2.2.2.2  
-----

```

Area Id : 2.2.2.2
Key Rollover Int.: 10
Virtual Links : 1
Active IFs : 0
Area Bdr Rtrs : 0
SPF Runs : 0
Router LSAs : 0
Summary LSAs : 0
Nssa ext LSAs : 0
Total LSAs : 1
Blackhole Range : True
Export database : False
Export Policies : None
Export Fltrd LSAs: 0
Import Policies : None
Import Fltrd LSAs: 0
Type : Standard
LFA : Include
Total Nbrs : 0
Total IFs : 0
AS Bdr Rtrs : 0
Last SPF Run : Never
Network LSAs : 0
Asbr-summ LSAs : 0
Area opaque LSAs : 1
LSA Cksum Sum : 0xd6af
Unknown LSAs : 0

```

=====

A:SetupCLI#

\*A:Dut-B# show router ospf area 0.0.0.0 detail

=====

Rtr Base OSPFv2 Instance 0 Area 0.0.0.0 (detail)

-----  
Configuration  
-----

```

Area Id : 0.0.0.0
Type : Standard

```

-----  
Statistics  
-----

```

Virtual Links : 0
Active IFs : 3
Area Bdr Rtrs : 0
SPF Runs : 7
Router LSAs : 3
Summary LSAs : 0
Nssa ext LSAs : 0
Total LSAs : 9
Blackhole Range : True
Export database : False
Export Policies : None
Export Fltrd LSAs: 0
Import Policies : None
Import Fltrd LSAs: 0
Total Nbrs : 2
Total IFs : 3
AS Bdr Rtrs : 0
Last SPF Run : 10/26/2006 10:09:18
Network LSAs : 3
Asbr-summ LSAs : 0
Area opaque LSAs : 3
LSA Cksum Sum : 0x28b62
Unknown LSAs : 0

```

=====

```
*A:Dut-B# show router ospf 0 area 0.0.0.0 lfa
=====
Rtr Base OSPFv2 Instance 0 Path Table
=====
Node Interface Nexthop
 LFA Interface LFA Nexthop

10.20.1.1 to_Dut-A1 10.20.1.1
 to_Dut-C1 10.20.1.3
10.20.1.3 to_Dut-C1 10.20.1.3
 to_Dut-A1 10.20.1.1
10.20.1.4 to_Dut-D1 10.20.1.4
10.20.1.6 to_Dut-D1 10.20.1.4
 to_Dut-C1 10.20.1.3
=====
*A:Dut-B#

*A:Dut-B# show router ospf 0 area 0.0.0.0 lfa detail
=====
Rtr Base OSPFv2 Instance 0 Path Table (detail)
=====
OSPF Area : 0.0.0.0

Node : 10.20.1.1 Metric : 10
Interface : to_Dut-A1 Nexthop : 10.20.1.1
LFA Interface : to_Dut-C1 LFA Metric : 20
LFA type : linkProtection LFA Nexthop : 10.20.1.3

Node : 10.20.1.3 Metric : 10
Interface : to_Dut-C1 Nexthop : 10.20.1.3
LFA Interface : to_Dut-A1 LFA Metric : 20
LFA type : linkProtection LFA Nexthop : 10.20.1.1

Node : 10.20.1.4 Metric : 10
Interface : to_Dut-D1 Nexthop : 10.20.1.4

Node : 10.20.1.6 Metric : 20
Interface : to_Dut-D1 Nexthop : 10.20.1.4
LFA Interface : to_Dut-C1 LFA Metric : 30
LFA type : nodeProtection LFA Nexthop : 10.20.1.3
=====
*A:Dut-B#
```

capabilities

|             |                                                                        |
|-------------|------------------------------------------------------------------------|
| Syntax      | capabilities [router-id]                                               |
| Context     | show>router>ospf<br>show>router>ospf3                                  |
| Description | This command displays the entries in the Router Information (RI) LSAs. |



**Parameters** *router-id* — Lists only the LSAs related to that router ID. If no *router-id* is specified, all database entries are listed.

**Output** OSPF Capabilities Output

The following table describes the standard and detailed command output fields for OSPF capabilities.

### Sample Output

```
*A:Dut-C# show router ospf capabilities
=====
Rtr Base OSPFv2 Instance 0 Capabilities
=====
scope Router Id Capabilities

Area 10.20.1.2 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 22000 range 1001
Area 10.20.1.3 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 23030 range 71
Area 10.20.1.4 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 24000 range 1001
Area 10.20.1.5 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 25000 range 1001
Area 10.20.1.1 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 21000 range 1001
Area 10.20.1.2 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 22000 range 1001
Area 10.20.1.3 0x14: Stub P2P-VLAN
 SR Algorithm: IGP-metric-based-SPF
 SR Label Range: start label 23030 range 71

No. of LSAs: 7
=====
```

## database

**Syntax** **database** [**type** {**router** | **network** | **summary** | **asbr-summary** | **external** | **nssa** | **all**}] [**area** *area-id*] [**adv-router** *router-id*] [*link-state-id*] [**detail**] [**filtered**]  
**database** [**type** {**router** | **network** | **inter-area-pfx** | **inter-area-rtr** | **external** | **nssa** | **intra-area-pfx** | **rtr-info** | **all**}] [**area** *area-id*] [**adv-router** *router-id*] [*link-state-id*] [**detail**] [**filtered**]

**Context** show>router>ospf  
show>router>ospf3

**Description** This command displays information about the OSPF link state database (LSDB).

When no command line options are specified, the command displays brief output for all database entries

- Parameters**
- type** *keyword* — Specifies to filter the OSPF LSDB information based on the type specified by *keyword*.
  - router** — Displays only router (Type 1) LSAs in the LSDB.
  - network** — Displays only network (Type 2) LSAs in the LSDB.
  - summary** — Displays only summary (Type 3) LSAs in the LSDB.
  - asbr-summary** — Displays only ASBR summary (Type 4) LSAs in the LSDB.
  - external** — Displays only AS external (Type 5) LSAs in the LSDB. External LSAs are maintained globally and not per area. If the display of external links is requested, the area parameter, if present, is ignored.
  - nssa** — Displays only NSSA area-specific AS external (Type 7) LSAs in the LSDB.
  - inter-area-pfx** — Displays inter-area prefix LSAs.
  - inter-area-rtr** — Displays inter-area router LSAs.
  - intra-area-pfx** — Displays intra-area prefix LSAs.
  - rtr-info** — Displays router info LSAs.
  - all** — Displays all LSAs in the LSDB. The all keyword is intended to be used with either the **area** *area-id* or the **adv-router** *router-id* [*link-state-id*] parameters.
  - area** *area-id* — Displays LSDB information associated with the specified OSPF *area-id*.
    - Values**
      - ip-address — a.b.c.d
      - area — 0 to 4294967295
  - adv-router** *router-id* [*link-state-id*] — Displays LSDB information associated with the specified advertising router. To further narrow the number of items displayed, the *link-state-id* can optionally be specified.
  - detail** — Displays detailed information on the LSDB entries.
  - filtered** — Displays LSDB entries that were filtered by an area import or export policy.

**Output** OSPF Database Output

The following table describes the standard and detailed command output fields for an OSPF database.

**Table 17** OSPF Database Output Fields

| Label   | Description               |
|---------|---------------------------|
| Area Id | The OSPF area identifier. |

**Table 17 OSPF Database Output Fields (Continued)**

| Label                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type<br>LSA Type            | Router — LSA type of router (OSPF)<br>Network — LSA type of network (OSPF)<br>Summary — LSA type of summary (OSPF)<br>ASBR Summary — LSA type of ASBR summary (OSPF)<br>Nssa-ext — LSA area-specific, NSSA external (OSPF)<br>Area opaque — LSA type of area opaque (OSPF)<br>router — LSA type of router (OSPF3)<br>Network — LSA type of network (OSPF3)<br>IE Pfx — LSA type of inter-area prefix (OSPF3) Newline IE Rtr<br>LSA type of inter-area router (OSPF3)<br>IA Pfx — LSA type of intra-area prefix (OSPF3)                                                   |
| Link State Id               | The link state Id is an LSA type specific field containing either a number to distinguish several LSAs from the same router, an interface ID, or a router-id; it identifies the piece of the routing domain being described by the advertisement.                                                                                                                                                                                                                                                                                                                        |
| Adv Rtr Id<br>Adv Router Id | The router identifier of the router advertising the LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Age                         | The age of the link state advertisement in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Sequence<br>Sequence No     | The signed 32-bit integer sequence number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Cksum<br>Checksum           | The 32-bit unsigned sum of the link-state advertisements' LS checksums.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| No. of LSAs                 | The number of LSAs displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Options                     | EA — External Attribute LSA Support<br>DC — Demand Circuit Support<br>R — If clear, a node can participates in OSPF topology distribution without being used to forward transit traffic.<br>N — Type 7 LSA Support<br>MC — Multicast Support<br>E — External Routes Support<br>V6 — V6 works in conjunction with R. If V6 is clear, a node can participate in OSPF topology distribution without being used to forward IPv6 datagrams. If R is set and V6 is clear, IPv6 datagrams are not forwarded but diagrams belonging to another protocol family may be forwarded. |

**Table 17** OSPF Database Output Fields (Continued)

| Label                    | Description                                                                                                                                                                                                                                                      |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prefix Options           | P — Propagate NSSA LSA.<br>MC — Multicast support.<br>LA — Local address capability. If set, the prefix is an IPv6 interface address of the advertising router.<br>NU — No unicast capability. If set, the prefix is excluded from IPv6 unicast calculations.    |
| Flags                    | None — No flags set<br>V — The router is an endpoint for one or more fully adjacent Virtual Links having the described area as the transit area<br>E — The router is an AS Boundary Router<br>B — The router is an Area Border Router                            |
| Link Count               | The number of links advertised in the LSA.                                                                                                                                                                                                                       |
| Link Type ( <i>n</i> )   | The link type of the <i>n</i> th link in the LSA.                                                                                                                                                                                                                |
| Network ( <i>n</i> )     | The network address of the <i>n</i> th link in the LSA.                                                                                                                                                                                                          |
| Mask ( <i>n</i> )        | The mask of the <i>n</i> th link in the LSA.                                                                                                                                                                                                                     |
| No of TOS ( <i>n</i> )   | The TOS of the <i>n</i> th link in the LSA.                                                                                                                                                                                                                      |
| Metric-0 ( <i>n</i> )    | The cost metric of the <i>n</i> th link in the LSA.                                                                                                                                                                                                              |
| Not Flooded ( <i>n</i> ) | Filtered on Import<br>The <i>n</i> th LSA was filtered by an area import policy.<br>Filtered on Export<br>The <i>n</i> th LSA was filtered by an area export policy.<br><b>Note:</b> this line does not appear if there is no filtering and the LSA was flooded. |

**Sample Output**

```

A:ALA-A# show router ospf 1 database
=====
Rtr Base OSPFv2 Instance 1 Link State Database (type : All)
=====
Area Id Type Link State Id Adv Rtr Id Age Sequence Cksum

0.0.0.0 Router 180.0.0.2 180.0.0.2 1800 0x800000b6 0xf54
0.0.0.0 Router 180.0.0.5 180.0.0.5 1902 0x8000009d 0xcb7c
0.0.0.0 Router 180.0.0.8 180.0.0.8 1815 0x8000009a 0x529b
0.0.0.0 Router 180.0.0.9 180.0.0.9 1156 0x80000085 0xd00f
0.0.0.0 Router 180.0.0.10 180.0.0.10 533 0x8000009d 0x3f1f
0.0.0.0 Router 180.0.0.11 180.0.0.11 137 0x80000086 0xc58f
0.0.0.0 Router 180.0.0.12 180.0.0.12 918 0x8000009d 0x4cf3
0.0.0.0 Router 180.0.0.13 180.0.0.13 1401 0x800000a2 0x879c

```

```

0.0.0.0 Network 180.0.53.28 180.0.0.28 149 0x80000083 0xe5cd
0.0.0.0 Network 180.0.54.28 180.0.0.28 1259 0x80000083 0xdad7
0.0.0.0 Summary 180.0.0.15 180.0.0.10 378 0x80000084 0xeba1
0.0.0.0 Summary 180.0.0.15 180.0.0.12 73 0x80000084 0xdfab
0.0.0.0 Summary 180.0.0.18 180.0.0.10 1177 0x80000083 0xcfbb
0.0.0.1 Summary 180.100.25.4 180.0.0.12 208 0x80000091 0x3049
0.0.0.1 AS Summ 180.0.0.8 180.0.0.10 824 0x80000084 0x3d07
0.0.0.1 AS Summ 180.0.0.8 180.0.0.12 1183 0x80000095 0x4bdf
0.0.0.1 AS Summ 180.0.0.9 180.0.0.10 244 0x80000082 0x73cb
n/a AS Ext 7.1.0.0 180.0.0.23 1312 0x80000083 0x45e7
n/a AS Ext 7.2.0.0 180.0.0.23 997 0x80000082 0x45e6
n/a AS Ext 10.20.0.0 180.0.0.23 238 0x80000081 0x2d81
...

```

No. of LSAs: 339

A:ALA-A# show router ospf 1 database detail

Rtr Base OSPFv2 Instance 1 Link State Database (type : All) (detail)

Router LSA for Area 0.0.0.0

```

Area Id : 0.0.0.0 Adv Router Id : 10.10.10.10
Link State Id : 10.10.10.10 (168430090)
LSA Type : Router
Sequence No : 0x80000002 Checksum : 0xd81
Age : 594 Length : 36
Options : E
Flags : None Link Count : 1
Link Type (1) : Stub Network
Network (1) : 192.168.0.0 Mask (1) : 255.255.255.0
No of TOS (1) : 0 Metric-0 (1) : 10

```

AS Ext LSA for Network 180.0.0.14

```

Area Id : N/A Adv Router Id : 180.0.0.10
Link State Id : 180.0.0.14 LSA Type : AS Ext
Sequence No : 0x80000083 Checksum : 0xa659
Age : 2033 Length : 36
Options : E
Network Mask : 255.255.255.255 Fwding Address : 180.1.6.15
Metric Type : Type 2 Metric-0 : 4
Ext Route Tag : 0

```

A:ALA-A# show router ospf database filtered

Rtr Base OSPFv2 Instance 0 Link State Database (filtered)

```

Type Area Id Link State Id Adv Rtr Id Filter
Router 0.0.0.0 10.0.0.2 10.0.0.2
Router 0.0.0.0 10.0.0.2 10.0.0.2
Summary 0.0.0.0 1.2.3.0 10.0.0.2
Summary 0.0.0.0 10.0.0.1 10.0.0.2 Export 0.0.0.1
AS Summ 0.0.0.0 10.0.0.1 10.0.0.2
Router 0.0.0.1 10.0.0.1 10.0.0.1
Router 0.0.0.1 10.0.0.2 10.0.0.2

```

```

Network 0.0.0.1 1.1.2.2 10.0.0.2
Summary 0.0.0.1 1.2.3.0 10.0.0.2 Import 0.0.0.1
Summary 0.0.0.1 1.2.4.0 10.0.0.2 Import 0.0.0.1
Summary 0.0.0.1 1.4.6.0 10.0.0.2 Import 0.0.0.1
Summary 0.0.0.1 1.5.6.0 10.0.0.2 Import 0.0.0.1
Summary 0.0.0.2 1.5.6.0 10.0.0.2
Summary 0.0.0.2 10.0.0.1 10.0.0.2 Export 0.0.0.1
Summary 0.0.0.2 10.0.6.0 10.0.0.2

```

```

No. of LSAs: 14
No. of Export Filtered LSAs: 2
No. of Import Filtered LSAs: 4
=====

```

A:ALA-A#

```

*A:Dut-C>config>router>ospf3# show router ospf3 0 database detail adv-
router 10.20.1.3

```

```

=====
Rtr Base OSPFv3 Instance 0 Link State Database (type: All) (detail)
=====

```

```

Router LSA for Area 0.0.0.1

```

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.3
Link State Id : 0.0.0.0 (0)
LSA Type : Router
Sequence No : 0x80000002 Checksum : 0x8749
Age : 343 Length : 56
Options : ---R--EV6
Flags : ASBR
Link Type (1) : P2P Link Link Count : 2
I/F Index (1) : 2 Nbr Rtr ID (1) : 10.20.1.1
Metric (1) : 100 Nbr I/F Index (1): 3
Link Type (2) : P2P Link Nbr Rtr ID (2) : 10.20.1.6
I/F Index (2) : 3 Nbr I/F Index (2): 4
Metric (2) : 100

```

```

IA Pfx LSA for Area 0.0.0.1

```

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.3
Link State Id : 0.0.0.0 (0)
LSA Type : IA Pfx
Sequence No : 0x80000004 Checksum : 0x285a
Age : 303 Length : 192
Ref Ls Type : 2001 Ref Ls Id : 0
Ref Adv Rtr : 10.20.1.3 No of Pfxs : 8
Prefix (1) : 3ffe::a14:103/128
Pfx Options (1) : LA Metric (1) : 0
Prefix (2) : 3ffe::3300:d00/120
Pfx Options (2) : Metric (2) : 100
Prefix (3) : 3ffe::100:d00/120
Pfx Options (3) : Metric (3) : 100
Prefix (4) : 3ffe::3300:2400/120
Pfx Options (4) : Metric (4) : 100
Prefix (5) : 3ffe::100:2400/120
Pfx Options (5) : Metric (5) : 100
Prefix (6) : 3ffe::100:900/120
Pfx Options (6) : Metric (6) : 200
Prefix (7) : 3ffe::3300:2103/128

```

```

Pfx Options (7) : LA Metric (7) : 0
Prefix (8) : 3ffe::100:2103/128
Pfx Options (8) : LA Metric (8) : 0

Rtr Inf LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.3
Link State Id : 0.0.0.0 (0)
LSA Type : Rtr Inf
Sequence No : 0x80000002 Checksum : 0x783a
Age : 343 Length : 68
TLV dump:
 Capabilities (1) Len 4 :
 0x28000000
 SR algorithm (8) Len 2 :
 0x2 0x0
 SR label range (9) Len 12 :
 Range-size=1000
 Sub-TLV SID/label(7) len 3 :
 label=40000
 SR Readable label depth capability: 15
 non-OSPF Func Capabilities (32768) Len 4 :
 0x80000000

E-Router LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.3
Link State Id : 0.0.0.0 (0)
LSA Type : E-Router
Sequence No : 0x80000004 Checksum : 0x44a8
Age : 303 Length : 88
Options : ----R--EV6
Flags : ASBR
Link Type (1) : P2P Link Nbr Rtr ID (1) : 10.20.1.1
I/F Index (1) : 2 Nbr I/F Index (1): 3
Metric (1) : 100
Adj SID (1) : 524287 Flags (1) : BVL--
 Weight (1) : 0
Link Type (2) : P2P Link Nbr Rtr ID (2) : 10.20.1.6
I/F Index (2) : 3 Nbr I/F Index (2): 4
Metric (2) : 100
Adj SID (2) : 524286 Flags (2) : BVL--
 Weight (2) : 0

E-IA Pfx LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.3
Link State Id : 0.0.0.0 (0)
LSA Type : E-IA Pfx
Sequence No : 0x80000004 Checksum : 0xb319
Age : 304 Length : 140
Ref Ls Type : a021 Ref Ls Id : 0
Ref Adv Rtr : 10.20.1.3 No of Pfxs : 3
Prefix (1) : 3ffe::a14:103/128
Pfx Options (1) : N LA Metric (1) : 0
Prefix SID (1) : 6 Flags (1) : NP-----
 Algorithm (1) : 0
Prefix (2) : 3ffe::3300:2103/128
Pfx Options (2) : N LA Metric (2) : 0

```

```

Prefix (3) : 3ffe::100:2103/128
Pfx Options (3) : N LA
Prefix SID (3) : 0
Metric (3) : 0
Flags (3) : NP-----
Algorithm (3) : 0

```

```

AS Ext LSA (0.0.0.2 (2))

```

```

Area Id : N/A
Link State Id : 0.0.0.2 (2)
LSA Type : AS Ext
Sequence No : 0x80000001
Age : 348
Options : E
Pfx Options :
Dest Prefix : 3ffe::6700:1/128
Checksum : 0x4775
Length : 44
E2-Metric : 1
Adv Router Id : 10.20.1.3

```

```

*A:Dut-C>config>router>ospf3# show router ospf3 0 database detail adv-
router 10.20.1.6

```

```
=====
Rtr Base OSPFv3 Instance 0 Link State Database (type: All) (detail)
=====
```

```

Router LSA for Area 0.0.0.1

```

```

Area Id : 0.0.0.1
Link State Id : 0.0.0.0 (0)
LSA Type : Router
Sequence No : 0x80000002
Age : 328
Options : ----R--EV6
Flags : ABR ASBR
Link Type (1) : P2P Link
I/F Index (1) : 2
Metric (1) : 100
Link Type (2) : P2P Link
I/F Index (2) : 4
Metric (2) : 100
Checksum : 0x2aa3
Length : 56
Link Count : 2
Nbr Rtr ID (1) : 10.20.1.2
Nbr I/F Index (1) : 4
Nbr Rtr ID (2) : 10.20.1.3
Nbr I/F Index (2) : 3
Adv Router Id : 10.20.1.6

```

```

IE Pfx LSA for Area 0.0.0.1

```

```

Area Id : 0.0.0.1
Link State Id : 0.0.0.1 (1)
LSA Type : IE Pfx
Sequence No : 0x80000001
Age : 335
Options :
Dest Prefix : 3ffe::a14:106/128
Checksum : 0x45df
Length : 44
Metric : 0
Adv Router Id : 10.20.1.6

```

```

IE Pfx LSA for Area 0.0.0.1

```

```

Area Id : 0.0.0.1
Link State Id : 0.0.0.3 (3)
LSA Type : IE Pfx
Sequence No : 0x80000001
Age : 328
Options :
Dest Prefix : 3ffe::100:3705/128
Checksum : 0xb9b
Length : 44
Metric : 100
Adv Router Id : 10.20.1.6

```



IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.4 (4)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x83ef
Age : 328 Length : 44
Pfx Options : Metric : 100
Dest Prefix : 3ffe::3300:3705/128

```

IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.5 (5)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x5d60
Age : 328 Length : 44
Pfx Options : Metric : 100
Dest Prefix : 3ffe::a14:105/128

```

IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.6 (6)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x1932
Age : 328 Length : 44
Pfx Options : Metric : 200
Dest Prefix : 3ffe::100:2c04/128

```

IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.8 (8)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x878f
Age : 328 Length : 44
Pfx Options : Metric : 200
Dest Prefix : 3ffe::3300:2c04/128

```

IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.9 (9)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x75e0
Age : 328 Length : 44
Pfx Options : Metric : 200
Dest Prefix : 3ffe::a14:104/128

```

IE Pfx LSA for Area 0.0.0.1

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.10 (10)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0xdcde
Age : 328 Length : 44
Pfx Options : Metric : 100
Dest Prefix : 3ffe::a14:102/128

```

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                       |               |             |
|---------------|-----------------------|---------------|-------------|
| Area Id       | : 0.0.0.1             | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.12 (12)       |               |             |
| LSA Type      | : IE Pfx              |               |             |
| Sequence No   | : 0x80000001          | Checksum      | : 0xc86a    |
| Age           | : 328                 | Length        | : 44        |
| Pfx Options   | :                     | Metric        | : 200       |
| Dest Prefix   | : 3ffe::3300:1800/120 |               |             |

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                      |               |             |
|---------------|----------------------|---------------|-------------|
| Area Id       | : 0.0.0.1            | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.13 (13)      |               |             |
| LSA Type      | : IE Pfx             |               |             |
| Sequence No   | : 0x80000001         | Checksum      | : 0x3c28    |
| Age           | : 328                | Length        | : 44        |
| Pfx Options   | :                    | Metric        | : 200       |
| Dest Prefix   | : 3ffe::100:1800/120 |               |             |

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                       |               |             |
|---------------|-----------------------|---------------|-------------|
| Area Id       | : 0.0.0.1             | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.14 (14)       |               |             |
| LSA Type      | : IE Pfx              |               |             |
| Sequence No   | : 0x80000001          | Checksum      | : 0x8c07    |
| Age           | : 328                 | Length        | : 44        |
| Pfx Options   | :                     | Metric        | : 100       |
| Dest Prefix   | : 3ffe::3300:1a00/120 |               |             |

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                      |               |             |
|---------------|----------------------|---------------|-------------|
| Area Id       | : 0.0.0.1            | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.15 (15)      |               |             |
| LSA Type      | : IE Pfx             |               |             |
| Sequence No   | : 0x80000001         | Checksum      | : 0xffc4    |
| Age           | : 328                | Length        | : 44        |
| Pfx Options   | :                    | Metric        | : 100       |
| Dest Prefix   | : 3ffe::100:1a00/120 |               |             |

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                       |               |             |
|---------------|-----------------------|---------------|-------------|
| Area Id       | : 0.0.0.1             | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.16 (16)       |               |             |
| LSA Type      | : IE Pfx              |               |             |
| Sequence No   | : 0x80000001          | Checksum      | : 0xaf6a    |
| Age           | : 328                 | Length        | : 44        |
| Pfx Options   | :                     | Metric        | : 200       |
| Dest Prefix   | : 3ffe::3300:2d00/120 |               |             |

-----  
IE Pfx LSA for Area 0.0.0.1  
-----

|               |                 |               |             |
|---------------|-----------------|---------------|-------------|
| Area Id       | : 0.0.0.1       | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.17 (17) |               |             |
| LSA Type      | : IE Pfx        |               |             |
| Sequence No   | : 0x80000001    | Checksum      | : 0x2328    |
| Age           | : 328           | Length        | : 44        |
| Pfx Options   | :               | Metric        | : 200       |

```

Dest Prefix : 3ffe::100:2d00/120

IE Pfx LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.18 (18)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0x551c
Age : 328 Length : 44
Pfx Options : Metric : 100
Dest Prefix : 3ffe::3300:3800/120

IE Pfx LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.19 (19)
LSA Type : IE Pfx
Sequence No : 0x80000001 Checksum : 0xc8d9
Age : 328 Length : 44
Pfx Options : Metric : 100
Dest Prefix : 3ffe::100:3800/120

IE Rtr LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 10.20.1.2 (169083138)
LSA Type : IE Rtr
Sequence No : 0x80000001 Checksum : 0xcf8b
Age : 328 Length : 32
Options : ---R--EV6 Metric : 100
ASB Rtr Id : 10.20.1.2

IE Rtr LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 10.20.1.4 (169083140)
LSA Type : IE Rtr
Sequence No : 0x80000001 Checksum : 0xc32f
Age : 328 Length : 32
Options : ----R--EV6 Metric : 200
ASB Rtr Id : 10.20.1.4

IE Rtr LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 10.20.1.5 (169083141)
LSA Type : IE Rtr
Sequence No : 0x80000001 Checksum : 0xdb79
Age : 328 Length : 32
Options : ----R--EV6 Metric : 100
ASB Rtr Id : 10.20.1.5

IA Pfx LSA for Area 0.0.0.1

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.0 (0)
LSA Type : IA Pfx
Sequence No : 0x80000003 Checksum : 0xb8c1
Age : 293 Length : 112

```

|                 |                       |            |       |
|-----------------|-----------------------|------------|-------|
| Ref Ls Type     | : 2001                | Ref Ls Id  | : 0   |
| Ref Adv Rtr     | : 10.20.1.6           | No of Pfxs | : 4   |
| Prefix (1)      | : 3ffe::3300:2400/120 |            |       |
| Pfx Options (1) | :                     | Metric (1) | : 100 |
| Prefix (2)      | : 3ffe::100:2400/120  |            |       |
| Pfx Options (2) | :                     | Metric (2) | : 100 |
| Prefix (3)      | : 3ffe::3300:4206/128 |            |       |
| Pfx Options (3) | : LA                  | Metric (3) | : 0   |
| Prefix (4)      | : 3ffe::100:4206/128  |            |       |
| Pfx Options (4) | : LA                  | Metric (4) | : 0   |

-----  
Rtr Inf LSA for Area 0.0.0.1  
-----

|                                            |               |               |             |
|--------------------------------------------|---------------|---------------|-------------|
| Area Id                                    | : 0.0.0.1     | Adv Router Id | : 10.20.1.6 |
| Link State Id                              | : 0.0.0.0 (0) |               |             |
| LSA Type                                   | : Rtr Inf     |               |             |
| Sequence No                                | : 0x80000002  | Checksum      | : 0x4abf    |
| Age                                        | : 329         | Length        | : 68        |
| TLV dump:                                  |               |               |             |
| Capabilities (1) Len 4 :                   |               |               |             |
| 0x28000000                                 |               |               |             |
| SR algorithm (8) Len 2 :                   |               |               |             |
| 0x2 0x0                                    |               |               |             |
| SR label range (9) Len 12 :                |               |               |             |
| Range-size=1000                            |               |               |             |
| Sub-TLV SID/label(7) len 3 :               |               |               |             |
| label=70000                                |               |               |             |
| SR Readable label depth capability: 15     |               |               |             |
| non-OSPF Func Capabilities (32768) Len 4 : |               |               |             |
| 0x80000000                                 |               |               |             |

-----  
E-Router LSA for Area 0.0.0.1  
-----

|               |               |                    |             |
|---------------|---------------|--------------------|-------------|
| Area Id       | : 0.0.0.1     | Adv Router Id      | : 10.20.1.6 |
| Link State Id | : 0.0.0.0 (0) |                    |             |
| LSA Type      | : E-Router    |                    |             |
| Sequence No   | : 0x80000003  | Checksum           | : 0xdf0f    |
| Age           | : 293         | Length             | : 88        |
| Options       | : ----R--EV6  |                    |             |
| Flags         | : ABR ASBR    | Link Count         | : 2         |
| Link Type (1) | : P2P Link    | Nbr Rtr ID (1)     | : 10.20.1.2 |
| I/F Index (1) | : 2           | Nbr I/F Index (1): | 4           |
| Metric (1)    | : 100         |                    |             |
| Adj SID (1)   | : 524285      | Flags (1)          | : BVL--     |
|               |               | Weight (1)         | : 0         |
| Link Type (2) | : P2P Link    | Nbr Rtr ID (2)     | : 10.20.1.3 |
| I/F Index (2) | : 4           | Nbr I/F Index (2): | 3           |
| Metric (2)    | : 100         |                    |             |
| Adj SID (2)   | : 524284      | Flags (2)          | : BVL--     |
|               |               | Weight (2)         | : 0         |

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

|               |               |               |             |
|---------------|---------------|---------------|-------------|
| Area Id       | : 0.0.0.1     | Adv Router Id | : 10.20.1.6 |
| Link State Id | : 0.0.0.2 (2) |               |             |
| LSA Type      | : E-IE Pfx    |               |             |
| Sequence No   | : 0x80000001  | Checksum      | : 0xf5de    |
| Age           | : 334         | Length        | : 60        |
| Pfx Options   | : N LA        | Metric        | : 0         |

```

Dest Prefix : 3ffe::a14:106/128
Prefix SID : 9
Flags : NP-----
Algorithm : 0

```

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.7 (7)
LSA Type : E-IE Pfx
Sequence No : 0x80000001 Checksum : 0xb351
Age : 328 Length : 60
Pfx Options : N Metric : 200
Dest Prefix : 3ffe::100:2c04/128
Prefix SID : 1
Flags : NP-----
Algorithm : 0

```

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.11 (11)
LSA Type : E-IE Pfx
Sequence No : 0x80000001 Checksum : 0x4b20
Age : 328 Length : 60
Pfx Options : N Metric : 100
Dest Prefix : 3ffe::a14:102/128
Prefix SID : 11
Flags : NP-----
Algorithm : 0

```

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.20 (20)
LSA Type : E-IE Pfx
Sequence No : 0x80000001 Checksum : 0xcd81
Age : 328 Length : 60
Pfx Options : N Metric : 100
Dest Prefix : 3ffe::100:3705/128
Prefix SID : 2
Flags : NP-----
Algorithm : 0

```

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.21 (21)
LSA Type : E-IE Pfx
Sequence No : 0x80000001 Checksum : 0xc29e
Age : 328 Length : 60
Pfx Options : N Metric : 100
Dest Prefix : 3ffe::a14:105/128
Prefix SID : 8
Flags : NP-----
Algorithm : 0

```

-----  
E-IE Pfx LSA for Area 0.0.0.1  
-----

```

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 0.0.0.22 (22)
LSA Type : E-IE Pfx
Sequence No : 0x80000001 Checksum : 0x5ca1
Age : 328 Length : 60

```

|                               |                       |               |             |
|-------------------------------|-----------------------|---------------|-------------|
| Pfx Options                   | : N                   | Metric        | : 200       |
| Dest Prefix                   | : 3ffe::a14:104/128   |               |             |
| Prefix SID                    | : 7                   | Flags         | : NP-----   |
|                               |                       | Algorithm     | : 0         |
| -----                         |                       |               |             |
| E-IA Pfx LSA for Area 0.0.0.1 |                       |               |             |
| -----                         |                       |               |             |
| Area Id                       | : 0.0.0.1             | Adv Router Id | : 10.20.1.6 |
| Link State Id                 | : 0.0.0.0 (0)         |               |             |
| LSA Type                      | : E-IA Pfx            |               |             |
| Sequence No                   | : 0x80000002          | Checksum      | : 0x54d0    |
| Age                           | : 293                 | Length        | : 100       |
| Ref Ls Type                   | : a021                | Ref Ls Id     | : 0         |
| Ref Adv Rtr                   | : 10.20.1.6           | No of Pfxs    | : 2         |
| Prefix (1)                    | : 3ffe::3300:4206/128 |               |             |
| Pfx Options (1)               | : N LA                | Metric (1)    | : 0         |
| Prefix (2)                    | : 3ffe::100:4206/128  |               |             |
| Pfx Options (2)               | : N LA                | Metric (2)    | : 0         |
| Prefix SID (2)                | : 3                   | Flags (2)     | : NP-----   |
|                               |                       | Algorithm (2) | : 0         |
| =====                         |                       |               |             |

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i>   <i>ip-address</i>   <b>area</b> <i>area-id</i> ] [ <b>detail</b> ]<br><b>interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ] <b>database</b> [ <b>detail</b> ]                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | Displays the details of the OSPF interface. This interface can be identified by its IP address or IP interface name. When neither is specified, all in-service interfaces are displayed.<br><br>The <b>detail</b> option produces a great amount of data. It is recommended to detail only when requesting a specific interface.                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>ip-int-name</i> — Displays only the interface identified by this interface name, up to a maximum of 32 characters.<br><i>ip-address</i> — Displays only the interface identified by this IPv4 or IPv6 address.<br><br><b>Values</b> ipv4-address:<br>• a.b.c.d<br>ipv6-address:<br>• x:x:x:x:x:x:x (eight 16-bit pieces)<br>• x:x:x:x:x:d.d.d.d<br>• x: [0 to FFFF]H<br>• d: [0 to 255]D<br><br><b>area</b> <i>area-id</i> — Displays all interfaces configured in this area.<br><br><b>Values</b> ip-address — a.b.c.d |

area — 0 to 4294967295

**detail** — Displays detailed information on the interface.

**database** — Displays database information for the interface.

**Output** Standard OSPF Interface Output

The following table describes the standard command output fields for an OSPF interface.

**Table 18** OSPF Interface Output Fields

| Label                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| If Name                | The interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Area Id                | A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| D Rtr Id               | The IP Interface address of the router identified as the Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| BD Rtr Id              | The IP Interface address of the router identified as the Backup Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Backup Designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Adm                    | Dn — OSPF on this interface is administratively shut down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                        | Up — OSPF on this interface is administratively enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Opr                    | Down — This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.<br>Wait — The router is trying to determine the identity of the (Backup) Designated router for the network.<br>PToP — The interface is operational, and connects either to a physical point-to-point network or to a virtual link.<br>DR — This router is the Designated Router for this network.<br>BDR — This router is the backup Designated Router for this network.<br>ODR — The interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the Designated Router. |
| No. of OSPF Interfaces | The number of interfaces listed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### Sample Output

```
*A:Dut-C# show router ospf interface "system"
```

```

=====
Rtr Base OSPFv2 Instance 0 Interface "system"
=====
If Name Area Id Designated Rtr Bkup Desig Rtr Adm Oper

system 0.0.0.0 10.20.1.3 0.0.0.0 Up DR

No. of OSPF Interfaces: 1
=====

*A:Dut-C# show router ospf interface "system" detail
=====
Rtr Base OSPFv2 Instance 0 Interface "system" (detail)
=====

Configuration

IP Address : 10.20.1.3
node sid index : 33
Area Id : 0.0.0.0 Priority : 1
Hello Intrvl : 5 sec Rtr Dead Intrvl : 15 sec
Retrans Intrvl : 5 sec Poll Intrvl : 120 sec
Cfg Metric : 0 Advert Subnet : True
Transit Delay : 1 Cfg IF Type : None
Passive : True Cfg MTU : 0
LSA-filter-out : None Adv Rtr Capab : Yes
LFA : Include LFA NH Template :
RIB-priority : None
Auth Type : None

State

Admin Status : Enabled Oper State : Designated Rtr
Designated Rtr : 10.20.1.3 Backup Desig Rtr : 0.0.0.0
IF Type : Broadcast Network Type : Stub
Oper MTU : 1500 Last Enabled : 05/27/2015 08:35:53
Oper Metric : 0 Bfd Enabled : No
Te Metric : 0 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0

Statistics

Nbr Count : 0 If Events : 4
Tot Rx Packets : 0 Tot Tx Packets : 0
Rx Hellos : 0 Tx Hellos : 0
Rx DBDs : 0 Tx DBDs : 0
Rx LSRS : 0 Tx LSRS : 0
Rx LSUs : 0 Tx LSUs : 0
Rx LS Acks : 0 Tx LS Acks : 0
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 0
Bad Dead Int. : 0 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0

```



```

LSA Count : 0 LSA Checksum : 0x0
=====
*A:Dut-C#

*A:Dut-C# show router ospf 1 interface "DUTC_TO_DUTB.1.0"
=====
Rtr Base OSPFv2 Instance 1 Interface "DUTC_TO_DUTB.1.0"
=====
If Name Area Id Designated Rtr Bkup Desig Rtr Adm Oper

DUTC_TO_DUTB.1.0 0.0.0.0 0.0.0.0 0.0.0.0 Up PToP

No. of OSPF Interfaces: 1
=====

*A:Dut-C# show router ospf 1 interface "DUTC_TO_DUTB.1.0" detail
=====
Rtr Base OSPFv2 Instance 1 Interface "DUTC_TO_DUTB.1.0" (detail)
=====

Configuration

IP Address : 1.0.23.3
Area Id : 0.0.0.0 Priority : 1
Hello Intrvl : 2 sec Rtr Dead Intrvl : 10 sec
Retrans Intrvl : 5 sec Poll Intrvl : 120 sec
Cfg Metric : 7000 Advert Subnet : True
Transit Delay : 1 Cfg IF Type : Point To Point
Passive : False Cfg MTU : 0
LSA-filter-out : None Adv Rtr Capab : Yes
LFA : Include LFA NH Template : template1
Auth Type : None

State

Admin Status : Enabled Oper State : Point To Point
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
IF Type : Point To Point Network Type : Transit
Oper MTU : 1500 Last Enabled : 01/14/2014 14:33:07
Oper Metric : 7000 Bfd Enabled : No
Te Metric : 7000 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0

Statistics

Nbr Count : 1 If Events : 1
Tot Rx Packets : 603 Tot Tx Packets : 602
Rx Hellos : 576 Tx Hellos : 577
Rx DBDs : 3 Tx DBDs : 2
Rx LSRS : 0 Tx LSRS : 1
Rx LSUs : 15 Tx LSUs : 16
Rx LS Acks : 9 Tx LS Acks : 6
Retransmits : 2 Discards : 2
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0

```

```

Bad Lengths : 0
Bad Dead Int. : 1
Bad Versions : 0
LSA Count : 0
Bad Hello Int. : 1
Bad Options : 0
Bad Checksums : 0
LSA Checksum : 0x0
=====
*A:Dut-C#

A:SetupCLI# show router ospf 1 interface "ip_if_1" detail
=====
Rtr Base OSPFv2 Instance 1 Interface "ip_if_1" (detail)
=====

Configuration

IP Address : 10.10.1.1
Area Id : 0.0.0.0
Hello Intrvl : 9 sec
Retrans Intrvl : 10 sec
Cfg Metric : 11
Transit Delay : 2
Passive : False
LFA : Exclude
IPsec InStatSA :
IPsec InStatSAmp:
IPsec OutStatSA :

State

Admin Status : Enabled
Designated Rtr : 0.0.0.0
IF Type : Secondary
Oper MTU : 1576
Oper Metric : 11
Te Metric : 16777215
Admin Groups : None
Ldp Sync : outOfService
Ldp Timer State : Disabled
Oper State : Down
Backup Desig Rtr : 0.0.0.0
Network Type : Stub
Last Enabled : Never
Bfd Enabled : No
Te State : Down
Ldp Sync Wait : Disabled
Ldp Tm Left : 0

Statistics

Nbr Count : 0
Tot Rx Packets : 0
Rx Hellos : 0
Rx DBDs : 0
Rx LSRs : 0
Rx LSUs : 0
Rx LS Acks : 0
Retransmits : 0
Bad Networks : 0
Bad Areas : 0
Bad Auth Types : 0
Bad Neighbors : 0
Bad Lengths : 0
Bad Dead Int. : 0
Bad Versions : 0
LSA Count : 0
If Events : 0
Tot Tx Packets : 0
Tx Hellos : 0
Tx DBDs : 0
Tx LSRs : 0
Tx LSUs : 0
Tx LS Acks : 0
Discards : 0
Bad Virt Links : 0
Bad Dest Addrs : 0
Auth Failures : 0
Bad Pkt Types : 0
Bad Hello Int. : 0
Bad Options : 0
Bad Checksums : 0
LSA Checksum : 0x0

Configuration

IP Address : 10.10.1.1

```

```

Area Id : 1.1.1.1 Priority : 10
Hello Intrvl : 9 sec Rtr Dead Intrvl : 45 sec
Retrans Intrvl : 10 sec Poll Intrvl : 120 sec
Cfg Metric : 11 Advert Subnet : False
Transit Delay : 2 Auth Type : MD5
Passive : False Cfg MTU : 9198
LFA : Exclude
IPsec InStatSA : IPsec OutStatSA :
IPsec InStatSATmp:

```

-----  
State

```

Admin Status : Enabled Oper State : Down
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
IF Type : Point To Point Network Type : Stub
Oper MTU : 1576 Last Enabled : Never
Oper Metric : 11 Bfd Enabled : No
Te Metric : 16777215 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0

```

-----  
Statistics

```

Nbr Count : 0 If Events : 0
Tot Rx Packets : 0 Tot Tx Packets : 0
Rx Hellos : 0 Tx Hellos : 0
Rx DBDs : 0 Tx DBDs : 0
Rx LSRs : 0 Tx LSRs : 0
Rx LSUs : 0 Tx LSUs : 0
Rx LS Acks : 0 Tx LS Acks : 0
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Adrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 0
Bad Dead Int. : 0 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0
LSA Count : 0 LSA Checksum : 0x0

```

=====

A:SetupCLI#

A:SetupCLI# show router ospf 1 interface area 1.1.1.1 detail

=====

Rtr Base OSPFv2 Instance 1 Interfaces in area 1.1.1.1 (detail)

-----

Interface : ip\_if\_1

```

IP Address : 10.10.1.1 Priority : 10
Area Id : 1.1.1.1 Rtr Dead Intrvl : 45 sec
Hello Intrvl : 9 sec Poll Intrvl : 120 sec
Retrans Intrvl : 10 sec Advert Subnet : False
Cfg Metric : 11 Auth Type : MD5
Transit Delay : 2 Cfg MTU : 9198
Passive : False LFA : Exclude
LFA : Exclude
IPsec InStatSA : IPsec OutStatSA :

```

```

IPsec InStatSATmp:
Admin Status : Enabled Oper State : Down
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
IF Type : Point To Point Network Type : Stub
Oper MTU : 1576 Last Enabled : Never
Oper Metric : 11 Bfd Enabled : No
Te Metric : 16777215 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0
Nbr Count : 0 If Events : 0
Tot Rx Packets : 0 Tot Tx Packets : 0
Rx Hellos : 0 Tx Hellos : 0
Rx DBDs : 0 Tx DBDs : 0
Rx LSRs : 0 Tx LSRs : 0
Rx LSUs : 0 Tx LSUs : 0
Rx LS Acks : 0 Tx LS Acks : 0
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 0
Bad Dead Int. : 0 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0
LSA Count : 0 LSA Checksum : 0x0
=====
A:SetupCLI#

A:SetupCLI# show router ospf 1 interface detail
=====
Rtr Base OSPFv2 Instance 1 Interfaces (detail)
=====

Interface : system

IP Address : 9.1.255.255
Area Id : 0.0.0.0 Priority : 1
Hello Intrvl : 10 sec Rtr Dead Intrvl : 40 sec
Retrans Intrvl : 5 sec Poll Intrvl : 120 sec
Cfg Metric : 0 Advert Subnet : True
Transit Delay : 1 Auth Type : None
Passive : True Cfg MTU : 0
Admin Status : Enabled Oper State : Designated Rtr
Designated Rtr : 2.2.2.2 Backup Desig Rtr : 0.0.0.0
IF Type : Broadcast Network Type : Transit
Oper MTU : 1500 Last Enabled : 05/14/2006 09:16:26
Oper Metric : 0 Bfd Enabled : No
Nbr Count : 0 If Events : 5
Tot Rx Packets : 0 Tot Tx Packets : 0
Rx Hellos : 0 Tx Hellos : 0
Rx DBDs : 0 Tx DBDs : 0
Rx LSRs : 0 Tx LSRs : 0
Rx LSUs : 0 Tx LSUs : 0
Rx LS Acks : 0 Tx LS Acks : 0
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0

```

```

Bad Neighbors : 0
Bad Lengths : 0
Bad Dead Int. : 0
Bad Versions : 0
LSA Count : 0
Bad Pkt Types : 0
Bad Hello Int. : 0
Bad Options : 0
Bad Checksums : 0
LSA Checksum : 0x0

Interface : sender

IP Address : 11.1.1.1
Area Id : 0.0.0.0
Hello Intrvl : 10 sec
Retrans Intrvl : 5 sec
Cfg Metric : 0
Transit Delay : 1
Passive : False
Priority : 1
Rtr Dead Intrvl : 40 sec
Poll Intrvl : 120 sec
Advert Subnet : True
Auth Type : None
Cfg MTU : 0
=====
A:SetupCLI#

```

#### Detailed OSPF Interface Output

The following table describes the detailed command output fields for an OSPF interface.

**Table 19 Detailed OSPF Interface Output Fields**

| Label           | Description                                                                                                                                                                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface       | The IP address of this OSPF interface.                                                                                                                                                                                                              |
| IP Address      | The IP address and mask of this OSPF interface.                                                                                                                                                                                                     |
| Interface Name  | The interface name.                                                                                                                                                                                                                                 |
| Area Id         | A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.                                                                                                                 |
| Priority        | The priority of this interface. Used in multi-access networks, this field is used in the designated router election algorithm.                                                                                                                      |
| Hello Intrvl    | The length of time, in seconds, between the Hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network.                                                                         |
| Rtr Dead Intrvl | The number of seconds that a router's Hello packets have not been seen before it's neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for all routers attached to a common network. |
| Retrans Intrvl  | The number of seconds between link-state advertisement retransmissions, for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets.                               |
| Poll Intrvl     | The larger time interval, in seconds, between the Hello packets sent to an inactive non-broadcast multi-access neighbor.                                                                                                                            |

**Table 19 Detailed OSPF Interface Output Fields (Continued)**

| Label         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metric        | The metric to be advertised for this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Advert Subnet | <p>False — When a point-to-point interface is configured as false, then the subnet is not advertised and the endpoints are advertised as host routes.</p> <p>True — When a point-to-point interface is configured to true, then the subnet is advertised.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| Transit Delay | The estimated number of seconds it takes to transmit a link state update packet over this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Auth Type     | <p>Identifies the authentication procedure to be used for the packet.</p> <p>None — Routing exchanges over the network/subnet are not authenticated.</p> <p>Simple — A 64-bit field is configured on a per-network basis. All packets sent on a particular network must have this configured value in their OSPF header 64-bit authentication field. This essentially serves as a “clear” 64-bit password.</p> <p>MD5 — A shared secret key is configured in all routers attached to a common network/subnet. For each OSPF protocol packet, the key is used to generate/verify a “message digest” that is appended to the end of the OSPF packet.</p> |
| Passive       | <p>False — This interfaces operates as a normal OSPF interface with regard to adjacency forming and network/link behavior.</p> <p>True — No OSPF HELLOs will be sent out on this interface and the router advertises this interface as a stub network/link in its router LSAs.</p>                                                                                                                                                                                                                                                                                                                                                                     |
| MTU           | The desired size of the largest packet which can be sent/received on this OSPF interface, specified in octets. This size does include the underlying IP header length, but not the underlying layer headers/trailers.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Admin Status  | Disabled — OSPF on this interface is administratively shut down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|               | Enabled — OSPF on this interface is administratively enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 19 Detailed OSPF Interface Output Fields (Continued)**

| Label        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oper State   | <p>Down — This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.</p> <p>Waiting — The router is trying to determine the identity of the (Backup) Designated router for the network.</p> <p>Point To Point — The interface is operational, and connects either to a physical point-to-point network or to a virtual link.</p> <p>Designated Rtr — This router is the Designated Router for this network.</p> <p>Other Desig Rtr — The interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the Designated Router.</p> <p>Backup Desig Rtr — This router is the Backup Designated Router for this network.</p> |
| DR-Id        | The IP Interface address of the router identified as the Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Designated router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| BDR-Id       | The IP Interface address of the router identified as the Backup Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Backup Designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| IF Type      | <p>Broadcast — LANs, such as Ethernet.</p> <p>Non-Broadcast — X.25, Frame Relay and similar technologies.</p> <p>Point-To-Point — Point-to-point links.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Network Type | Stub — OSPF has not established a neighbor relationship with any other OSPF router on this network as such only traffic sourced or destined to this network will be routed to this network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|              | Transit — OSPF has established at least one neighbor relationship with any other OSPF router on this network as such traffic en route to other networks may be routed via this network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Oper MTU     | The operational size of the largest packet which can be sent/received on this OSPF interface, specified in octets. This size does include the underlying IP header length, but not the underlying layer headers/trailers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Last Enabled | The time that this interface was last enabled to run OSPF on this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Nbr Count    | The number of OSPF neighbors on the network for this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| If Events    | The number of times this OSPF interface has changed its state, or an error has occurred since this interface was last enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 19 Detailed OSPF Interface Output Fields (Continued)**

| Label          | Description                                                                                                                                                |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tot Rx Packets | The total number of OSPF packets received on this interface since this interface was last enabled.                                                         |
| Tot Tx Packets | The total number of OSPF packets transmitted on this interface since this interface was last enabled.                                                      |
| Rx Hellos      | The total number of OSPF Hello packets received on this interface since this interface was last enabled.                                                   |
| Tx Hellos      | The total number of OSPF Hello packets transmitted on this interface since this interface was last enabled.                                                |
| Rx DBDs        | The total number of OSPF database description packets received on this interface since this interface was last enabled.                                    |
| Tx DBDs        | The total number of OSPF database description packets transmitted on this interface since this interface was last enabled.                                 |
| Rx LSRs        | The total number of Link State Requests (LSRs) received on this interface since this interface was last enabled.                                           |
| Tx LSRs        | The total number of Link State Requests (LSRs) transmitted on this interface since this interface was last enabled.                                        |
| Rx LSUs        | The total number of Link State Updates (LSUs) received on this interface since this interface was last enabled.                                            |
| Tx LSUs        | The total number of Link State Updates (LSUs) transmitted on this interface since this interface was last enabled.                                         |
| Rx LS Acks     | The total number of Link State Acknowledgments received on this interface since this interface was last enabled.                                           |
| Tx LS Acks     | The total number of Link State Acknowledgments transmitted on this interface since this interface was last enabled.                                        |
| Retransmits    | The total number of OSPF Retransmits sent on this interface since this interface was last enabled.                                                         |
| Discards       | The total number of OSPF packets discarded on this interface since this interface was last enabled.                                                        |
| Bad Networks   | The total number of OSPF packets received with invalid network or mask since this interface was last enabled.                                              |
| Bad Virt Links | The total number of OSPF packets received on this interface that are destined to a virtual link that does not exist since this interface was last enabled. |
| Bad Areas      | The total number of OSPF packets received with an area mismatch since this interface was last enabled.                                                     |



**Table 19 Detailed OSPF Interface Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bad Dest Addrs | The total number of OSPF packets received with the incorrect IP destination address since this interface was last enabled.                                                                                                                               |
| Bad Auth Types | The total number of OSPF packets received with an invalid authorization type since this interface was last enabled.                                                                                                                                      |
| Auth Failures  | The total number of OSPF packets received with an invalid authorization key since this interface was last enabled.                                                                                                                                       |
| Bad Neighbors  | The total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since this interface was last enabled.                                                                          |
| Bad Pkt Types  | The total number of OSPF packets received with an invalid OSPF packet type since this interface was last enabled.                                                                                                                                        |
| Bad Lengths    | The total number of OSPF packets received on this interface with a total length not equal to the length given in the packet itself since this interface was last enabled.                                                                                |
| Bad Hello int. | The total number of OSPF packets received where the hello interval given in packet was not equal to that configured on this interface since this interface was last enabled.                                                                             |
| Bad Dead Int.  | The total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since this interface was last enabled.                                                                          |
| Bad Options    | The total number of OSPF packets received with an option that does not match those configured for this interface or area since this interface was last enabled.                                                                                          |
| Bad Versions   | The total number of OSPF packets received with bad OSPF version numbers since this interface was last enabled.                                                                                                                                           |
| Te Metric      | Indicates the TE metric configured for this interface. This metric is flooded out in the TE metric sub-tlv in the OSPF TE LSAs. Depending on the configuration, either the TE metric value or the native OSPF metric value is used in CSPF computations. |
| Te State       | Indicates the MPLS interface TE status from OSPF standpoint.                                                                                                                                                                                             |
| Admin Groups   | Indicates the bit-map inherited from MPLS interface that identifies the admin groups to which this interface belongs.                                                                                                                                    |
| Ldp Sync       | Specifies whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol.                                                                                                               |
| Ldp Sync Wait  | Indicates the time to wait for the LDP adjacency to come up.                                                                                                                                                                                             |

**Table 19 Detailed OSPF Interface Output Fields (Continued)**

| Label           | Description                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------|
| Ldp Timer State | Indicates the state of the LDP sync time left on the OSPF interface.                              |
| Ldp Tm Left     | Indicates the time left before OSPF reverts back to advertising normal metric for this interface. |

**Sample Output**

```

*A:JC-NodeA# show router ospf 1 interface area 1 detail
=====
Rtr Base OSPFv2 Instance 1 Interfaces in area 1 (detail)
=====
Interface : ip-10.10.1.1

IP Address : 10.10.1.1
Area Id : 0.0.0.1 Priority : 1
Hello Intrvl : 5 sec Rtr Dead Intrvl : 15 sec
Retrans Intrvl : 5 sec Poll Intrvl : 120 sec
Cfg Metric : 0 Advert Subnet : True
Transit Delay : 1 Auth Type : None
Passive : False Cfg MTU : 0
Admin Status : Enabled Oper State : Designated Rtr
Designated Rtr : 10.20.1.1 Backup Desig Rtr : 0.0.0.0
IF Type : Broadcast Network Type : Transit
Oper MTU : 1500 Last Enabled : 04/11/2007 16:06:27
Oper Metric : 1000 Bfd Enabled : No
Nbr Count : 0 If Events : 5
Tot Rx Packets : 0 Tot Tx Packets : 1116
Rx Hellos : 0 Tx Hellos : 1116
Rx DBDs : 0 Tx DBDs : 0
Rx LSRS : 0 Tx LSRS : 0
Rx LSUs : 0 Tx LSUs : 0
Rx LS Acks : 0 Tx LS Acks : 0
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 0
Bad Dead Int. : 0 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0
LSA Count : 0 LSA Checksum : 0x0
TE Metric : 678
=====
*A:JC-NodeA#

```

lfa-coverage

**Syntax lfa-coverage**

**Context** show>router>ospf  
show>router>ospf3

**Description** This command displays OSPFv2 or OSPFv3 Loop-Free Alternate (LFA) next-hop information.

**Output**

### Sample Output

```
*A:Dut-A# show router ospf 1 lfa-coverage
=====
Rtr Base OSPFv2 Instance 1 LFA Coverage
=====
Area Node Prefix

0.0.0.0 4/4 (100%) 8/8 (100%)
=====
*A:Dut-A#

*A:Dut-C>config>router>ospf3# show router ospf3 0 lfa-coverage
=====
Rtr Base OSPFv3 Instance 0 LFA Coverage
=====
Area Node Prefix

0.0.0.1 1/3 (33%) 9/27 (33%)
=====
```

## neighbor

**Syntax** neighbor [*ip-int-name* | *ip-address*] [**detail**]  
neighbor [*ip-int-name*] [*router-id*] [**detail**]  
neighbor overview  
neighbor [**remote** *ip-address*] [**detail**]

**Context** show>router>ospf  
show>router>ospf3

**Description** This command will display all neighbor information. To reduce the amount of output the user may opt to select the neighbors on a given interface by address or name.

The **detail** option produces a large amount of data. It is recommended to use **detail** only when requesting a specific neighbor.

**Parameters** *ip-int-name* — Displays neighbor information only for neighbors of the interface identified by the interface name.  
*ip-address* — Displays neighbor information only for neighbors of the interface identified by the IP address.

*router-id* — Displays neighbor information for the neighbor identified by the specified router ID in the **show>router>ospf3** context.

**remote** — Displays neighbor information for the neighbor identified by the specified remote IP address in the **show>router>ospf** context.

**Output** Standard OSPF Neighbor Output

The following table describes the standard command output fields for an OSPF neighbor.

**Table 20 OSPF Neighbor Output Fields**

| Label       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nbr IP Addr | The IP address this neighbor is using in its IP Source Address. On addressless links, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Nbr Rtr Id  | A 32-bit integer uniquely identifying the neighboring router in the Autonomous System.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Nbr State   | <p>Down — This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor.</p> <p>Attempt — This state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor.</p> <p>Init — In this state, an Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (the router itself did not appear in the neighbor's Hello packet).</p> <p>Two Way — In this state, communication between the two routers is bidirectional.</p> <p>ExchStart — This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Descriptor sequence number.</p> <p>Exchange — In this state the router is describing its entire link state database by sending Database Description packets to the neighbor.</p> <p>Loading — In this state, Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state.</p> <p>Full — In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.</p> |

**Table 20 OSPF Neighbor Output Fields (Continued)**

| Label            | Description (Continued)                                                                                                                                                                        |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority         | The priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network. |
| RetxQ Len        | The current length of the retransmission queue.                                                                                                                                                |
| Dead Time        | The time until this neighbor is declared down, this timer is set to the dead router interval when a valid hello packet is received from the neighbor.                                          |
| No. of Neighbors | The number of adjacent OSPF neighbors on this interface.                                                                                                                                       |

### Sample Output

```

A:ALA-A# show router ospf 1 neighbor
=====
Rtr Base OSPFv2 Instance 1 Neighbors
=====
Interface-Name Rtr Id State Pri RetxQ TTL

pc157-2/1 10.13.8.158 Full 1 0 37
pc157-2/2 10.13.7.165 Full 100 0 33
pc157-2/3 10.13.6.188 Full 1 0 38

No. of Neighbors: 3
=====
A:ALA-A#

```

### Detailed OSPF Neighbor Output

The following table describes the detailed command output fields for an OSPF neighbor.

**Table 21 Detailed OSPF Neighbor Output Fields**

| Label            | Description                                                                                                                                                                        |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Neighbor IP Addr | The IP address this neighbor is using in its IP source address. On addressless links, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.           |
| Local IF IP Addr | The IP address of this OSPF interface.                                                                                                                                             |
| Area Id          | A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone                                                 |
| Designated Rtr   | The IP Interface address of the router identified as the Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no Designated router. |

**Table 21 Detailed OSPF Neighbor Output Fields (Continued)**

| Label            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Neighbor Rtr Id  | A 32-bit integer uniquely identifying the neighboring router in the AS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Neighbor State   | <p>Down — This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor</p> <p>Attempt — This state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor.</p> <p>Init — In this state, an Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (the router itself did not appear in the neighbor's Hello packet).</p> <p>Two Way — In this state, communication between the two routers is bidirectional.</p> <p>Exchange start — This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Descriptor sequence number.</p> <p>Exchange — In this state the router is describing its entire link state database by sending Database Description packets to the neighbor.</p> <p>Loading — In this state, Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state.</p> <p>Full — In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.</p> |
| Priority         | The priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Retrans Q Length | The current length of the retransmission queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Options          | <p>E — External Routes Support</p> <p>MC — Multicast Support</p> <p>N/P — Type 7 LSA Support</p> <p>EA — External Attribute LSA Support</p> <p>DC — Demand Circuit Support</p> <p>O — Opaque LSA Support</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 21 Detailed OSPF Neighbor Output Fields (Continued)**

| Label             | Description                                                                                                                                                                                      |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Desig Rtr  | The IP Interface address of the router identified as the Backup Designated Router for the network in which this interface is configured. Set to 0.0.0.0 if there is no backup designated router. |
| Events            | The number of times this neighbor relationship has changed state, or an error has occurred.                                                                                                      |
| Last Event Time   | The time when the last event occurred that affected the adjacency to the neighbor.                                                                                                               |
| Up Time           | This value represents the uninterrupted time, in hundredths of seconds, the adjacency to this neighbor has been up. To evaluate when the last state change occurred see last event time.         |
| Time Before Dead  | The time until this neighbor is declared down, this timer is set to the dead router interval when a valid hello packet is received from the neighbor.                                            |
| Bad Nbr States    | The total number of OSPF packets received when the neighbor state was not expecting to receive this packet type since this interface was last enabled.                                           |
| LSA Inst fails    | The total number of times an LSA could not be installed into the LSDB due to a resource allocation issue since this interface was last enabled.                                                  |
| Bad Seq Nums      | The total number of times when a database description packet was received with a sequence number mismatch since this interface was last enabled.                                                 |
| Bad MTUs          | The total number of times when the MTU in a received database description packet was larger than the MTU of the receiving interface since this interface was last enabled.                       |
| Bad Packets       | The total number of times when an LS update was received with an illegal LS type or an option mismatch since this interface was last enabled.                                                    |
| LSA not in LSDB   | The total number of times when an LS request was received for an LSA not installed in the LSDB of this router since this interface was last enabled.                                             |
| Option Mismatches | The total number of times when a LS update was received with an option mismatch since this interface was last enabled.                                                                           |
| Nbr Duplicates    | The total number of times when a duplicate database description packet was received during the exchange state since this interface was last enabled.                                             |

**Sample Output**

```

*A:Dut-C# show router ospf neighbor detail
=====
Rtr Base OSPFv2 Instance 0 Neighbors (detail)
=====

Neighbor Rtr Id : 10.20.1.1 Interface: to_Dut-A

Neighbor IP Addr : 1.1.3.1
Local IF IP Addr : 1.1.3.3
Area Id : 0.0.0.1 Adj SR SID : Label 262141
Designated Rtr : 10.20.1.3 Backup Desig Rtr : 10.20.1.1
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : - E - - - - O --
Events : 5 Last Event Time : 05/27/2015 08:36:02
Up Time : 0d 00:11:01 Time Before Dead : 8 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.20.1.2 Interface: to_Dut-B1

Neighbor IP Addr : 1.2.3.2
Local IF IP Addr : 1.2.3.3
Area Id : 0.0.0.1 Adj SR SID : Label 262139
Designated Rtr : 10.20.1.3 Backup Desig Rtr : 10.20.1.2
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : - E - - - - O --
Events : 6 Last Event Time : 05/27/2015 08:36:03
Up Time : 0d 00:11:03 Time Before Dead : 10 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.20.1.2 Interface: to_Dut-B2

Neighbor IP Addr : 2.2.3.2
Local IF IP Addr : 2.2.3.3
Area Id : 0.0.0.0 Adj SR SID : Label 262138
Designated Rtr : 10.20.1.3 Backup Desig Rtr : 10.20.1.2
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : - E - - - - O --
Events : 5 Last Event Time : 05/27/2015 08:36:03
Up Time : 0d 00:11:01 Time Before Dead : 9 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0

```



```

Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.20.1.5 Interface: to_Dut-E

Neighbor IP Addr : 1.3.5.5
Local IF IP Addr : 1.3.5.3
Area Id : 0.0.0.0 Adj SR SID : Label 262140
Designated Rtr : 10.20.1.5 Backup Desig Rtr : 10.20.1.3
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : - E - - - - O --
Events : 7 Last Event Time : 05/27/2015 08:36:04
Up Time : 0d 00:11:01 Time Before Dead : 8 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 0 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never
=====
*A:Dut-C#

A:ALA-A# show router ospf 1 neighbor detail
=====
Rtr Base OSPFv2 Instance 1 Neighbors (detail)

Neighbor Rtr Id : 10.13.8.158 Interface: pc157-2/1

Neighbor IP Addr : 10.16.1.8
Local IF IP Addr : 10.16.1.7
Area Id : 0.0.0.0
Designated Rtr : 0.0.0.0 Backup Desig Rtr : 0.0.0.0
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : -E--O-
Events : 4 Last Event Time : 05/06/2006 00:11:16
Up Time : 1d 18:20:20 Time Before Dead : 38 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.13.7.165 Interface: pc157-2/2

Neighbor IP Addr : 10.12.1.3
Local IF IP Addr : 10.12.1.7
Area Id : 0.0.0.0
Designated Rtr : 10.13.9.157 Backup Desig Rtr : 10.13.7.165
Neighbor State : Full Priority : 100
Retrans Q Length : 0 Options : -E--O-
Events : 4 Last Event Time : 05/05/2006 01:39:13
Up Time : 0d 16:52:27 Time Before Dead : 33 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 0 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0

```

```

Option Mismatches: 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never

Neighbor Rtr Id : 10.13.6.188 Interface: pc157-2/3

Neighbor IP Addr : 10.14.1.4
Local IF IP Addr : 10.14.1.7
Area Id : 0.0.0.0
Designated Rtr : 10.13.9.157 Backup Desig Rtr : 10.13.6.188
Neighbor State : Full Priority : 1
Retrans Q Length : 0 Options : -E--O-
Events : 4 Last Event Time : 05/05/2006 08:35:18
Up Time : 0d 09:56:25 Time Before Dead : 38 sec
GR Helper : Not Helping GR Helper Age : 0 sec
GR Exit Reason : None GR Restart Reason: Unknown
Bad Nbr States : 1 LSA Inst fails : 0
Bad Seq Nums : 0 Bad MTUs : 0
Bad Packets : 0 LSA not in LSDB : 0
Option Mismatches : 0 Nbr Duplicates : 0
Num Restarts : 0 Last Restart at : Never
=====
A:ALA-A#

```

## opaque-database

- Syntax** **opaque-database** [**area** *area-id* | **as**] [**adv-router** *router-id*] [*/s-id*] [**detail**]
- Context** show>router>ospf
- Description** This command displays OSPF opaque database information.
- Parameters** **area** *area-id* — Displays all opaque databases configured in this area.
- Values** ip-address — a.b.c.d  
area — 0 to 4294967295
- as** — Displays opaque databases configured in the autonomous system (AS).
- adv-router** *router-id* [*/s-id*] — Displays opaque database information associated with the specified advertising router. To further narrow the number of items displayed, the */s-id* can optionally be specified.
- Output** OSPF Opaque Database Output
- The following table describes the OSPF opaque database output fields.

**Table 22** OSPF Opaque Database Output Fields

| Label   | Description                                                                                   |
|---------|-----------------------------------------------------------------------------------------------|
| Area Id | A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone. |

**Table 22 OSPF Opaque Database Output Fields (Continued)**

| Label         | Description                                                                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type          | NSSA — This area is configured as an NSSA area.<br>Area — This area is configured as a standard area (not NSSA or stub).<br>Stub — This area is configured as an NSSA area.         |
| Link State Id | The link state ID is an LSA type specific field containing either a Router-Id or an IP Address; it identifies the piece of the routing domain being described by the advertisement. |
| Adv Rtr Id    | The router identifier of the router advertising the LSA.                                                                                                                            |
| Age           | The age of the link state advertisement in seconds.                                                                                                                                 |
| Sequence      | The signed 32-bit integer sequence number.                                                                                                                                          |
| Cksum         | The 32-bit unsigned sum of the link-state advertisements' LS checksums.                                                                                                             |

### Sample Output

```
*A:Dut-C# show router ospf opaque-database
=====
Rtr Base OSPFv2 Instance 0 Opaque Link State Database (type : All)
=====
Type Id Link State Id Adv Rtr Id Age Sequence Cksum

Area 0.0.0.0 4.0.0.0 10.20.1.2 740 0x80000002 0x5653
Area 0.0.0.0 7.0.0.2 10.20.1.2 745 0x80000001 0xee35
Area 0.0.0.0 7.16.0.6 10.20.1.2 725 0x80000002 0xe434
Area 0.0.0.0 8.0.0.6 10.20.1.2 730 0x80000002 0x5f1d
Area 0.0.0.0 8.0.0.7 10.20.1.2 731 0x80000002 0xbeb8
Area 0.0.0.0 4.0.0.0 10.20.1.3 739 0x80000002 0x6dd6
Area 0.0.0.0 7.0.0.2 10.20.1.3 744 0x80000001 0x1601
Area 0.0.0.0 7.16.0.2 10.20.1.3 734 0x80000001 0x914
Area 0.0.0.0 8.0.0.6 10.20.1.3 728 0x80000002 0x8ac1
Area 0.0.0.0 8.0.0.7 10.20.1.3 729 0x80000002 0xf57b
Area 0.0.0.0 4.0.0.0 10.20.1.4 740 0x80000002 0x15ba
Area 0.0.0.0 7.0.0.2 10.20.1.4 745 0x80000001 0x3dcc
Area 0.0.0.0 7.16.0.3 10.20.1.4 736 0x80000001 0xda04
Area 0.0.0.0 8.0.0.4 10.20.1.4 732 0x80000002 0xfe4a
Area 0.0.0.0 8.0.0.5 10.20.1.4 732 0x80000002 0x4f1f
Area 0.0.0.0 4.0.0.0 10.20.1.5 738 0x80000002 0x746e
Area 0.0.0.0 7.0.0.2 10.20.1.5 744 0x80000001 0x6498
Area 0.0.0.0 7.16.0.6 10.20.1.5 730 0x80000001 0xb624
Area 0.0.0.0 8.0.0.4 10.20.1.5 729 0x80000002 0x50f1
Area 0.0.0.0 8.0.0.5 10.20.1.5 730 0x80000002 0xc279
Area 0.0.0.1 4.0.0.0 10.20.1.1 740 0x80000002 0xf5a0
Area 0.0.0.1 7.0.0.2 10.20.1.1 745 0x80000001 0xc769
Area 0.0.0.1 8.0.0.4 10.20.1.1 730 0x80000002 0x3f46
Area 0.0.0.1 8.0.0.5 10.20.1.1 731 0x80000002 0x7e02
Area 0.0.0.1 4.0.0.0 10.20.1.2 739 0x80000002 0x5653
Area 0.0.0.1 7.16.0.1 10.20.1.2 744 0x80000001 0x46cc
```

```

Area 0.0.0.1 7.16.0.2 10.20.1.2 735 0x80000001 0x9663
Area 0.0.0.1 7.16.0.3 10.20.1.2 734 0x80000001 0xe6f9
Area 0.0.0.1 7.16.0.4 10.20.1.2 725 0x80000002 0xad3d
Area 0.0.0.1 7.16.0.5 10.20.1.2 725 0x80000002 0x49b8
Area 0.0.0.1 8.0.0.4 10.20.1.2 730 0x80000002 0x3324
Area 0.0.0.1 8.0.0.5 10.20.1.2 731 0x80000002 0x89f3
Area 0.0.0.1 4.0.0.0 10.20.1.3 739 0x80000002 0x6dd6
Area 0.0.0.1 7.16.0.1 10.20.1.3 743 0x80000001 0x6d98
Area 0.0.0.1 7.16.0.3 10.20.1.3 723 0x80000002 0xdef9
Area 0.0.0.1 7.16.0.4 10.20.1.3 729 0x80000001 0xa941
Area 0.0.0.1 7.16.0.5 10.20.1.3 724 0x80000002 0x7084
Area 0.0.0.1 7.16.0.6 10.20.1.3 724 0x80000002 0xcff
Area 0.0.0.1 8.0.0.4 10.20.1.3 730 0x80000002 0xada2
Area 0.0.0.1 8.0.0.5 10.20.1.3 730 0x80000002 0x9bb2

```

```

No. of Opaque LSAs: 40
=====

```

```

A:ALA-A# show router ospf 1 opaque-database

```

```

=====
Rtr Base OSPFv2 Instance 1 Opaque Link State Database (type : All)
=====

```

| Area Id | Type | Link State Id | Adv Rtr Id | Age  | Sequence   | Cksun  |
|---------|------|---------------|------------|------|------------|--------|
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.2  | 205  | 0x8000007e | 0xb1b2 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.5  | 617  | 0x80000084 | 0xb1a6 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.8  | 1635 | 0x80000081 | 0xc391 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.9  | 1306 | 0x80000082 | 0xc58c |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.10 | 53   | 0x80000082 | 0xc986 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.11 | 577  | 0x8000007e | 0xd57c |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.12 | 1628 | 0x80000080 | 0xd578 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.13 | 581  | 0x80000080 | 0xd972 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.22 | 1006 | 0x80000080 | 0xfd3c |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.23 | 1238 | 0x80000083 | 0xfb39 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.27 | 55   | 0x80000083 | 0xc21  |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.28 | 389  | 0x80000083 | 0x101b |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.29 | 1658 | 0x80000082 | 0x1614 |
| 0.0.0.0 | Area | 1.0.0.1       | 180.0.0.30 | 976  | 0x80000083 | 0x180f |
| 0.0.0.0 | Area | 1.0.0.2       | 180.0.0.2  | 45   | 0x800000a0 | 0x2f60 |
| 0.0.0.0 | Area | 1.0.0.2       | 180.0.0.5  | 1357 | 0x80000084 | 0x7038 |
| 0.0.0.0 | Area | 1.0.0.2       | 180.0.0.8  | 1960 | 0x80000084 | 0x3472 |
| ...     |      |               |            |      |            |        |

```

No. of Opaque LSAs: 88
=====

```

```

A:ALA-A#

```

```

*A:Dut-C# show router ospf opaque-database adv-router 10.20.1.5

```

```

=====
Rtr Base OSPFv2 Instance 0 Opaque Link State Database (type : All)
=====

```

| Type | Id      | Link State Id | Adv Rtr Id | Age | Sequence   | Cksun  |
|------|---------|---------------|------------|-----|------------|--------|
| Area | 0.0.0.0 | 4.0.0.0       | 10.20.1.5  | 750 | 0x80000002 | 0x746e |
| Area | 0.0.0.0 | 7.0.0.2       | 10.20.1.5  | 756 | 0x80000001 | 0x6498 |
| Area | 0.0.0.0 | 7.16.0.6      | 10.20.1.5  | 742 | 0x80000001 | 0xb624 |
| Area | 0.0.0.0 | 8.0.0.4       | 10.20.1.5  | 741 | 0x80000002 | 0x50f1 |
| Area | 0.0.0.0 | 8.0.0.5       | 10.20.1.5  | 742 | 0x80000002 | 0xc279 |

```

No. of Opaque LSAs: 5
=====

*A:Dut-C# show router ospf opaque-database adv-router 10.20.1.5 detail
=====
Rtr Base OSPFv2 Instance 0 Opaque Link State Database (type : All) (detail)
=====

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.5
Link State Id : 4.0.0.0 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0x746e
Age : 752 Length : 52
Options : E
Advertisement : Router Info
 Capabilities (1) Len 4 :
 0x14
 SR algorithm (8) Len 1 :
 0x0
 SR label range (9) Len 12 :
 Range-size=1001
 Sub-TLV SID/label(1) len 3 :
 label=25000

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.5
Link State Id : 7.0.0.2 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x6498
Age : 758 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=1 pfxLen=32 AF=0 pfx=10.20.1.5
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=55

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.5
Link State Id : 7.16.0.6 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xb624
Age : 744 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.6
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=66

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.5
Link State Id : 8.0.0.4 LSA Type : Area Opaque

```

```

Sequence No : 0x80000002 Checksum : 0x50f1
Age : 743 Length : 52
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 28 :
 link Type=Transit (2) Id=1.3.5.5 Data=1.3.5.5
 Sub-TLV LAN-Adj-SID (3) len 11 :
 Flags=Value Local (0x60)
 MT-ID=0 Weight=0 Neighbor-ID=10.20.1.3
 SID/Index/Label=262139

```

-----  
Opaque LSA  
-----

```

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.5
Link State Id : 8.0.0.5 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0xc279
Age : 744 Length : 52
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 28 :
 link Type=Transit (2) Id=1.4.5.5 Data=1.4.5.5
 Sub-TLV LAN-Adj-SID (3) len

```

```

=====
*A:Dut-A# show router ospf 1 opaque-database adv-router 10.20.1.1 detail

```

```

=====
Rtr Base OSPFv2 Instance 1 Opaque Link State Database (type : All) (detail)
=====

```

-----  
Opaque LSA  
-----

```

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.1
Link State Id : 1.0.0.1 LSA Type : Area Opaque
Sequence No : 0x800000028 Checksum : 0xb136
Age : 192 Length : 28
Options : E
Advertisement :
 ROUTER-ID TLV (0001) Len 4 : 10.20.1.1

```

-----  
Opaque LSA  
-----

```

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.1
Link State Id : 1.0.0.2 LSA Type : Area Opaque
Sequence No : 0x80000000d Checksum : 0x17f3
Age : 678 Length : 164
Options : E
Advertisement :
 LINK INFO TLV (0002) Len 140 :
 Sub-TLV: 1 Len: 1 LINK_TYPE : 2
 Sub-TLV: 2 Len: 4 LINK_ID : 10.10.1.2
 Sub-TLV: 3 Len: 4 LOC_IP_ADDR : 10.10.1.1
 Sub-TLV: 4 Len: 4 REM_IP_ADDR : 0.0.0.0
 Sub-TLV: 5 Len: 4 TE_METRIC : 1000
 Sub-TLV: 6 Len: 4 MAX_BDWTH : 100000 Kbps
 Sub-TLV: 7 Len: 4 RSRVBL_BDWTH : 800000 Kbps
 Sub-TLV: 8 Len: 32 UNRSRVD_CLS0 :
 P0: 80000 Kbps P1: 320000 Kbps P2: 320000 Kbps P3: 320000 Kbps
 P4: 400000 Kbps P5: 400000 Kbps P6: 400000 Kbps P7: 80000 Kbps
 Sub-TLV: 9 Len: 4 ADMIN_GROUP : 0 None

```

```

Sub-TLV: 17 Len: 36 TELK_BW_CONST:
 BW Model : MAM
 BC0: 80000 Kbps BC1: 0 Kbps BC2: 320000 Kbps BC3: 0 Kbps
 BC4: 0 Kbps BC5: 400000 Kbps BC6: 0 Kbps BC7: 0 Kbps

=====
*A:Dut-A#

*A:Dut-F# show router ospf opaque-database adv-router 10.20.1.6 detail
=====
Rtr Base OSPFv2 Instance 0 Opaque Link State Database (type: All) (detail)
=====

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 4.0.0.0 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0x590e
Age : 288 Length : 52
Options : E
Advertisement : Router Info
 Capabilities (1) Len 4 :
 0x14
 SR algorithm (8) Len 2 :
 0x0 0x2
 SR label range (9) Len 12 :
 Range-size=1000
 Sub-TLV SID/label(1) len 3 :
 label=70000

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.0.0.7 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x899a
Age : 292 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=1 pfxLen=32 AF=0 pfx=10.20.1.6
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=9

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.2 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x6d0b
Age : 292 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.66.6
 Flags=Att Node (0xc0)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=3

```

```

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.9 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xfale
Age : 288 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.33.3
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=0

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.10 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x40d4
Age : 288 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.3
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=6

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.11 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xcb52
Age : 288 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.22.2
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=5

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.12 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x7898
Age : 288 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.1
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)

```



```

MT-ID=0 Algorithm=0 SID/Index/Label=10

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.13 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xff29
Age : 288 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.11.1
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=4

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 8.0.0.2 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0x3098
Age : 289 Length : 48
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 24 :
 link Type=P2P (1) Id=10.20.1.2 Data=1.0.26.6
 Sub-TLV Adj-SID (2) len 7 :
 Flags=Backup Value Local (0xe0)
 MT-ID=0 Weight=0 SID/Index/Label=262143

Opaque LSA

Area Id : 0.0.0.0 Adv Router Id : 10.20.1.6
Link State Id : 8.0.0.4 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0xc9b
Age : 289 Length : 48
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 24 :
 link Type=P2P (1) Id=10.20.1.5 Data=1.0.56.6
 Sub-TLV Adj-SID (2) len 7 :
 Flags=Backup Value Local (0xe0)
 MT-ID=0 Weight=0 SID/Index/Label=262141

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 4.0.0.0 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0x590e
Age : 290 Length : 52
Options : E
Advertisement : Router Info
 Capabilities (1) Len 4 :
 0x14
 SR algorithm (8) Len 2 :
 0x0 0x2
 SR label range (9) Len 12 :
 Range-size=1000

```

```

Sub-TLV SID/label(1) len 3 :
label=70000

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.0.0.6 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xf214
Age : 295 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=1 pfxLen=32 AF=0 pfx=1.0.66.6
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=3

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.1 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x7cc6
Age : 290 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=1 pfxLen=32 AF=0 pfx=1.0.22.2
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP Backup (0x42)
 MT-ID=0 Algorithm=2 SID/Index/Label=996

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.3 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x491
Age : 296 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.6
 Flags=Att Node (0xc0)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=9

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.4 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x9c67
Age : 291 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.55.5
 Flags=Node (0x40)

```

```

Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=2

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.5 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xc253
Age : 291 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.5
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=8

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.6 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xd03e
Age : 291 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=1.0.44.4
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=1

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.7 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0x868f
Age : 291 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.4
 Flags=Node (0x40)
 Sub-TLV Prefix SID (2) len 8 :
 Flags=noPHP (0x40)
 MT-ID=0 Algorithm=0 SID/Index/Label=7

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 7.16.0.8 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xc84a
Age : 291 Length : 44
Options : E
Advertisement : Extended Prefix
 TLV Extended prefix (1) Len 20 :
 rtType=3 pfxLen=32 AF=0 pfx=10.20.1.2

```

```
Flags=Node (0x40)
Sub-TLV Prefix SID (2) len 8 :
Flags=noPHP (0x40)
MT-ID=0 Algorithm=0 SID/Index/Label=11

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 8.0.0.3 LSA Type : Area Opaque
Sequence No : 0x80000002 Checksum : 0xc1fb
Age : 292 Length : 48
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 24 :
 link Type=P2P (1) Id=10.20.1.3 Data=1.0.36.6
 Sub-TLV Adj-SID (2) len 7 :
 Flags=Backup Value Local (0xe0)
 MT-ID=0 Weight=0 SID/Index/Label=262142

Opaque LSA

Area Id : 0.0.0.1 Adv Router Id : 10.20.1.6
Link State Id : 8.0.0.5 LSA Type : Area Opaque
Sequence No : 0x80000001 Checksum : 0xbc0d
Age : 298 Length : 48
Options : E
Advertisement : Extended Link
 TLV Extended link (1) Len 24 :
 link Type=P2P (1) Id=10.20.1.2 Data=1.0.26.6
 Sub-TLV Adj-SID (2) len 7 :
 Flags=Backup Value Local (0xe0)
 MT-ID=0 Weight=0 SID/Index/Label=262140
=====
*A:Dut-F#
```

prefix-sids

|             |                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax      | <b>prefix-sids</b> [ <i>ip-prefix[/prefix-length]</i> ] [ <i>sid sid</i> ] [ <i>adv-router router-id</i> ]                                                                                                                                                                                                                                            |
| Context     | show>router>ospf<br>show>router>ospf3                                                                                                                                                                                                                                                                                                                 |
| Description | This command displays OSPF prefix SIDs.                                                                                                                                                                                                                                                                                                               |
| Parameters  | <i>ip-prefix[/prefix-length]</i> — Displays information about the specified IP prefix and length, up to 64 characters.<br><br><i>sid</i> — Displays information for the specific segment identifier.<br><br><b>Values</b> 0 to 524287<br><br><i>router-id</i> — Displays information for the specific advertising router identified by its router ID. |
| Output      |                                                                                                                                                                                                                                                                                                                                                       |

### Sample Output

```
*A:Dut-F# show router ospf prefix-sids
=====
Rtr Base OSPFv2 Instance 0 Prefix-Sids
=====
```

| Prefix       | Area<br>Adv-Rtr | RtType     | SID<br>Flags |
|--------------|-----------------|------------|--------------|
| 1.0.11.1/32  | 0.0.0.0         | INTER-AREA | 4            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.11.1/32  | 0.0.0.1         | INTRA-AREA | 4            |
|              | 10.20.1.1       |            | NnP          |
| 1.0.11.1/32  | 0.0.0.1         | INTRA-AREA | 999          |
|              | 10.20.1.3       |            | NnPB         |
| 1.0.22.2/32  | 0.0.0.0         | INTER-AREA | 5            |
|              | 10.20.1.2       |            | NnPA         |
| 1.0.22.2/32  | 0.0.0.1         | INTRA-AREA | 5            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.22.2/32  | 0.0.0.1         | INTRA-AREA | 996          |
|              | 10.20.1.6       |            | NnPB         |
| 1.0.33.3/32  | 0.0.0.0         | INTER-AREA | 0            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.33.3/32  | 0.0.0.1         | INTRA-AREA | 0            |
|              | 10.20.1.3       |            | NnP          |
| 1.0.33.3/32  | 0.0.0.1         | INTRA-AREA | 998          |
|              | 10.20.1.1       |            | NnPB         |
| 1.0.44.4/32  | 0.0.0.0         | INTRA-AREA | 1            |
|              | 10.20.1.4       |            | NnP          |
| 1.0.44.4/32  | 0.0.0.0         | INTRA-AREA | 994          |
|              | 10.20.1.5       |            | NnPB         |
| 1.0.44.4/32  | 0.0.0.1         | INTER-AREA | 1            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.55.5/32  | 0.0.0.0         | INTRA-AREA | 2            |
|              | 10.20.1.5       |            | NnP          |
| 1.0.55.5/32  | 0.0.0.0         | INTRA-AREA | 995          |
|              | 10.20.1.4       |            | NnPB         |
| 1.0.55.5/32  | 0.0.0.1         | INTER-AREA | 2            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.66.6/32  | 0.0.0.0         | INTER-AREA | 3            |
|              | 10.20.1.2       |            | NnP          |
| 1.0.66.6/32  | 0.0.0.1         | INTRA-AREA | 3            |
|              | 10.20.1.6       |            | NnP          |
| 1.0.66.6/32  | 0.0.0.1         | INTRA-AREA | 997          |
|              | 10.20.1.2       |            | NnPB         |
| 10.20.1.1/32 | 0.0.0.0         | INTER-AREA | 10           |
|              | 10.20.1.2       |            | NnP          |
| 10.20.1.1/32 | 0.0.0.1         | INTRA-AREA | 10           |
|              | 10.20.1.1       |            | NnP          |
| 10.20.1.2/32 | 0.0.0.0         | INTRA-AREA | 11           |
|              | 10.20.1.2       |            | NnP          |
| 10.20.1.2/32 | 0.0.0.1         | INTER-AREA | 11           |
|              | 10.20.1.2       |            | NnPA         |
| 10.20.1.3/32 | 0.0.0.0         | INTER-AREA | 6            |
|              | 10.20.1.2       |            | NnP          |
| 10.20.1.3/32 | 0.0.0.1         | INTRA-AREA | 6            |
|              | 10.20.1.3       |            | NnP          |
| 10.20.1.4/32 | 0.0.0.0         | INTRA-AREA | 7            |
|              | 10.20.1.4       |            | NnP          |

```

10.20.1.4/32 0.0.0.1 INTER-AREA 7
 10.20.1.2 NnP
10.20.1.5/32 0.0.0.0 INTRA-AREA 8
 10.20.1.5 NnP
10.20.1.5/32 0.0.0.1 INTER-AREA 8
 10.20.1.2 NnP
10.20.1.6/32 0.0.0.0 INTRA-AREA 9
 10.20.1.6 NnP
10.20.1.6/32 0.0.0.1 INTER-AREA 9
 10.20.1.2 NnP

```

-----

No. of Prefix/SIDs: 30

SID Flags : N = Node-SID

nP = no penultimate hop POP

M = Mapping server

E = Explicit-Null

V = Prefix-SID carries a value

L = value/index has local significance

I = Inter Area flag

A = Attached flag

B = Backup flag

=====

\*A:Dut-F#

\*A:Dut-C# show router ospf prefix-sids sid 66

=====

Rtr Base OSPFv2 Instance 0 Prefix-Sids

```

=====
Prefix Area RtType SID
 Adv-Rtr Active Flags

10.20.1.6/32 0.0.0.0 INTER-AREA 66
 10.20.1.4 N NnP
10.20.1.6/32 0.0.0.0 INTER-AREA 66
 10.20.1.5 Y NnP
10.20.1.6/32 0.0.0.1 INTER-AREA 66
 10.20.1.2 N NnP

```

-----

No. of Prefix/SIDs: 3

Flags: N = Node-SID

nP = no penultimate hop POP

M = Mapping server

E = Explicit-Null

V = Prefix-SID carries a value

L = value/index has local significance

I = Inter Area flag

A = Attached flag

=====

\*A:Dut-C#

\*A:Dut-C>config>router>ospf3# show router ospf3 0 prefix-sids

=====

Rtr Base OSPFv3 Instance 0 Prefix-Sids

```

=====
Prefix Area RtType SID
 Adv-Rtr SRMS Flags

3ffe::100:b01/128 0.0.0.1 INTRA-AREA 4

```

```

10.20.1.1 N NnPA
3ffe::100:1602/128 0.0.0.1 INTRA-AREA 5
10.20.1.2 N NnPA
3ffe::100:2103/128 0.0.0.1 INTRA-AREA 0
10.20.1.3 N NnPA
3ffe::100:2c04/128 0.0.0.1 INTER-AREA 1
10.20.1.2 N NnP
3ffe::100:2c04/128 0.0.0.1 INTER-AREA 1
10.20.1.6 N NnP
3ffe::100:3705/128 0.0.0.1 INTER-AREA 2
10.20.1.2 N NnP
3ffe::100:3705/128 0.0.0.1 INTER-AREA 2
10.20.1.6 N NnP
3ffe::100:4206/128 0.0.0.1 INTRA-AREA 3
10.20.1.6 N NnPA
3ffe::a14:101/128 0.0.0.1 INTRA-AREA 10
10.20.1.1 N NnPA
3ffe::a14:102/128 0.0.0.1 INTER-AREA 11
10.20.1.2 N NnPA
3ffe::a14:102/128 0.0.0.1 INTER-AREA 11
10.20.1.6 N NnP
3ffe::a14:103/128 0.0.0.1 INTRA-AREA 6
10.20.1.3 N NnPA
3ffe::a14:104/128 0.0.0.1 INTER-AREA 7
10.20.1.2 N NnP
3ffe::a14:104/128 0.0.0.1 INTER-AREA 7
10.20.1.6 N NnP
3ffe::a14:105/128 0.0.0.1 INTER-AREA 8
10.20.1.2 N NnP
3ffe::a14:105/128 0.0.0.1 INTER-AREA 8
10.20.1.6 N NnP
3ffe::a14:106/128 0.0.0.1 INTER-AREA 9
10.20.1.2 N NnP
3ffe::a14:106/128 0.0.0.1 INTER-AREA 9
10.20.1.6 N NnPA

No. of Prefix/SIDs: 18
SRMS : Y/N = prefix SID advertised by SR Mapping Server (Y) or not (N)
 S = SRMS prefix SID is selected to be programmed
SID Flags : N = Node-SID
 nP = no penultimate hop POP
 M = Mapping server
 E = Explicit-Null
 V = Prefix-SID carries a value
 L = value/index has local significance
 I = Inter Area flag
 A = Attached flag
 B = Backup flag
=====

```

Table 23 describes the OSPF prefix SIDs output fields.

**Table 23** OSPF Prefix SIDs Output Fields

| Label  | Description                        |
|--------|------------------------------------|
| Prefix | Displays the IP prefix for the SID |

**Table 23 OSPF Prefix SIDs Output Fields (Continued)**

| Label              | Description                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| Area               | Displays the OSPF area                                                                                           |
| RtType             | Displays the route type (INTRA-AREA or INTER-AREA)                                                               |
| SID                | Displays the SID value                                                                                           |
| Adv-Rtr            | Displays the IP address of the advertised router                                                                 |
| SRMS               | Displays whether the prefix SID is advertised by the SR mapping server (Y or N) or selected to be programmed (S) |
| Flags              | Displays the SID flags                                                                                           |
| No. of Prefix/SIDs | Displays the number of Prefix/SIDs                                                                               |

## range

**Syntax** `range [area-id]`

**Context** `show>router>ospf`  
`show>router>ospf3`

**Description** This command displays ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression.

**Parameters** `area-id` — Displays the configured ranges for the specified area.

**Values** `ip-address` — a.b.c.d  
`area` — 0 to 4294967295

**Output** OSPF Range Output

The following table describes the OSPF range output fields.

**Table 24 OSPF Range Output Fields**

| Label        | Description                                                                                      |
|--------------|--------------------------------------------------------------------------------------------------|
| Area Id      | A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.    |
| Address/Mask | The mask for the range expressed as a decimal integer mask length or in dotted decimal notation. |



**Table 24 OSPF Range Output Fields (Continued)**

| Label     | Description                                                                                                                                                                                                                                                                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advertise | False — The specified address/mask is not advertised outside the area.<br>True — The specified address/mask is advertised outside the area.                                                                                                                                                                                                        |
| LSDB Type | NSSA — This range was specified in the NSSA context, and specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.<br>Summary — This range was not specified in the NSSA context, the range applies to summary LSAs even if the area is an NSSA. |

### Sample Output

```
A:ALA-A# show router ospf 1 range
=====
Rtr Base OSPFv2 Instance 1 Ranges
=====
Area Id Address/Mask Advertise LSDB Type

No. of Ranges: 0
=====
A:ALA-A#

A:ALA-A# show router ospf range 180.0.7.9
=====
Rtr Base OSPFv2 Instance 0 Ranges for Area-Id 180.0.7.9
=====
Area Id Address/Mask Advertise LSDB Type

No. of Ranges: 0
=====
A:ALA-A#
```

## routes

|                    |                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>routes</b> [ <i>ip-prefix[/prefix-length]</i> ] [ <b>type</b> ] [ <b>detail</b> ] [ <b>alternative</b> ] [ <b>summary</b> ] [ <b>exclude-shortcut</b> ]                                                                                        |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3                                                                                                                                                                                                             |
| <b>Description</b> | This command displays information about OSPF routes.                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>ip-prefix[/prefix-length]</i> — Displays information about the specified IP prefix and length.<br><b>type</b> — Displays information about the specified type.<br><b>Values</b> intra-area, inter-area, external-1, external-2, nssa-1, nssa-2 |

**detail** — Displays detailed information about the routes.

**alternative** — Displays the level of protection per prefix (ref. show router OSPF routes alternative)

**summary** — Displays summarized information about the routes.

**exclude-shortcut** — Displays routes without shortcut, in the **show>router>ospf** context.

**Output** The following output shows an example of OSPF information.

### Sample Output

```
*A:Dut-C# show router ospf routes
=====
Rtr Base OSPFv2 Instance 0 Route Table (alternative)
=====
Destination Type (Dest) Stat SID SIDflgs
 NHIP NHIF Cost [E2]

1.1.1.1/32 IA (HOST) N (R)
 1.1.3.1 3 1000
1.1.2.0/24 IA (NET) N (R)
 1.1.3.1 3 2000
 1.2.3.2 4 2000
1.1.3.0/24 IA (NET) D (F)
 DIRECT 3 1000
1.2.3.0/24 IA (NET) D (F)
 DIRECT 4 1000
1.2.4.0/24 IA (NET) N (R)
 2.2.3.2 5 2000
1.3.5.0/24 IA (NET) D (F)
 DIRECT 6 1000
1.4.5.0/24 IA (NET) N (R)
 1.3.5.5 6 2000
1.4.6.0/24 IE (NET) N (R)
 2.2.3.2 5 3000
 1.3.5.5 6 3000
1.5.6.0/24 IE (NET) N (R)
 1.3.5.5 6 2000
2.2.2.2/32 IA (HOST) N (R)
 2.2.3.2 5 1000
2.2.3.0/24 IA (NET) D (F)
 DIRECT 5 1000
3.3.3.3/32 IA (HOST) D (F)
 DIRECT 2 0
4.4.4.4/32 IA (HOST) N (R)
 2.2.3.2 5 2000
 1.3.5.5 6 2000
5.5.5.5/32 IA (HOST) N (R)
 1.3.5.5 6 1000
6.6.6.6/32 IE (HOST) N (R)
 1.3.5.5 6 2000
10.20.1.1/32 IA (HOST) N (R) 11 NnP
 1.1.3.1 3 1000
10.20.1.2/32 IA (HOST) N (R) 22 NnP
 2.2.3.2 5 1000
```

```

10.20.1.3/32 IA (HOST) D (F) 33 NnP
 DIRECT 1 0
10.20.1.4/32 IA (HOST) N (R) 44 NnP
 2.2.3.2 5 2000
 1.3.5.5 6 2000
10.20.1.5/32 IA (HOST) N (R) 55 NnP
 1.3.5.5 6 1000
10.20.1.6/32 IE (HOST) N (R) 66 NnP
 1.3.5.5 6 2000
10.20.1.1/0 IA (RTR) N (N)
 1.1.3.1 3 1000
10.20.1.2/0 IA (AB-AS) N (N)
 2.2.3.2 5 1000
10.20.1.2/0 IA (AB-AS) N (N)
 1.2.3.2 4 1000
10.20.1.4/0 IA (AB-AS) N (N)
 2.2.3.2 5 2000
 1.3.5.5 6 2000
10.20.1.5/0 IA (AB-AS) N (N)
 1.3.5.5 6 1000

```

```

No. of routes found: 26 (31 paths)
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
 (N) = not added (D) = policy discarded
SID Flags : N = Node-SID
 nP = no penultimate hop POP
 M = Mapping server
 E = Explicit-Null
 V = Prefix-SID carries a value
 L = value/index has local significance
 I = Inter Area flag
 A = Attached flag
=====

```

\*A:Dut-C# show router ospf routes

Rtr Base OSPFv2 Instance 0 Routing Table

```

=====
Destination Type(Dest) Stat SID SIDflgs
 NHIP NHIF Cost[E2]

1.1.1.1/32 IA (HOST) N (R)
 1.1.3.1 3 1000
1.1.2.0/24 IA (NET) N (R)
 1.1.3.1 3 2000
 1.2.3.2 4 2000
1.1.3.0/24 IA (NET) D (F)
 DIRECT 3 1000
1.2.3.0/24 IA (NET) D (F)
 DIRECT 4 1000
1.2.4.0/24 IA (NET) N (R)
 2.2.3.2 5 2000
1.3.5.0/24 IA (NET) D (F)
 DIRECT 6 1000
1.4.5.0/24 IA (NET) N (R)
 1.3.5.5 6 2000
1.4.6.0/24 IE (NET) N (R)
 2.2.3.2 5 3000

```

```

1.3.5.5 6 3000
1.5.6.0/24 IE (NET) N (R)
1.3.5.5 6 2000
2.2.2.2/32 IA (HOST) N (R)
2.2.3.2 5 1000
2.2.3.0/24 IA (NET) D (F)
DIRECT 5 1000
3.3.3.3/32 IA (HOST) D (F)
DIRECT 2 0
4.4.4.4/32 IA (HOST) N (R)
2.2.3.2 5 2000
1.3.5.5 6 2000
5.5.5.5/32 IA (HOST) N (R)
1.3.5.5 6 1000
6.6.6.6/32 IE (HOST) N (R)
1.3.5.5 6 2000
10.20.1.1/32 IA (HOST) N (R) 11 NnP
1.1.3.1 3 1000
10.20.1.2/32 IA (HOST) N (R) 22 NnP
2.2.3.2 5 1000
10.20.1.3/32 IA (HOST) D (F) 33 NnP
DIRECT 1 0
10.20.1.4/32 IA (HOST) N (R) 44 NnP
2.2.3.2 5 2000
1.3.5.5 6 2000
10.20.1.5/32 IA (HOST) N (R) 55 NnP
1.3.5.5 6 1000
10.20.1.6/32 IE (HOST) N (R) 66 NnP
1.3.5.5 6 2000
10.20.1.1/0 IA (RTR) N (N)
1.1.3.1 3 1000
10.20.1.2/0 IA (AB-AS) N (N)
2.2.3.2 5 1000
10.20.1.2/0 IA (AB-AS) N (N)
1.2.3.2 4 1000
10.20.1.4/0 IA (AB-AS) N (N)
2.2.3.2 5 2000
1.3.5.5 6 2000
10.20.1.5/0 IA (AB-AS) N (N)
1.3.5.5 6 1000

```

-----  
No. of routes found: 26 (31 paths)

Stat: D = direct    N = not direct

(RTM stat):(R) = added            (F) = add failed

(N) = not added    (D) = policy discarded

SID Flags : N = Node-SID

nP = no penultimate hop POP

M = Mapping server

E = Explicit-Null

V = Prefix-SID carries a value

L = value/index has local significance

I = Inter Area flag

A = Attached flag

=====

\*A:Dut-C# show router ospf routes alternative

=====

Rtr Base OSPFv2 Instance 0 Route Table (alternative)

=====

| Destination<br>NHIF<br>A-NHIF (L) | Type (Dest)<br>NHIF<br>A-NHIF | Stat<br>Cost [E2]<br>A-Cost [E2] | SID | SIDflgs |
|-----------------------------------|-------------------------------|----------------------------------|-----|---------|
| 1.1.1.1/32                        | IA (HOST)                     | N (R)                            |     |         |
| 1.1.3.1                           | 3                             | 1000                             |     |         |
| 1.2.3.2 (L)                       | 4                             | 2000                             |     |         |
| 1.1.2.0/24                        | IA (NET)                      | N (R)                            |     |         |
| 1.1.3.1                           | 3                             | 2000                             |     |         |
| 1.2.3.2                           | 4                             | 2000                             |     |         |
| 1.1.3.0/24                        | IA (NET)                      | D (F)                            |     |         |
| DIRECT                            | 3                             | 1000                             |     |         |
| 1.2.3.0/24                        | IA (NET)                      | D (F)                            |     |         |
| DIRECT                            | 4                             | 1000                             |     |         |
| 1.2.4.0/24                        | IA (NET)                      | N (R)                            |     |         |
| 2.2.3.2                           | 5                             | 2000                             |     |         |
| 1.3.5.0/24                        | IA (NET)                      | D (F)                            |     |         |
| DIRECT                            | 6                             | 1000                             |     |         |
| 1.4.5.0/24                        | IA (NET)                      | N (R)                            |     |         |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 1.4.6.0/24                        | IE (NET)                      | N (R)                            |     |         |
| 2.2.3.2                           | 5                             | 3000                             |     |         |
| 1.3.5.5                           | 6                             | 3000                             |     |         |
| 1.5.6.0/24                        | IE (NET)                      | N (R)                            |     |         |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 2.2.2.2/32                        | IA (HOST)                     | N (R)                            |     |         |
| 2.2.3.2                           | 5                             | 1000                             |     |         |
| 2.2.3.0/24                        | IA (NET)                      | D (F)                            |     |         |
| DIRECT                            | 5                             | 1000                             |     |         |
| 3.3.3.3/32                        | IA (HOST)                     | D (F)                            |     |         |
| DIRECT                            | 2                             | 0                                |     |         |
| 4.4.4.4/32                        | IA (HOST)                     | N (R)                            |     |         |
| 2.2.3.2                           | 5                             | 2000                             |     |         |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 5.5.5.5/32                        | IA (HOST)                     | N (R)                            |     |         |
| 1.3.5.5                           | 6                             | 1000                             |     |         |
| 6.6.6.6/32                        | IE (HOST)                     | N (R)                            |     |         |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 10.20.1.1/32                      | IA (HOST)                     | N (R)                            | 11  | NnP     |
| 1.1.3.1                           | 3                             | 1000                             |     |         |
| 1.2.3.2 (L)                       | 4                             | 2000                             |     |         |
| 10.20.1.2/32                      | IA (HOST)                     | N (R)                            | 22  | NnP     |
| 2.2.3.2                           | 5                             | 1000                             |     |         |
| 10.20.1.3/32                      | IA (HOST)                     | D (F)                            | 33  | NnP     |
| DIRECT                            | 1                             | 0                                |     |         |
| 10.20.1.4/32                      | IA (HOST)                     | N (R)                            | 44  | NnP     |
| 2.2.3.2                           | 5                             | 2000                             |     |         |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 10.20.1.5/32                      | IA (HOST)                     | N (R)                            | 55  | NnP     |
| 1.3.5.5                           | 6                             | 1000                             |     |         |
| 10.20.1.6/32                      | IE (HOST)                     | N (R)                            | 66  | NnP     |
| 1.3.5.5                           | 6                             | 2000                             |     |         |
| 10.20.1.1/0                       | IA (RTR)                      | N (N)                            |     |         |
| 1.1.3.1                           | 3                             | 1000                             |     |         |
| 10.20.1.2/0                       | IA (AB-AS)                    | N (N)                            |     |         |
| 2.2.3.2                           | 5                             | 1000                             |     |         |
| 10.20.1.2/0                       | IA (AB-AS)                    | N (N)                            |     |         |
| 1.2.3.2                           | 4                             | 1000                             |     |         |
| 10.20.1.4/0                       | IA (AB-AS)                    | N (N)                            |     |         |

```

 2.2.3.2 5 2000
 1.3.5.5 6 2000
10.20.1.5/0 IA (AB-AS) N (N)
 1.3.5.5 6 1000

```

```

No. of routes found: 26 (31 paths)
Flags: L = Loop-Free Alternate nexthop
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
 (N) = not added (D) = policy discarded
SID Flags : N = Node-SID
 nP = no penultimate hop POP
 M = Mapping server
 E = Explicit-Null
 V = Prefix-SID carries a value
 L = value/index has local significance
 I = Inter Area flag
 A = Attached flag

```

```

=====
*A:Dut-C# show router ospf routes alternative detail
=====

```

```

Rtr Base OSPFv2 Instance 0 Route Table (alternative) (detail)
=====

```

| Destination | Type(Dest) | Stat        | SID     | SIDflgs | Type     | Weight:Cfg/Norm |
|-------------|------------|-------------|---------|---------|----------|-----------------|
| NHIF        | NHIF       | Cost [E2]   | Area    |         |          |                 |
| A-NHIF(L)   | A-NHIF     | A-Cost [E2] | A-Type  |         | PGID     |                 |
| 1.1.1.1/32  | IA (HOST)  | N (R)       |         |         |          |                 |
| 1.1.3.1     | 3          | 1000        | 0.0.0.1 |         |          |                 |
| 1.2.3.2 (L) | 4          | 2000        | LINK    |         | 0x410079 |                 |
| 1.1.2.0/24  | IA (NET)   | N (R)       |         |         |          |                 |
| 1.1.3.1     | 3          | 2000        | 0.0.0.1 |         |          |                 |
| 1.2.3.2     | 4          | 2000        | 0.0.0.1 |         |          |                 |
| 1.1.3.0/24  | IA (NET)   | D (F)       |         |         |          |                 |
| DIRECT      | 3          | 1000        | 0.0.0.1 |         |          |                 |
| 1.2.3.0/24  | IA (NET)   | D (F)       |         |         |          |                 |
| DIRECT      | 4          | 1000        | 0.0.0.1 |         |          |                 |
| 1.2.4.0/24  | IA (NET)   | N (R)       |         |         |          |                 |
| 2.2.3.2     | 5          |             |         |         |          |                 |

```

=====
*A:Dut-C# show router ospf 1 routes exclude-shortcut alternative detail
=====

```

```

Rtr Base OSPFv2 Instance 1 Route Table (excl-shortcut) (alternative) (detail)
=====

```

| Destination | Type(Dest) | Stat        | Area    | Tunnel-Information |
|-------------|------------|-------------|---------|--------------------|
| NHIF        | NHIF       | Cost [E2]   | Area    |                    |
| A-NHIF(L)   | A-NHIF     | A-Cost [E2] | A-Type  | PGID               |
| 1.1.2.0/24  | IA (NET)   | N (R)       |         |                    |
| 1.1.3.1     | 3          | 20          | 0.0.0.0 |                    |
| 1.2.3.2     | 4          | 20          | 0.0.0.0 |                    |
| 1.1.3.0/24  | IA (NET)   | D (F)       |         |                    |
| DIRECT      | 3          | 10          | 0.0.0.0 |                    |
| 1.2.3.0/24  | IA (NET)   | D (F)       |         |                    |
| DIRECT      | 4          | 10          | 0.0.0.0 |                    |
| 1.2.4.0/24  | IA (NET)   | N (R)       |         |                    |
| 1.2.3.2     | 4          | 20          | 0.0.0.0 |                    |

```

1.3.5.0/24 IA (NET) D (F)
 DIRECT 5 10 0.0.0.0
1.4.5.0/24 IA (NET) N (R)
 1.3.5.5 5 20 0.0.0.0
1.4.6.0/24 IA (NET) N (R)
 1.2.3.2 4 30 0.0.0.0
 1.3.5.5 5 30 0.0.0.0
10.20.1.1/32 IA (HOST) N (R)
 1.1.3.1 3 10 0.0.0.0
10.20.1.2/32 IA (HOST) N (R)
 1.2.3.2 4 10 0.0.0.0
10.20.1.3/32 IA (HOST) D (F)
 DIRECT 1 0 0.0.0.0
10.20.1.4/32 IA (HOST) N (R)
 1.2.3.2 4 20 0.0.0.0
 1.3.5.5 5 20 0.0.0.0
10.20.1.5/32 IA (HOST) N (R)
 1.3.5.5 5 10 0.0.0.0
10.20.1.6/32 IA (HOST) N (R)
 1.2.3.2 4 30 0.0.0.0
 1.3.5.5 5 30 0.0.0.0
10.20.1.1/0 IA (RTR) N (N)
 1.1.3.1 3 10 0.0.0.0
10.20.1.2/0 IA (RTR) N (N)
 1.2.3.2 4 10 0.0.0.0
10.20.1.4/0 IA (RTR) N (N)
 1.2.3.2 4 20 0.0.0.0
 1.3.5.5 5 20 0.0.0.0
10.20.1.5/0 IA (RTR) N (N)
 1.3.5.5 5 10 0.0.0.0
10.20.1.6/0 IA (RTR) N (N)
 1.2.3.2 4 30 0.0.0.0
 1.3.5.5 5 30 0.0.0.0

```

```

No. of routes found: 18 (24 paths)
Flags: L = Loop-Free Alternate nexthop
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
 (N) = not added (D) = policy discarded
=====

```

\*A:Dut-C#

\*A:Dut-A# show router ospf routes alternative detail

```

=====
Rtr Base OSPFv2 Instance 0 Route Table (alternative) (detail)
=====
Destination Type(Dest) Stat
 NHIP NHIF Cost [E2] Area Tunnel-Information
 A-NHIP(L) A-NHIF A-Cost [E2] A-Type PGID

1.1.2.0/24 IA (NET) D (F)
 DIRECT 2 10 0.0.0.0
1.1.3.0/24 IA (NET) D (F)
 DIRECT 3 10 0.0.0.0
1.2.3.0/24 IA (NET) N (R)
 1.1.2.2 2 20 0.0.0.0
 1.1.3.3 3 20 0.0.0.0
1.2.4.0/24 IA (NET) N (R)
 1.1.2.2 2 20 0.0.0.0

```

```

 1.1.3.3(L) 3 30 LINK 0x130015
1.3.5.0/24 IA (NET) N (R)
 1.1.3.3 3 20 0.0.0.0
 1.1.2.2(L) 2 30 LINK 0x130016
1.4.5.0/24 IA (NET) N (R)
 1.1.2.2 2 30 0.0.0.0
 1.1.3.3 3 30 0.0.0.0
1.4.6.0/24 IA (NET) N (R)
 1.1.2.2 2 30 0.0.0.0
 1.1.3.3(L) 3 40 LINK 0x130015
1.5.6.0/24 IA (NET) N (R)
 1.1.3.3 3 30 0.0.0.0
 1.1.2.2(L) 2 40 LINK 0x130016
10.20.1.1/32 IA (HOST) D (F)
 DIRECT 1 0 0.0.0.0
10.20.1.2/32 IA (HOST) N (R)
 1.1.2.2 2 10 0.0.0.0
 1.1.3.3(L) 3 20 LINK 0x130015
10.20.1.3/32 IA (HOST) N (R)
 1.1.3.3 3 10 0.0.0.0
 1.1.2.2(L) 2 20 LINK 0x130016
10.20.1.4/32 IA (HOST) N (R)
 1.1.2.2 2 20 0.0.0.0
 1.1.3.3(L) 3 30 LINK 0x130015
10.20.1.5/32 IA (HOST) N (R)
 1.1.3.3 3 20 0.0.0.0
 1.1.2.2(L) 2 30 LINK 0x130016
10.20.1.3/0 IA (RTR) N (N)
 1.1.3.3 3 10 0.0.0.0
10.20.1.4/0 IA (RTR) N (N)
 1.1.2.2 2 20 0.0.0.0
10.20.1.5/0 IA (RTR) N (N)
 1.1.3.3 3 20 0.0.0.0
10.20.1.6/0 IA (RTR) N (N)
 1.1.3.3 3 30 0.0.0.0
 1.1.2.2 2 30 0.0.0.0

```

-----

19 OSPFv2 routes found (23 paths)

Flags: L = Loop-Free Alternate nexthop

=====

\*A:Dut-C# show router ospf 1 routes 10.0.0.2/32 detail

=====

Rtr Base OSPFv2 Instance 1 Routing Table (detail)

=====

| Destination<br>NHIP | Type(Dest)<br>NHIF | Stat<br>Cost [E2] | Area    | Type     | Weight:Cfg/Norm |
|---------------------|--------------------|-------------------|---------|----------|-----------------|
| 10.0.0.2/32         | E2 (HOST)          | N (R)             |         |          |                 |
| 1.0.0.3             | RSVP:94            | 9:10              | 0.0.0.0 | Shortcut | 40/20           |
| 1.0.0.3             | RSVP:61442         | 9:10              | 0.0.0.0 | Shortcut | 2/ 1            |

-----

No. of routes found: 1 (2 paths)

Stat: D = direct N = not direct

(RTM stat):(R) = added (F) = add failed

(N) = not added (D) = policy discarded

=====

\*A:Dut-C>config>router>ospf3# show router ospf3 0 routes



| Rtr Base OSPFv3 Instance 0 Route Table |           |             |           |
|----------------------------------------|-----------|-------------|-----------|
| =====                                  |           |             |           |
| Destination                            |           | Type (Dest) | Stat      |
| SID                                    | SID flags |             |           |
| NHIP                                   |           | NHIF        | Cost [E2] |
| -----                                  |           |             |           |
| 3ffe::100:900/120                      |           | IA (STUB)   | D (F)     |
| DIRECT                                 |           | 4           | 200       |
| 3ffe::100:b01/128                      |           | IA (HOST)   | N (R)     |
| 4                                      | NnPA      |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 100       |
| 3ffe::100:c00/120                      |           | IA (STUB)   | N (R)     |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 200       |
| 3ffe::100:d00/120                      |           | IA (STUB)   | D (F)     |
| DIRECT                                 |           | 2           | 100       |
| 3ffe::100:1602/128                     |           | IA (HOST)   | N (R)     |
| 5                                      | NnPA      |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 200       |
| 3ffe::100:1800/120                     |           | IE (NET)    | N (R)     |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 300       |
| 3ffe::100:1a00/120                     |           | IE (NET)    | N (R)     |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 200       |
| 3ffe::100:2103/128                     |           | IA (HOST)   | D (F)     |
| 0                                      | NnPA      |             |           |
| DIRECT                                 |           | 5           | 0         |
| 3ffe::100:2400/120                     |           | IA (STUB)   | D (F)     |
| DIRECT                                 |           | 3           | 100       |
| 3ffe::100:2c04/128                     |           | IE (HOST)   | N (R)     |
| 1                                      | NnP       |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 300       |
| 3ffe::100:2d00/120                     |           | IE (NET)    | N (R)     |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 300       |
| 3ffe::100:3705/128                     |           | IE (HOST)   | N (R)     |
| 2                                      | NnP       |             |           |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 200       |
| 3ffe::100:3800/120                     |           | IE (NET)    | N (R)     |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 200       |
| 3ffe::100:4206/128                     |           | IA (HOST)   | N (R)     |
| 3                                      | NnPA      |             |           |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 100       |
| 3ffe::a14:101/128                      |           | IA (HOST)   | N (R)     |
| 10                                     | NnPA      |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 100       |
| 3ffe::a14:102/128                      |           | IE (HOST)   | N (R)     |
| 11                                     | NnPA      |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 200       |
| 3ffe::a14:103/128                      |           | IA (HOST)   | D (F)     |
| 6                                      | NnPA      |             |           |
| DIRECT                                 |           | 1           | 0         |
| 3ffe::a14:104/128                      |           | IE (HOST)   | N (R)     |
| 7                                      | NnP       |             |           |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" |           | 2           | 300       |
| 3ffe::a14:105/128                      |           | IE (HOST)   | N (R)     |
| 8                                      | NnP       |             |           |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" |           | 3           | 200       |
| 3ffe::a14:106/128                      |           | IE (HOST)   | N (R)     |
| 9                                      | NnPA      |             |           |
| fe80::200:ff:fe00:6-"DUTC TO DUTF.1.0" |           | 3           | 100       |

|                                        |            |       |
|----------------------------------------|------------|-------|
| 3ffe::3300:b01/128                     | IA (HOST)  | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 100   |
| 3ffe::3300:c00/120                     | IA (STUB)  | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 200   |
| 3ffe::3300:d00/120                     | IA (STUB)  | D (F) |
| DIRECT                                 | 2          | 100   |
| 3ffe::3300:1602/128                    | IA (HOST)  | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 200   |
| 3ffe::3300:1800/120                    | IE (NET)   | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 300   |
| 3ffe::3300:1a00/120                    | IE (NET)   | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 200   |
| 3ffe::3300:2103/128                    | IA (HOST)  | D (F) |
| DIRECT                                 | 5          | 0     |
| 3ffe::3300:2400/120                    | IA (STUB)  | D (F) |
| DIRECT                                 | 3          | 100   |
| 3ffe::3300:2c04/128                    | IE (HOST)  | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 300   |
| 3ffe::3300:2d00/120                    | IE (NET)   | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 300   |
| 3ffe::3300:3705/128                    | IE (HOST)  | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 200   |
| 3ffe::3300:3800/120                    | IE (NET)   | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 200   |
| 3ffe::3300:4206/128                    | IA (HOST)  | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 100   |
| 3ffe::6800:1/128                       | E2 (HOST)  | N (R) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 300:1 |
| 3ffe::6900:1/128                       | E2 (HOST)  | N (R) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 200:1 |
| 10.20.1.1/0                            | IA (RTR)   | N (N) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 100   |
| 10.20.1.2/0                            | IA (AB-AS) | N (N) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 200   |
| 10.20.1.4/0                            | IE (ASBR)  | N (N) |
| fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0" | 2          | 300   |
| 10.20.1.5/0                            | IE (ASBR)  | N (N) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 200   |
| 10.20.1.6/0                            | IA (AB-AS) | N (N) |
| fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0" | 3          | 100   |

-----

No. of routes found: 40 (40 paths)

Stat: D = direct N = not direct

(RTM stat): (R) = added (F) = add failed

(N) = not added (D) = policy discarded

SID Flags : N = Node-SID

nP = no penultimate hop POP

M = Mapping server

E = Explicit-Null

V = Prefix-SID carries a value

L = value/index has local significance

I = Inter Area flag

A = Attached flag

B = Backup flag

=====

\*A:Dut-C# show router ospf3 31 routes detail

=====

Rtr Base OSPFv3 Instance 31 Route Table (detail)

```

=====
Destination Type (Dest) Stat
 NHIP NHIF Cost [E2] Area

3ffe::10:10:2:0/120 IA (NET) D (F)
 DIRECT 2 1000 0.0.0.0
3ffe::10:10:3:0/120 IA (NET) D (F)
 DIRECT 3 1000 0.0.0.0
3ffe::10:10:4:0/120 IA (NET) N (R)
 fe80::2052:1ff:fe01:3-"ip-10.10.3.3" 3 2000 0.0.0.0
3ffe::10:10:5:0/120 IA (NET) D (F)
 DIRECT 4 1000 0.0.0.0
3ffe::10:10:6:0/120 IA (NET) N (R)
 fe80::205e:1ff:fe01:1-"ip-10.10.5.3" 4 2000 0.0.0.0
3ffe::10:10:9:0/120 IA (NET) N (R)
 10.20.1.4 (SC) RSVP:1 3000 0.0.0.0
3ffe::10:10:13:0/120 IA (STUB) N (R)
 fe80::6628:1ff:fe01:2-"ip-10.10.2.3" 2 2000 0.0.0.0
3ffe::10:10:14:0/120 IA (STUB) N (R)
 10.20.1.4 (SC) RSVP:1 4000 0.0.0.0
3ffe::10:20:1:1/128 IA (HOST) N (R)
 fe80::6628:1ff:fe01:2-"ip-10.10.2.3" 2 1000 0.0.0.0
3ffe::10:20:1:2/128 IA (HOST) N (R)
 fe80::2052:1ff:fe01:3-"ip-10.10.3.3" 3 1000 0.0.0.0
3ffe::10:20:1:3/128 IA (HOST) D (F)
 DIRECT 1 0 0.0.0.0
3ffe::10:20:1:4/128 IA (HOST) N (R)
 10.20.1.4 (SC) RSVP:1 2000 0.0.0.0
3ffe::10:20:1:5/128 IA (HOST) N (R)
 fe80::205e:1ff:fe01:1-"ip-10.10.5.3" 4 1000 0.0.0.0
3ffe::10:20:1:6/128 IA (HOST) N (R)
 10.20.1.4 (SC) RSVP:1 3000 0.0.0.0
3ffe::100:100:100:4/128 IA (HOST) N (R)
 10.20.1.4 (SC) RSVP:1 2000 0.0.0.0
3ffe::100:100:100:6/128 IA (HOST) N (R)
 10.20.1.4 (SC) RSVP:1 3000 0.0.0.0
10.20.1.1/0 IA (RTR) N (N)
 fe80::6628:1ff:fe01:2-"ip-10.10.2.3" 2 1000 0.0.0.0
10.20.1.2/0 IA (RTR) N (N)
 fe80::2052:1ff:fe01:3-"ip-10.10.3.3" 3 1000 0.0.0.0
10.20.1.4/0 IA (RTR) N (N)
 10.20.1.4 (SC) RSVP:1 2000 0.0.0.0
10.20.1.5/0 IA (RTR) N (N)
 fe80::205e:1ff:fe01:1-"ip-10.10.5.3" 4 1000 0.0.0.0
10.20.1.6/0 IA (RTR) N (N)
 10.20.1.4 (SC) RSVP:1 3000 0.0.0.0

No. of routes found: 21 (21 paths)
Stat: D = direct N = not direct
(RTM stat):(R) = added (F) = add failed
 (N) = not added (D) = policy discarded
=====
*A:Dut-C#
*A:Dut-C# show router ospf3 20 status
=====
Rtr Base OSPFv3 Instance 20 Status
=====
OSPF Cfg Router Id : 10.20.1.3
OSPF Oper Router Id : 10.20.1.3

```

---

```

OSPF Version : 3
OSPF Admin Status : Enabled
OSPF Oper Status : Enabled
Graceful Restart : Disabled
GR Helper Mode : Disabled
GR Strict LSA Checking : Enabled (operational down)
Preference : 10
External Preference : 150
Backbone Router : True
Area Border Router : False
AS Border Router : False
Traffic Engineering Support : False
Demand Exts Support : False
In Overload State : False
In External Overflow State : False
Exit Overflow Interval : 0
Last Overflow Entered : Never
Last Overflow Exit : Never
External LSA Limit : -1
Reference Bandwidth : 100,000,000 Kbps
Init SPF Delay : 1000 msec
Sec SPF Delay : 1000 msec
Max SPF Delay : 10000 msec
Min LS Arrival Interval : 1000 msec
Init LSA Gen Delay : 5000 msec
Sec LSA Gen Delay : 5000 msec
Max LSA Gen Delay : 5000 msec
Lsa accumulate : 1000 msec
Redistribute delay : 1000 msec
Incremental SPF wait : 1000 msec
Last Ext SPF Run : Never
Ext LSA Cksum Sum : 0x0
OSPF Last Enabled : 10/10/2018 20:22:05
Unicast Import : True
Multicast Import : False
Export Policies : None
Import Policies : None
Lfa Policies : None
OSPF Ldp Sync Admin Status : Enabled
LDP-over-RSVP : Disabled
IGP-Shortcut : Enabled
IPv6 IGP SC Tunn-Nhop : Res-Filter(RSVP)
Advertise-Tunnel-Link : Disabled
LFA : Disabled
Remote-LFA : Disabled
Max PQ Cost : 65535
TI-LFA : Disabled
Max SR FRR Labels : 2
Export Limit : 0
Export Limit Log Percent : 0
Total Exp Routes : 0
RIB-priority-high prefix list : None
Segment Routing : Disabled
Extended LSA : sparse
=====

```

## sham-link

|                    |                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sham-link</b> [ <i>interface-name</i> ] [ <b>detail</b> ]<br><b>sham-link</b> <i>interface-name</i> <b>remote</b> <i>ip-address</i> [ <b>detail</b> ]                                                                                                                                                                      |
| <b>Context</b>     | show>router>ospf                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command displays OSPFv2 sham-link information.                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>interface-name</i> — Displays only the sham-link identified by this interface name, up to a maximum of 32 characters.<br><b>remote</b> — Displays sham-link information for the sham-link identified by the specified remote IP address.<br><i>ip-address</i> — Displays only the sham-link identified by this IP address. |

## sham-link-neighbor

|                    |                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sham-link-neighbor</b> [ <b>detail</b> ]<br><b>sham-link-neighbor</b> <i>interface-name</i> <b>remote</b> <i>ip-address</i> [ <b>detail</b> ]                                                                                                                                                                                                        |
| <b>Context</b>     | show>router>ospf                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command displays OSPFv2 sham-link neighbor information.                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>interface-name</i> — Displays only the sham-link neighbor identified by this interface name, up to a maximum of 32 characters.<br><b>remote</b> — Displays sham-link neighbor information for the neighbor identified by the specified remote IP address.<br><i>ip-address</i> — Displays only the sham-link neighbor identified by this IP address. |

## spf

|                    |                                                                             |
|--------------------|-----------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>spf</b> [ <i>lfa</i> ]                                                   |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3                                       |
| <b>Description</b> | This command displays statistics of shortest-path-first (SPF) calculations. |
| <b>Parameters</b>  | <b>lfa</b> — Displays Loop-Free Alternate (LFA) next-hop information.       |
| <b>Output</b>      | SPF Output Fields<br><br>The following table describes SPF output fields.   |

**Table 25**      **SPF Output Fields**

| Label                     | Description                                                                                                                                                                                                                                                          |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total SPF Runs            | The total number of SPF runs triggered by new or updated LSAs.                                                                                                                                                                                                       |
| Last Full SPF run @       | The date and time when the last full SPF was run.                                                                                                                                                                                                                    |
| Last Full SPF Time        | The duration, in seconds, of the last full SPF calculation. Last Full SPF Time = Intra SPF Time + Inter SPF Time + Extern SPF Time + RTM Updt Time.                                                                                                                  |
| Intra SPF Time            | The duration, in seconds, of the Intra-area phase of the last SPF calculation.                                                                                                                                                                                       |
| Inter SPF Time            | The duration, in seconds, of the Inter-area phase of the last SPF calculation.                                                                                                                                                                                       |
| Extern SPF Time           | The duration, in seconds, of the External phase of the last SPF calculation.                                                                                                                                                                                         |
| RTM Updt Time             | The duration, in seconds, of the RTM update phase of the last SPF calculation.                                                                                                                                                                                       |
| Min/Avg/Max Full SPF Time | Min<br>The minimum time, in seconds, used to perform a full SPF calculation.<br>Avg<br>The average time, in seconds, of all the full SPF calculations performed by this OSPF router.<br>Max<br>The maximum time, in seconds, used to perform a full SPF calculation. |
| Total Sum Incr SPF Runs   | The total number of incremental SPF runs triggered by new or updated type-3 and type-4 summary LSAs.                                                                                                                                                                 |
| Total Ext Incr SPF Runs   | The total number of incremental SPF runs triggered by new or updated type-5 external LSAs.                                                                                                                                                                           |

**Sample Output**

```

*A:Dut-C# show router ospf spf
=====
Rtr Base OSPFv2 Instance 0 SPF Statistics
=====
Total SPF Runs : 6
Last Full SPF run @ : 05/27/2015 08:45:25
Last Full SPF Time : < 0.01 secs
 Intra SPF Time : < 0.01 secs
 Inter SPF Time : < 0.01 secs
 Extern SPF Time : < 0.01 secs

```

```

RTM Updt Time : < 0.01 secs

Min/Avg/Max Full SPF Times : 0.00/0.00/0.00 secs
Min/Avg/Max RTM Updt Times : 0.00/0.00/0.00 secs

Total Sum Incr SPF Runs : 9
Last Sum Incr SPF run @ : 05/27/2015 08:36:14
Last Sum Incr Calc Time : < 0.01 secs

Total Ext Incr SPF Runs : 0
=====

*A:Dut-C# show router ospf spf lfa
=====
Rtr Base OSPFv2 Instance 0 SPF Statistics
=====
Total SPF Runs : 6
Last Full SPF run @ : 05/27/2015 08:45:25
Last Full SPF Time : < 0.01 secs
 Intra SPF Time : < 0.01 secs
 Inter SPF Time : < 0.01 secs
 Extern SPF Time : < 0.01 secs
 RTM Updt Time : < 0.01 secs

Min/Avg/Max Full SPF Times : 0.00/0.00/0.00 secs
Min/Avg/Max RTM Updt Times : 0.00/0.00/0.00 secs

Total Sum Incr SPF Runs : 9
Last Sum Incr SPF run @ : 05/27/2015 08:36:14
Last Sum Incr Calc Time : < 0.01 secs

Total Ext Incr SPF Runs : 0

Total LFA SPF Runs : 1
Last LFA SPF run @ : 05/27/2015 08:45:25
Last LFA SPF Time : < 0.01 secs
Min/Avg/Max LFA SPF Times : 0.00/0.00/0.00 secs
=====

A:Dut-A# show router ospf 1 spf lfa
=====
Rtr Base OSPFv2 Instance 1 SPF Statistics
=====
Total SPF Runs : 6
Last Full SPF run @ : 02/20/2012 09:19:35
Last Full SPF Time : < 0.01 secs
 Intra SPF Time : < 0.01 secs
 Inter SPF Time : < 0.01 secs
 Extern SPF Time : < 0.01 secs
 RTM Updt Time : < 0.01 secs

Min/Avg/Max Full SPF Times : 0.00/0.00/0.00 secs
Min/Avg/Max RTM Updt Times : 0.00/0.00/0.00 secs

Total Sum Incr SPF Runs : 0

Total Ext Incr SPF Runs : 0

Total LFA SPF Runs : 5

```

```

Last LFA SPF run @ : 02/20/2012 09:19:35
Last LFA SPF Time : < 0.01 secs
Min/Avg/Max LFA SPF Times : 0.00/0.00/0.00 secs
=====

A:ALA-A# show router ospf 1 spf
=====
Rtr Base OSPFv2 Instance 1 SPF Statistics
=====
Total SPF Runs : 109
Last Full SPF run @ : 11/07/2006 18:43:07
Last Full SPF Time : < 0.01 secs
 Intra SPF Time : < 0.01 secs
 Inter SPF Time : < 0.01 secs
 Extern SPF Time : < 0.01 secs
 RTM Updt Time : < 0.01 secs

Min/Avg/Max Full SPF Times : 0.02/0.00/0.06 secs
Min/Avg/Max RTM Updt Times : 0.02/0.00/0.06 secs

Total Sum Incr SPF Runs : 333
Last Sum Incr SPF run @ : 11/07/2006 18:43:09
Last Sum Incr Calc Time : < 0.01 secs

Total Ext Incr SPF Runs : 0
=====
A:ALA-A#

```

## sr-adj-sets

- Syntax** `sr-adj-sets [adjacency-set-id] [detail]`
- Context** `show>router>ospf`
- Description** This command displays OSPF SR adjacency set log information.
- Parameters** *adjacency-set-id* — Displays SR adjacency set ID information.  
**Values** 1 to 4294967295  
**detail** — Displays detailed SR adjacency set information.
- Output** The following output is an example of **sr-adj-sets** log information.

### Sample Output

```

*A:Dut-C# show router ospf sr-adj-sets
=====
Rtr Base OSPF Instance 0 Adjacency Sets
=====
ID State Label Label Members Members Neighbor
 State Value Type Total Active System Id

1003 active 524285 DL 1 1 0100.2000.1001
1006 active 524284 DL 1 1 0100.2000.1005

```



```

1015 active 18448 SL 1 1 0100.2000.1001
1018 active 18451 SL 1 1 0100.2000.1005
1027 active 524283 DL 1 1 0100.2000.1001
1030 active 524282 DL 1 1 0100.2000.1005
1039 active 18460 SL 1 1 0100.2000.1001
1042 active 18463 SL 1 1 0100.2000.1005
1050 active 524281 DL 2 2 none
1056 active 18471 SL 2 2 none
1063 active 524280 DL 1 1 0100.2000.1001
1066 active 524279 DL 1 1 0100.2000.1005
1075 active 18478 SL 1 1 0100.2000.1001
1078 active 18481 SL 1 1 0100.2000.1005
1087 active 524278 DL 1 1 0100.2000.1001
1090 active 524277 DL 1 1 0100.2000.1005
1099 active 18490 SL 1 1 0100.2000.1001
1102 active 18493 SL 1 1 0100.2000.1005
1110 active 524276 DL 2 2 none
1116 active 18501 SL 2 2 none

```

-----  
DL: dynamic local

SL: static local  
-----

No. of SR Adjacency Sets: 20  
=====

The following table describes the OSPF **sr-adj-sets** output fields.

**Table 26** OSPF SR Adjacency Set Output Fields

| Label                    | Description                                                      |
|--------------------------|------------------------------------------------------------------|
| ID                       | Displays the ID                                                  |
| State                    | Displays the state (active or inactive)                          |
| Label Value              | Displays the label value                                         |
| Label Type               | Displays the label type, Dynamic Local (DL) or Static Local (SL) |
| Members Total            | Displays the number of total members                             |
| Members Active           | Displays the number of active members                            |
| Neighbor System Id       | Displays the neighbor system ID                                  |
| No. of SR Adjacency Sets | Displays the number of SR adjacency sets                         |

## sr-lfa-coverage

**Syntax** **sr-lfa-coverage** [area *area-id*]

**Context** show>router>ospf  
show>router>ospf3

**Description** This command displays the details of the OSPF SR LFA coverage.

**Parameters** **area** *area-id* — Displays all interfaces configured in this area.

**Values** ip-address — a.b.c.d  
area — 0 to 4294967295

### Output

#### Sample Output

```
*A:Dut-C>config>router>ospf3# show router ospf3 0 sr-lfa-coverage
=====
Rtr Base OSPFv3 Instance 0 SR LFA Coverage
=====
SidType Area LFA RLFA TILFA Coverage

node-sid 0.0.0.1 4 (40%) 6 (60%) 0 (0%) 10/10 (100%)
adj-sid 0.0.0.1 0 (0%) 2 (100%) 0 (0%) 2/2 (100%)
=====
```

[Table 27](#) describes the OSPF SR LFA coverage output fields.

**Table 27** OSPF SR LFA Coverage Output Fields

| Label    | Description                         |
|----------|-------------------------------------|
| SidType  | Displays the SID type               |
| Area     | Displays the OSPF area              |
| LFA      | Displays the LFA                    |
| RLFA     | Displays the RLFA                   |
| TIFLA    | Displays the TI-LFA                 |
| Coverage | Displays the percentage of coverage |

## statistics

**Syntax** **statistics**

**Context** show>router>ospf  
show>router>ospf3

**Description** This command displays the global OSPF statistics.

**Output** The following output is an example of OSPF statistics information.

#### Sample Output

```
*A:Dut-C# show router ospf statistics
=====
Rtr Base OSPFv2 Instance 0 Statistics
=====
Rx Packets : 2394 Tx Packets : 2418
Rx Hellos : 2264 Tx Hellos : 2269
Rx DBDs : 11 Tx DBDs : 9
Rx LSRs : 5 Tx LSRs : 4
Rx LSUs : 88 Tx LSUs : 110
Rx LS Acks : 26 Tx LS Acks : 26
New LSAs Recvd : 0 New LSAs Orig : 159
Ext LSAs Count : 0 No of Areas : 2
No of Interfaces : 6 No of Neighbors : 4
Retransmits : 1 Discards : 3
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Adrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 2
Bad Dead Int. : 1 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0
SID range errors : 0 SID errors : 0
Failed SPF Attempts: 0 Bad MTUs : 0
CSPF Requests : 0 CSPF Request Drops : 0
CSPF Path Found : 0 CSPF Path Not Found: 0
Total SPF Runs : 6 Total LFA SPF Runs : 1
Total RLFA SPF Runs: 0
=====

A:ALA-A# show router ospf 1 statistics
=====
Rtr Base OSPFv2 Instance 1 Statistics
=====
Rx Packets : 308462 Tx Packets : 246800
Rx Hellos : 173796 Tx Hellos : 149062
Rx DBDs : 67 Tx DBDs : 48
Rx LSRs : 21 Tx LSRs : 19
Rx LSUs : 105672 Tx LSUs : 65530
Rx LS Acks : 28906 Tx LS Acks : 32141
New LSAs Recvd : 38113 New LSAs Orig : 21067
Ext LSAs Count : 17 No of Areas : 3
Total SPF Runs : 327 Ext SPF Runs : 0
Retransmits : 46 Discards : 0
Bad Networks : 0 Bad Virt Links : 0
Bad Areas : 0 Bad Dest Adrs : 0
Bad Auth Types : 0 Auth Failures : 0
Bad Neighbors : 0 Bad Pkt Types : 0
Bad Lengths : 0 Bad Hello Int. : 0
Bad Dead Int. : 0 Bad Options : 0
Bad Versions : 0 Bad Checksums : 0
Failed SPF Attempts: 0 CSPF Request Drops : 0
CSPF Requests : 0 CSPF Path Not Found: 0
CSPF Path Found : 0
=====
A:ALA-A#
```

The following table describes the command output fields for OSPF statistics.

**Table 28 OSPF Statistics Output Fields**

| Label          | Description                                                                                         |
|----------------|-----------------------------------------------------------------------------------------------------|
| Rx Packets     | The total number of OSPF packets received on all OSPF enabled interfaces.                           |
| Tx Packets     | The total number of OSPF packets transmitted on all OSPF enabled interfaces.                        |
| Rx Hellos      | The total number of OSPF Hello packets received on all OSPF enabled interfaces.                     |
| Tx Hellos      | The total number of OSPF Hello packets transmitted on all OSPF enabled interfaces.                  |
| Rx DBDs        | The total number of OSPF database description packets received on all OSPF enabled interfaces.      |
| Tx DBDs        | The total number of OSPF database description packets transmitted on all OSPF enabled interfaces    |
| Rx LSRs        | The total number of OSPF Link State Requests (LSRs) received on all OSPF enabled interfaces.        |
| Tx LSRs        | The total number of OSPF Link State Requests (LSRs) transmitted on all OSPF enabled interfaces.     |
| Rx LSUs        | The total number of OSPF Link State Update (LSUs) received on all OSPF enabled interfaces.          |
| Tx LSUs        | The total number of OSPF Link State Update (LSUs) transmitted on all OSPF enabled interfaces.       |
| Rx LS Acks     | The total number of OSPF Link State Acknowledgments (LSAs) received on all OSPF enabled interfaces. |
| New LSAs Recvd | The total number of new OSPF Link State Advertisements received on all OSPF enabled interfaces.     |
| New LSAs Orig  | The total number of new OSPF Link State Advertisements originated on all OSPF enabled interfaces.   |
| Ext LSAs Count | The total number of OSPF External Link State Advertisements.                                        |
| No of Areas    | The number of areas configured for this OSPF instance.                                              |
| Total SPF Runs | The total number of incremental SPF runs triggered by new or updated LSAs.                          |
| Ext SPF Runs   | The total number of incremental SPF runs triggered by new or updated type-5 external LSAs.          |

**Table 28 OSPF Statistics Output Fields (Continued)**

| Label          | Description                                                                                                                                                                         |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retransmits    | The total number of OSPF Retransmits transmitted on all OSPF enabled interfaces.                                                                                                    |
| Discards       | The total number of OSPF packets discarded on all OSPF enabled interfaces.                                                                                                          |
| Bad Networks   | The total number of OSPF packets received on all OSPF enabled interfaces with invalid network or mask.                                                                              |
| Bad Virt Links | The total number of OSPF packets received on all OSPF enabled interfaces that are destined to a virtual link that does not exist.                                                   |
| Bad Areas      | The total number of OSPF packets received on all OSPF enabled interfaces with an area mismatch                                                                                      |
| Bad Dest Addr  | The total number of OSPF packets received on all OSPF enabled interfaces with the incorrect IP destination address.                                                                 |
| Bad Auth Types | The total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization type.                                                                        |
| Auth Failures  | The total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization key.                                                                         |
| Bad Neighbors  | The total number of OSPF packets received on all OSPF enabled interfaces where the neighbor information does not match the information this router has for the neighbor.            |
| Bad Pkt Types  | The total number of OSPF packets received on all OSPF enabled interfaces with an invalid OSPF packet type.                                                                          |
| Bad Lengths    | The total number of OSPF packets received on all OSPF enabled interfaces with a total length not equal to the length given in the packet itself.                                    |
| Bad Hello Int. | The total number of OSPF packets received on all OSPF enabled interfaces where the hello interval given in packet was not equal to that configured for the respective interface.    |
| Bad Dead Int.  | The total number of OSPF packets received on all OSPF enabled interfaces where the dead interval given in the packet was not equal to that configured for the respective interface. |
| Bad Options    | The total number of OSPF packets received on all OSPF enabled interfaces with an option that does not match those configured for the respective interface or area.                  |
| Bad Versions   | The total number of OSPF packets received on all OSPF enabled interfaces with bad OSPF version numbers.                                                                             |

---

**status**

|                    |                                       |
|--------------------|---------------------------------------|
| <b>Syntax</b>      | <b>status</b>                         |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3 |
| <b>Description</b> | Displays the general status of OSPF.  |
| <b>Output</b>      | OSPF Status Output Fields             |

The following table describes the command output fields for OSPF status.

**Table 29 OSPF Status Output Fields**

| Label               | Description                                                                                                                                                                                                 |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OSPF Router Id      | A 32-bit integer uniquely identifying the router in the Autonomous System. The SR-OS system defaults to the System IP address or if not configured the 32 least significant bits of the system MAC address. |
| OSPF Version        | The current version number of the OSPF protocol is 2.                                                                                                                                                       |
| OSPF Admin Status   | Disabled — Denotes that the OSPF process is disabled on all interfaces.<br>Enabled — Denotes that the OSPF process is active on at least one interface.                                                     |
| OSPF Oper Status    | Disabled — Denotes that the OSPF process is not operational on all interfaces.<br>Enabled — Denotes that the OSPF process is operational on at least one interface.                                         |
| Preference          | The route preference for OSPF internal routes.                                                                                                                                                              |
| External Preference | The route preference for OSPF external routes.                                                                                                                                                              |
| Backbone Router     | False — This variable indicates that this router is not configured as an OSPF back bone router.<br>True — This variable indicates that this router is configured as an OSPF back bone router.               |
| Area Border Router  | False — This router is not an area border router.<br>True — This router is an area border router.                                                                                                           |
| AS Border Router    | False — This router is not configured as an Autonomous System border router.<br>True — This router is configured as an Autonomous System border router.                                                     |

**Table 29 OSPF Status Output Fields (Continued)**

| Label                      | Description (Continued)                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| OSPF Ldp Sync Admin Status | Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol. |

### Sample Output

```
*A:Dut-A>config>service# show router ospf status
=====
Rtr Base OSPFv2 Instance 0 Status
=====
OSPF Cfg Router Id : 0.0.0.0
OSPF Oper Router Id : 10.20.1.1
OSPF Version : 2
OSPF Admin Status : Enabled
OSPF Oper Status : Enabled
Graceful Restart : Disabled
GR Helper Mode : Disabled
GR Strict LSA Checking : Enabled
Preference : 10
External Preference : 150
Backbone Router : True
Area Border Router : False
AS Border Router : False
Opaque LSA Support : True
Traffic Engineering Support : True
RFC 1583 Compatible : True
Demand Exts Support : False
In Overload State : False
In External Overflow State : False
Exit Overflow Interval : 0
Last Overflow Entered : Never
Last Overflow Exit : Never
External LSA Limit : -1
Reference Bandwidth : 100,000,000 Kbps
Init SPF Delay : 1000 msec
Sec SPF Delay : 1000 msec
Max SPF Delay : 1000 msec
Min LS Arrival Interval : 1000 msec
Init LSA Gen Delay : 5000 msec
Sec LSA Gen Delay : 5000 msec
Max LSA Gen Delay : 5000 msec
Lsa accumulate : 1000 msec
Redistribute delay : 1000 msec
Incremental SPF wait : 1000 msec
Last Ext SPF Run : Never
Ext LSA Csum Sum : 0x0
OSPF Last Enabled : 07/06/2017 18:51:52
Unicast Import : True
Multicast Import : False
Export Policies : None
Import Policies : None
Lfa Policies : None
OSPF Ldp Sync Admin Status : Enabled
```

```

LDP-over-RSVP : Disabled
IGP-Shortcut : Disabled
IPv4 IGP SC Tunn-Nhop : Res-Disabled
Advertise-Tunnel-Link : Disabled
LFA : Disabled
Remote-LFA : Disabled
Max PQ Cost : 65535
TI-LFA : Disabled
Max SR FRR Labels : 2
Export Limit : 0
Export Limit Log Percent : 0
Total Exp Routes : 0
RIB-priority-high prefix list: None
Segment Routing : Enabled
Database export : Disabled
ASN : n/a
Entropy Label : Enabled
Override ELC : Enabled
=====

```

```
*A:Dut-C>config>router>ospf3# show router ospf3 0 status
```

```
=====
Rtr Base OSPFv3 Instance 0 Status
```

```

=====
OSPF Cfg Router Id : 10.20.1.3
OSPF Oper Router Id : 10.20.1.3
OSPF Version : 3
OSPF Admin Status : Enabled
OSPF Oper Status : Enabled
Graceful Restart : Disabled
GR Helper Mode : Disabled
GR Strict LSA Checking : Enabled (operational down)
Preference : 10
External Preference : 150
Backbone Router : False
Area Border Router : False
AS Border Router : True
Traffic Engineering Support : False
Demand Exts Support : False
In Overload State : False
In External Overflow State : False
Exit Overflow Interval : 0
Last Overflow Entered : Never
Last Overflow Exit : Never
External LSA Limit : -1
Reference Bandwidth : 100,000,000 Kbps
Init SPF Delay : 1000 msec
Sec SPF Delay : 1000 msec
Max SPF Delay : 10000 msec
Min LS Arrival Interval : 1000 msec
Init LSA Gen Delay : 5000 msec
Sec LSA Gen Delay : 5000 msec
Max LSA Gen Delay : 5000 msec
Lsa accumulate : 1000 msec
Redistribute delay : 1000 msec
Incremental SPF wait : 1000 msec
Last Ext SPF Run : 10/11/2018 07:39:48
Ext LSA Cksum Sum : 0x10926
OSPF Last Enabled : 10/11/2018 07:39:45

```



```
Unicast Import : True
Multicast Import : False
Export Policies : static
Import Policies : None
Lfa Policies : None
OSPF Ldp Sync Admin Status : Enabled
LDP-over-RSVP : Disabled
IGP-Shortcut : Disabled
IPv6 IGP SC Tunn-Nhop : Res-Disabled
Advertise-Tunnel-Link : Disabled
LFA : Enabled
Remote-LFA : Enabled
Max PQ Cost : 65535
TI-LFA : Disabled
Max SR FRR Labels : 2
Export Limit : 0
Export Limit Log Percent : 0
Total Exp Routes : 1
RIB-priority-high prefix list: None
Segment Routing : Enabled
Extended LSA : sparse
=====
```

virtual-link

- Syntax

**virtual-link [detail]**  
**virtual-link database [detail]**
- Context

show>router>ospf  
show>router>ospf3
- Description

This command displays information for OSPF virtual links.
- Parameters

**database** — Displays the virtual link database.  
**detail** — Provides operational and statistical information about virtual links associated with this router.
- Output

OSPF Virtual Link Output  
  
The following table describes OSPF virtual-link output fields.

Table 30 OSPF Virtual Link Output Fields

| Label           | Description                                                                                                 |
|-----------------|-------------------------------------------------------------------------------------------------------------|
| Nbr Rtr ID      | The router ID(s) of neighboring routers.                                                                    |
| Area Id         | A 32-bit integer which identifies an area.                                                                  |
| Local Interface | The IP address of the local egress interface used to maintain the adjacency to reach this virtual neighbor. |

**Table 30 OSPF Virtual Link Output Fields (Continued)**

| Label           | Description                                                                                                                                                                                                                                                                                                                 |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metric          | The metric value associated with the route. This value is used when importing this static route into other protocols. When the metric is configured as zero then the metric configured in OSPF, default-import-metric, applies. This value is also used to determine which static route to install in the forwarding table. |
| State           | The operational state of the virtual link to the neighboring router.                                                                                                                                                                                                                                                        |
| Authentication  | Specifies whether authentication is enabled for the interface or virtual link.                                                                                                                                                                                                                                              |
| Hello Intrval   | Specifies the length of time, in seconds, between the Hello packets that the router sends on the interface.                                                                                                                                                                                                                 |
| Rtr Dead Intrvl | Specifies the total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since the OSPF admin status was enabled.                                                                                                                                 |
| Tot Rx Packets  | Specifies the total number of OSPF packets received on this interface since the OSPF admin status was enabled.                                                                                                                                                                                                              |
| Rx Hellos       | Specifies the total number of OSPF Hello packets received on this interface since the OSPF admin status was enabled.                                                                                                                                                                                                        |
| Rx DBDs         | Specifies the total number of OSPF Database Description packets received on this interface since the OSPF administrative status was enabled.                                                                                                                                                                                |
| Rx LSRs         | Specifies the total number of Link State Requests (LSRs) received on this interface since the OSPF admin status was enabled.                                                                                                                                                                                                |
| Rx LSUs         | Specifies the total number of Link State Updates (LSUs) received on this interface since the OSPF admin status was enabled.                                                                                                                                                                                                 |
| Rx LS Acks      | Specifies the total number of Link State Acknowledgments received on this interface since the OSPF admin status was enabled.                                                                                                                                                                                                |
| Tot Tx Packets  | Specifies the total number of OSPF packets transmitted on this virtual interface since it was created.                                                                                                                                                                                                                      |
| Tx Hellos       | Specifies the total number of OSPF Hello packets transmitted on this virtual interface since it was created.                                                                                                                                                                                                                |
| Tx DBDs         | Specifies the total number of OSPF database description packets transmitted on this virtual interface.                                                                                                                                                                                                                      |
| Tx LSRs         | Specifies the total number of OSPF Link State Requests (LSRs) transmitted on this virtual interface.                                                                                                                                                                                                                        |

**Table 30 OSPF Virtual Link Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tx LSUs        | Specifies the total number of OSPF Hello packets transmitted on this interface since the OSPF admin status was enabled.                                                                          |
| Tx LS Acks     | Specifies the total number of OSPF Link State Acknowledgments (LSA) transmitted on this virtual interface.                                                                                       |
| Retransmits    | Specifies the total number of OSPF retransmits sent on this interface since the OSPF admin status was last enabled.                                                                              |
| Discards       | Specifies the total number of OSPF packets discarded on this interface since the OSPF admin status was last enabled.                                                                             |
| Bad Networks   | Specifies the total number of OSPF packets received with invalid network or mask since the OSPF admin status was last enabled.                                                                   |
| Bad Versions   | Specifies the total number of OSPF packets received with bad OSPF version numbers since the OSPF admin status was last enabled.                                                                  |
| Bad Areas      | Specifies the total number of OSPF packets received with an area mismatch since the OSPF admin status was last enabled.                                                                          |
| Bad Dest Addrs | Specifies the total number of OSPF packets received with the incorrect IP destination address since the OSPF admin status was last enabled.                                                      |
| Bad Auth Types | Specifies the total number of OSPF packets received with an invalid authorization type since the OSPF admin status was last enabled.                                                             |
| Auth Failures  | Specifies the total number of OSPF packets received with an invalid authorization key since the OSPF admin status was last enabled.                                                              |
| Bad Neighbors  | Specifies the total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled. |
| Bad Pkt Types  | Specifies the total number of OSPF packets received with an invalid OSPF packet type since the OSPF admin status was last enabled.                                                               |
| Bad Lengths    | Specifies the total number of OSPF packets received on this interface with a total length not equal to the length given in the packet itself since the OSPF admin status was last enabled.       |
| Bad Hello Int. | Specifies the total number of OSPF packets received where the hello interval given in packet was not equal to that configured on this interface since the OSPF admin status was last enabled.    |

**Table 30 OSPF Virtual Link Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bad Dead Int.  | Specifies the total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since the OSPF admin status was last enabled. |
| Bad Options    | Specifies the total number of OSPF packets received with an option that does not match those configured for this interface or area since the OSPF admin status was last enabled.                 |
| Retrans Intrvl | Specifies the length of time, in seconds, that OSPF waits before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.                                            |
| Transit Delay  | Specifies the time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link.                                                                     |
| Last Event     | Specifies the date and time when an event was last associated with this OSPF interface.                                                                                                          |

**Sample Output**

```
A:ALA-A# show router ospf 1 virtual-link
=====
Rtr Base OSPFv2 Instance 1 Virtual Links
=====
Nbr Rtr Id Area Id Local Interface Metric State

180.0.0.10 0.0.0.1 180.1.7.12 300 PToP
180.0.0.10 0.0.0.2 180.2.7.12 300 PToP

No. of OSPF Virtual Links: 2
=====
```

```
A:ALA-A# show router ospf virtual-link detail
=====
Rtr Base OSPFv2 Instance 1 Virtual Links (detail)
=====
Neighbor Router Id : 180.0.0.10

Nbr Router Id : 180.0.0.10 Area Id : 0.0.0.1
Local Interface: 180.1.7.12 Metric : 300
State : Point To Point Admin State : Up
Hello Intrvl : 10 sec Rtr Dead Intrvl : 60 sec
Tot Rx Packets : 43022 Tot Tx Packets : 42964
Rx Hellos : 24834 Tx Hellos : 24853
Rx DBDs : 3 Tx DBDs : 2
Rx LSRS : 0 Tx LSRS : 0
Rx LSUs : 15966 Tx LSUs : 16352
Rx LS Acks : 2219 Tx LS Acks : 1757
Retransmits : 0 Discards : 0
Bad Networks : 0 Bad Versions : 0
Bad Areas : 0 Bad Dest Addrs : 0
Bad Auth Types : 0 Auth Failures : 0
```

```
Bad Neighbors : 0
Bad Lengths : 0
Bad Dead Int. : 0
Retrans Intrvl : 5 sec
Last Event : 11/07/2006 17:11:56
Bad Pkt Types : 0
Bad Hello Int. : 0
Bad Options : 0
Transit Delay : 1 sec
Authentication : None

Neighbor Router Id : 180.0.0.10

Nbr Router Id : 180.0.0.10
Local Interface: 180.2.7.12
State : Point To Point
Hello Intrvl : 10 sec
Tot Rx Packets : 43073
Rx Hellos : 24851
Rx DBDs : 3
Rx LSRs : 1
Rx LSUs : 18071
Rx LS Acks : 147
Retransmits : 0
Bad Networks : 0
Bad Areas : 0
Bad Auth Types : 0
Bad Neighbors : 0
Bad Lengths : 0
Bad Dead Int. : 0
Retrans Intrvl : 5 sec
Last Event : 11/07/2006 17:12:00
Area Id : 0.0.0.2
Metric : 300
Admin State : Up
Rtr Dead Intrvl: 60 sec
Tot Tx Packets : 43034
Tx Hellos : 24844
Tx DBDs : 2
Tx LSRs : 1
Tx LSUs : 17853
Tx LS Acks : 334
Discards : 0
Bad Versions : 0
Bad Dest Addrs : 0
Auth Failures : 0
Bad Pkt Types : 0
Bad Hello Int. : 0
Bad Options : 0
Transit Delay : 1 sec
Authentication : MD5
=====
A:ALA-A#
```

## virtual-neighbor

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>virtual-neighbor</b> [ <i>remote ip-address</i> ] [ <i>detail</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | show>router>ospf<br>show>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command displays virtual neighbor information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <b>remote ip-address</b> — Displays the specified IPv4 or IPv6 address. This reduces the amount of output displayed.<br><br><b>Values</b><br>ipv4-address: <ul style="list-style-type: none"><li>• a.b.c.d</li></ul> ipv6-address: <ul style="list-style-type: none"><li>• x:x:x:x:x:x:x (eight 16-bit pieces)</li><li>• x:x:x:x:x:d.d.d.d</li><li>• x: [0 to FFFF]H</li><li>• d: [0 to 255]D</li></ul><br><b>detail</b> — Produces detailed information on the virtual neighbor. This option produces a large amount of data. It is recommended to use <b>detail</b> only when requesting information for a specific neighbor. |

**Output**    OSPF Virtual Neighbor Output

The following table describes OSPF virtual neighbor output fields.

**Table 31      OSPF Virtual Neighbor Output Fields**

| Label            | Description                                                                                                                                                                                      |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nbr IP Addr      | The IP address this neighbor is using in its IP source address. On addressless links, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.                         |
| Nbr Rtr ID       | Specifies the router ID(s) of neighboring routers.                                                                                                                                               |
| Transit Area     | Specifies the transit area ID that links the backbone area with the area that has no physical connection with the backbone.                                                                      |
| Retrans Q Length | The current length of the retransmission queue.                                                                                                                                                  |
| No. of Neighbors | Specifies the total number of OSPF neighbors adjacent on this interface, in a state of INIT or greater, since the OSPF admin status was enabled.                                                 |
| Nbr State        | Specifies the operational state of the virtual link to the neighboring router.                                                                                                                   |
| Options          | Specifies the total number of OSPF packets received with an option that does not match those configured for this virtual interface or transit area since the OSPF admin status was enabled.      |
| Events           | Specifies the total number of events that have occurred since the OSPF admin status was enabled.                                                                                                 |
| Last Event Time  | Specifies the date and time when an event was last associated with this OSPF interface.                                                                                                          |
| Up Time          | Specifies the uninterrupted time, in hundredths of seconds, the adjacency to this neighbor has been up.                                                                                          |
| Time Before Dead | Specifies the amount of time, in seconds, until the dead router interval expires.                                                                                                                |
| Bad Nbr States   | Specifies the total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled. |
| LSA Inst fails   | Specifies the total number of times an LSA could not be installed into the LSDB due to a resource allocation issue since the OSPF admin status was last enabled.                                 |

**Table 31 OSPF Virtual Neighbor Output Fields (Continued)**

| Label             | Description (Continued)                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bad Seq Nums      | Specifies the total number of times when a database description packet was received with a sequence number mismatch since the OSPF admin status was last enabled.                      |
| Bad MTUs          | Specifies the total number of times when the MTU in a received database description packet was larger than the MTU of the receiving interface since the OSPF admin status was enabled. |
| Bad Packets       | Specifies the total number of times when an LS update was received with an illegal LS type or an option mismatch since the OSPF admin status was enabled.                              |
| LSA not in LSDB   | Specifies the total number of times when an LS request was received for an LSA not installed in the LSDB of this router since the OSPF admin status was enabled.                       |
| Option Mismatches | Specifies the total number of times when a LS update was received with an option mismatch since the OSPF admin status was enabled.                                                     |
| Nbr Duplicates    | Specifies the total number of times when a duplicate database description packet was received during the Exchange state since the OSPF admin status was enabled.                       |

### Sample Output

```

A:ALA-A# show router ospf 1 virtual-neighbor
=====
Rtr Base OSPFv2 Instance 1 Virtual Neighbors
=====
Nbr IP Addr Nbr Rtr Id Nbr State Transit Area RetxQ Len Dead Time

180.1.6.10 180.0.0.10 Full 0.0.0.1 0 58
180.2.9.10 180.0.0.10 Full 0.0.0.2 0 52

No. of Neighbors: 2
=====
A:ALA-A#

A:ALA-A# show router ospf virtual-neighbor detail
=====
Rtr Base OSPFv2 Instance 0 Virtual Neighbors (detail)
=====
Virtual Neighbor Router Id : 180.0.0.10

Neighbor IP Addr : 180.1.6.10 Neighbor Rtr Id : 180.0.0.10
Neighbor State : Full Transit Area : 0.0.0.1
Retrans Q Length : 0 Options : -E-
Events : 4 Last Event Time : 11/07/2006 17:11:56
Up Time : 2d 17:47:17 Time Before Dead : 57 sec
Bad Nbr States : 1 LSA Inst fails : 0

```

```

Bad Seq Nums : 0
Bad Packets : 0
Option Mismatches: 0
Bad MTUs : 0
LSA not in LSDB : 0
Nbr Duplicates : 0

Virtual Neighbor Router Id : 180.0.0.10

Neighbor IP Addr : 180.2.9.10
Neighbor State : Full
Retrans Q Length : 0
Events : 4
Up Time : 2d 17:47:14
Bad Nbr States : 1
Bad Seq Nums : 0
Bad Packets : 0
Option Mismatches: 0
Neighbor Rtr Id : 180.0.0.10
Transit Area : 0.0.0.2
Options : -E--
Last Event Time : 11/07/2006 17:11:59
Time Before Dead : 59 sec
LSA Inst fails : 0
Bad MTUs : 0
LSA not in LSDB : 0
Nbr Duplicates : 0
=====
A:ALA-A#

```

### 3.15.2.2 Clear Commands

#### ospf

**Syntax** `ospf [ospf-instance]`

**Context** `clear>router`

**Description** This command clears and resets OSPF protocol entities.

**Parameters** *ospf-instance* — Clears the specified OSPF instance.

**Values** 0 to 31

#### ospf3

**Syntax** `ospf3 [ospf-instance]`

**Context** `clear>router`

**Description** This command clears and resets OSPF3 protocol entities.

**Parameters** *ospf-instance* — Clears the specified OSPF3 instance.

**Values** 0 to 31 | 64 to 95  
 0 to 31 — IPv6-unicast address-family  
 64 to 95 — IPv4-unicast address-family



## database

|                    |                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database [purge]</b>                                                                                                                             |
| <b>Context</b>     | clear>router>ospf<br>clear>router>ospf3                                                                                                             |
| <b>Description</b> | This command clears all LSAs received from other nodes, sets all adjacencies other than two-way to one-way, and refreshes all self-originated LSAs. |
| <b>Parameters</b>  | <b>purge</b> — Clears all self-originated LSAs and re-originates all self-originated LSAs.                                                          |

## export

|                    |                                                          |
|--------------------|----------------------------------------------------------|
| <b>Syntax</b>      | <b>export</b>                                            |
| <b>Context</b>     | clear>router>ospf<br>clear>router>ospf3                  |
| <b>Description</b> | This command re-evaluates all effective export policies. |

## neighbor

|                    |                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>neighbor [ip-int-name   ip-address]</b><br><b>neighbor [ip-int-name] [router-id]</b>                                                                                                                                                                                             |
| <b>Context</b>     | clear>router>ospf<br>clear>router>ospf3                                                                                                                                                                                                                                             |
| <b>Description</b> | This command marks the neighbor as dead and re-initiates the affected adjacencies.                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>ip-int-name</i> — Clears all neighbors for the interface specified by this interface name.<br><i>ip-address</i> — Clears all neighbors for the interface specified by this IP address.<br><i>router-id</i> — Clears all neighbors for the interface specified by this router ID. |

## overload

|                    |                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>overload {rtm   fib   rtr-adv-lsa-limit}</b>                                                                                                                           |
| <b>Context</b>     | clear>router>ospf<br>clear>router>ospf3                                                                                                                                   |
| <b>Description</b> | This command clears the OSPF or OSPF3 overload.                                                                                                                           |
| <b>Parameters</b>  | <b>rtm</b> — Clears the overload because OSPF or OSPF3 reached the configured maximum route limit set with <b>maximum-routes</b> or <b>maximum-ipv6-routes</b> in a VPRN. |

- fib** — Clears the overload because adding routes to the hardware FIB failed.
- rtr-adv-lsa-limit** — Clears the overload because OSPF or OSPF3 exceeded the configured maximum limit on LSAs advertised by another router, which was set with **rtr-adv-lsa-limit**.

statistics

|                    |                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                                                                                     |
| <b>Context</b>     | clear>router>ospf<br>clear>router>ospf3                                                               |
| <b>Description</b> | This command clears all neighbor, router, interface, SPF and global statistics of this OSPF instance. |

3.15.2.3 Debug Commands

ospf

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ospf</b> [ <i>ospf-instance</i> ]<br><b>no ospf</b> [ <i>ospf-instance</i> ] |
| <b>Context</b>     | debug>router                                                                    |
| <b>Description</b> | Indicates the OSPF instance for debugging purposes.                             |
| <b>Parameters</b>  | <i>ospf-instance</i> — The OSPF instance.<br><b>Values</b> 0 to 31              |

ospf3

|                    |                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ospf3</b> [ <i>ospf-instance</i> ]<br><b>no ospf3</b> [ <i>ospf-instance</i> ]                                                                                                  |
| <b>Context</b>     | debug>router                                                                                                                                                                       |
| <b>Description</b> | Indicates the OSPF3 instance for debugging purposes.                                                                                                                               |
| <b>Parameters</b>  | <i>ospf-instance</i> — Clears the specified OSPF3 instance.<br><b>Values</b> 0 to 31   64 to 95<br>0 to 31 — IPv6-unicast address-family<br>64 to 95 — IPv4-unicast address-family |

## area

|                    |                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>area</b> [ <i>area-id</i> ]<br><b>no area</b>                                                                 |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                          |
| <b>Description</b> | This command enables debugging for an OSPF area.                                                                 |
| <b>Parameters</b>  | <i>area-id</i> — Specifies the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer. |
| <b>Values</b>      | ip-address — a.b.c.d<br>area — 0 to 4294967295                                                                   |

## area-range

|                    |                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>area-range</b> [ <i>ip-address</i> ]<br><b>no area-range</b>                                                                                      |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                              |
| <b>Description</b> | This command enables debugging for an OSPF area range.                                                                                               |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the IPv4 or IPv6 address for the range used by the ABR to advertise the area into another area.                        |
| <b>Values</b>      | ipv4-address:<br>• a.b.c.d<br>ipv6-address:<br>• x:x:x:x:x:x:x (eight 16-bit pieces)<br>• x:x:x:x:x:d.d.d.d<br>• x: [0 to FFFF]H<br>• d: [0 to 255]D |

## cspf

|                    |                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cspf</b> [ <i>ip-address</i> ]<br><b>no cspf</b>                                     |
| <b>Context</b>     | debug>router>ospf                                                                       |
| <b>Description</b> | This command enables debugging for an OSPF constraint-based shortest path first (CSPF). |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the IP address for the range used for CSPF.               |

---

## graceful-restart

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] graceful-restart</b>                                        |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                             |
| <b>Description</b> | This command enables debugging for OSPF and OSPF3 graceful restart. |

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ]<br><b>interface</b> [ <i>interface-name</i> ]<br><b>no interface</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command enables debugging for an OSPF and OSPF3 interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the IP interface name, in the <b>debug&gt;router&gt;ospf</b> context. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.<br><br><i>ip-address</i> — Specifies the interface's IP address, in the <b>debug&gt;router&gt;ospf</b> context.<br><br><i>interface-name</i> — Specifies the interface name, in the <b>debug&gt;router&gt;ospf3</b> context. |

## leak

|                    |                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>leak</b> [ <i>ip-address</i> ]<br><b>no leak</b>                                                                                                                                                                                  |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                                                                              |
| <b>Description</b> | This command enables debugging for OSPF leaks.                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the IPv4 or IPv6 address to debug OSPF leaks.<br><br><b>Values</b><br>ipv4-address:<br>• a.b.c.d<br>ipv6-address:<br>• x:x:x:x:x:x:x (eight 16-bit pieces)<br>• x:x:x:x:x:d.d.d.d<br>• x: [0 to FFFF]H |

- d: [0 to 255]D

## lsdb

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsdb</b> [ <i>type</i> ] [ <i>ls-id</i> ] [ <i>adv-rtr-id</i> ] [ <b>area</b> <i>area-id</i> ]<br><b>no lsdb</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command enables debugging for an OSPF link-state database (LSDB).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>type</i> — Specifies the OSPF link-state database (LSDB) type.</p> <p><b>Values</b>      in the <b>ospf</b> context — router, network, summary, asbr, extern, nssa, area-opaque, as-opaque, link-opaque<br/>                 in the <b>ospf3</b> context — router, network, inter-area-pfx, inter-area-rtr, external, nssa, intra-area-pfx, rtr-info-link, rtr-info-area, rtr-info-as</p> <p><i>ls-id</i> — Specifies an LSA type specific field containing either a router ID or an IP address. It identifies the piece of the routing domain being described by the advertisement.</p> <p><i>adv-rtr-id</i> — Specifies the router identifier of the router advertising the LSA.</p> <p><b>area</b> <i>area-id</i> — Specifies a 32-bit integer uniquely identifying an area.</p> <p><b>Values</b>      ip-address — a.b.c.d<br/>                 area — 0 to 4294967295</p> |

## misc

|                    |                                                               |
|--------------------|---------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] misc</b>                                              |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                       |
| <b>Description</b> | This command enables debugging for miscellaneous OSPF events. |

## neighbor

|                    |                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>neighbor</b> [ <i>ip-int-name</i>   <i>ip-address</i> ]<br><b>neighbor</b> [ <i>ip-int-name</i> ] [ <i>router-id</i> ]<br><b>no neighbor</b> |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                         |
| <b>Description</b> | This command enables debugging for an OSPF or OSPF3 neighbor.                                                                                   |

- Parameters**
- ip-int-name* — Specifies the neighbor interface name.
  - ip-address* — Specifies neighbor information for the neighbor identified by the specified IP address, in the **debug>router>ospf** context.
  - router-id* — Specifies neighbor information for the neighbor identified by the specified router ID, in the **debug>router>ospf3** context.

## nssa-range

- Syntax** **nssa-range** [*ip-address*]  
**no nssa-range**
- Context** debug>router>ospf  
debug>router>ospf3
- Description** This command enables debugging for an NSSA range.
- Parameters** *ip-address* — Specifies the IPv4 or IPv6 address range to debug OSPF or OSPF3 leaks.
- Values**
- ipv4-address:
    - a.b.c.d
  - ipv6-address:
    - x:x:x:x:x:x:x (eight 16-bit pieces)
    - x:x:x:x:x:d.d.d.d
    - x: [0 to FFFF]H
    - d: [0 to 255]D

## packet

- Syntax** **packet** [*packet-type*] [*interface-name*] [**ingress** | **egress**] [**detail**]  
**packet** [*packet-type*] [*interface-name*] [**ingress** | **egress** | **drop**] [**detail**]  
**no packet**
- Context** debug>router>ospf  
debug>router>ospf3
- Description** This command enables debugging for OSPF packets.
- Parameters** *packet-type* — Specifies the OSPF packet type to debug.
- Values** hello, dbdescr, lsrequest, lsupdate, lsack
- interface-name* — Specifies the interface to debug, up to a maximum of 32 characters.
- ingress** — Specifies to display ingress packets.
- egress** — Specifies to display egress packets.
- drop** — Specifies to display dropped packets.

## rsvp-shortcut

|                    |                                                                       |
|--------------------|-----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rsvp-shortcut</b> [ <i>ip-address</i> ]<br><b>no rsvp-shortcut</b> |
| <b>Context</b>     | debug>router>ospf                                                     |
| <b>Description</b> | This command debugs the OSPFv2 RSVP shortcut.                         |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the IP address to debug.                |

## rtm

|                    |                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rtm</b> [ <i>ip-address</i> ]<br><b>no rtm</b>                                                                                                                                                                                                                 |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                                                                                                           |
| <b>Description</b> | This command enables debugging for OSPF RTM.                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the IP address to debug.                                                                                                                                                                                                            |
| <b>Values</b>      | ipv4-address: <ul style="list-style-type: none"> <li>• a.b.c.d</li> </ul> ipv6-address: <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> <li>• d: [0 to 255]D</li> </ul> |

## sham-neighbor

|                    |                                                                                  |
|--------------------|----------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sham-neighbor</b> [ <i>ip-address</i> ]<br><b>no sham-neighbor</b>            |
| <b>Context</b>     | show>router>ospf                                                                 |
| <b>Description</b> | This command enables debugging of the OSPFv2 sham-link neighbor.                 |
| <b>Parameters</b>  | <i>ip-address</i> — Debugs the sham-link neighbor identified by this IP address. |

## spf

|               |                                                 |
|---------------|-------------------------------------------------|
| <b>Syntax</b> | <b>spf</b> [ <i>type</i> ] [ <i>dest-addr</i> ] |
|---------------|-------------------------------------------------|

---

**no spf**

|                    |                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                                                                                             |
| <b>Description</b> | This command enables debugging for OSPF SPF. Information regarding overall SPF start and stop times will be shown. To see detailed information regarding the SPF calculation of a given route, the route must be specified as an optional argument. |
| <b>Parameters</b>  | <i>type</i> — Specifies the area to debug<br><b>Values</b> intra-area, inter-area, external<br><i>dest-addr</i> — Specifies the destination IP address to debug.                                                                                    |

**tunnel-endpoint**

|                    |                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-endpoint</b> [ <b>tunnel-spf</b> <i>ip-address</i> ] [ <b>tunnel-leak</b> <i>ip-address</i> ]                                                                 |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                                                                                                                 |
| <b>Description</b> | This command enables debugging for OSPF tunnel endpoints.                                                                                                               |
| <b>Parameters</b>  | <b>tunnel-spf</b> — Specifies the tunnel SPF IP address.<br><b>tunnel-leak</b> — Specifies the tunnel leak IP address.<br><i>ip-address</i> — Specifies the IP address. |

**virtual-neighbor**

|                    |                                                                            |
|--------------------|----------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>virtual-neighbor</b> [ <i>router-id</i> ]<br><b>no virtual-neighbor</b> |
| <b>Context</b>     | debug>router>ospf<br>debug>router>ospf3                                    |
| <b>Description</b> | This command enables debugging for an OSPF virtual neighbor.               |
| <b>Parameters</b>  | <i>router-id</i> — Specifies the router ID of the virtual neighbor.        |



### 3.15.2.4 Tools Commands

#### sr-adjacencies

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sr-adjacencies</b> [ <i>ip-int-name</i>   <i>ip-address</i> ] [ <b>detail</b> ]<br><b>sr-adjacencies</b> [ <b>remote</b> <i>ip-address</i> ] [ <b>detail</b> ]                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | tools>dump>router>ospf<br>tools>dump>router>ospf3                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command displays OSPFv2 or OSPFv3 SR adjacency information.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>ip-int-name</i> — Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.<br><br><i>ip-address</i> — Specifies the interface's IP address.<br><br><b>Values</b> a.b.c.d |

#### Output

##### Sample Output

```
*A:Dut-C>config>router>ospf3# /tools dump router ospf3 0 sr-adjacencies detail
=====
Rtr Base OSPFv3 Instance 0 Segment Routing Adjacencies (detail)
Legend:
label stack is ordered from bottom-most to top-most
=====

Interface DUTC_TO_DUTA.1.0

Nbr RouterID : 10.20.1.1 Nbr State : FULL
Label : 524287 PgID : 9
SID Protection : Enabled Curr Hold Time : 15
Area ID : 0.0.0.1 SR MTU : 1492
Neighbor IP : fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0"
LFA Labels : 30004/70005
LFA Next hop : fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0"
LFA Glob Itf id: 26 LFA Loc Itf id : 3

Interface DUTC_TO_DUTF.1.0

Nbr RouterID : 10.20.1.6 Nbr State : FULL
Label : 524286 PgID : 10
SID Protection : Enabled Curr Hold Time : 15
Area ID : 0.0.0.1 SR MTU : 1492
Neighbor IP : fe80::200:ff:fe00:6-"DUTC_TO_DUTF.1.0"
LFA Labels : 30003/20005
LFA Next hop : fe80::200:ff:fe00:1-"DUTC_TO_DUTA.1.0"
```

```
LFA Glob Itf id: 25 LFA Loc Itf id : 2

No. of SR Adjacencies: 2
=====
```

sr-database

- Syntax

sr-database [sid sid] [prefix ip-address] [nh-type {sr | ldp | rsvp}] [detail]
- Context

tools>dump>router>ospf  
tools>dump>router>ospf3
- Description

This command displays OSPFv2 or OSPFv3 SR database information.
- Parameters

sid — Displays information for the specific segment identifier.

Values

0 to 524287

ip-address — Specifies the IP address in dotted decimal notation.

Values

ipv4-address:
  - a.b.c.dipv6-address:
  - x:x:x:x:x:x:x [eight 16-bit pieces]
  - x:x:x:x:x:d.d.d.d
  - x: [0 to FFFF]H
  - d: [0 to 255]D

Output

Sample Output

```
A:Dut-E# tools dump router ospf sr-database
=====
Rtr Base OSPFv2 Instance 0 Segment Routing Database
=====
SID Label St Type Prefix AdvRtr Area Flags Stitching
 FRR

0 60000 +R T1 1.0.33.3/32 10.20.1.3 0.0.0.0 [NnP]] T(R) -
1 60001 +R T1 1.0.44.4/32 10.20.1.4 0.0.0.0 [NnP]] Tn(R) -
2 60002 +R LT1 1.0.55.5/32 10.20.1.5 0.0.0.0 [NnP]] - -
3 60003 +R T1 1.0.66.6/32 10.20.1.6 0.0.0.0 [NnP]] T(R) -
4 60004 +R T1 1.0.11.1/32 10.20.1.1 0.0.0.0 [NnP]] Tn(R) -
5 60005 +R T1 1.0.22.2/32 10.20.1.2 0.0.0.0 [NnP]] Tn(R) -
6 60006 +R T1 10.20.1.3/32 10.20.1.3 0.0.0.0 [NnP]] T(R) -
```

```

7 60007 +R T1 10.20.1.4/32
 10.20.1.4 0.0.0.0 [NnP] Tn(R) -
8 60008 +R LT1 10.20.1.5/32
 10.20.1.5 0.0.0.0 [NnP] -
9 60009 +R T1 10.20.1.6/32
 10.20.1.6 0.0.0.0 [NnP] T(R) -
10 60010 +R T1 10.20.1.1/32
 10.20.1.1 0.0.0.0 [NnP] Tn(R) -
11 60011 +R T1 10.20.1.2/32
 10.20.1.2 0.0.0.0 [NnP] Tn(R) -

```

-----  
No. of Entries: 12  
-----

St: R:reported I:incomplete W:wrong N:not reported F:failed

+:SR-ack -:no route

Type: L:local M: mapping Srv Tx: route type

FRR: L:Lfa R:RLfa T:TLfa (R):Reported (F):Failed

Rn: RLfa + node-protect Tn: TLfa + node-protect

=====

\*A:Dut-F# tools dump router ospf sr-database detail

=====

Rtr Base OSPFv2 Instance 0 Segment Routing Database (detail)

=====

| SID  | Label | Sts  | Prefix | Type | AdvRtr | Area | Flags |
|------|-------|------|--------|------|--------|------|-------|
| TnID | FRR   | lsId | AdvRtr | Area | Type   |      |       |

-----

|    |       |    |            |              |         |              |  |
|----|-------|----|------------|--------------|---------|--------------|--|
| 0  | 70000 | +R | 1.0.33.3/  |              |         |              |  |
| 32 | T1    |    | 10.20.1.3  | 0.0.0.1 [NnP | ]       | 2 R(R)       |  |
|    |       |    | 7.16.0.9   | 10.20.1.2    | 0.0.0.0 | T3           |  |
|    |       |    | 7.0.0.4    | 10.20.1.3    | 0.0.0.1 | T1 <<<< Best |  |
| 1  | 70001 | +R | 1.0.44.4/  |              |         |              |  |
| 32 | T1    |    | 10.20.1.4  | 0.0.0.0 [NnP | ]       | 2 L(R)       |  |
|    |       |    | 7.16.0.5   | 10.20.1.2    | 0.0.0.1 | T3           |  |
|    |       |    | 7.0.0.4    | 10.20.1.4    | 0.0.0.0 | T1 <<<< Best |  |
| 2  | 70002 | +R | 1.0.55.5/  |              |         |              |  |
| 32 | T1    |    | 10.20.1.5  | 0.0.0.0 [NnP | ]       | 2 R(R)       |  |
|    |       |    | 7.16.0.3   | 10.20.1.2    | 0.0.0.1 | T3           |  |
|    |       |    | 7.0.0.4    | 10.20.1.5    | 0.0.0.0 | T1 <<<< Best |  |
| 3  | 70003 | +R | 1.0.66.6/  |              |         |              |  |
| 32 | LT1   |    | 10.20.1.6  | 0.0.0.1 [NnP | ]       | 0 -          |  |
|    |       |    | 7.16.0.8   | 10.20.1.2    | 0.0.0.0 | T3           |  |
|    |       |    | 7.0.0.6    | 10.20.1.6    | 0.0.0.1 | T1 <<<< Best |  |
| 4  | 70004 | +R | 1.0.11.1/  |              |         |              |  |
| 32 | T1    |    | 10.20.1.1  | 0.0.0.1 [NnP | ]       | 2 L(R)       |  |
|    |       |    | 7.16.0.13  | 10.20.1.2    | 0.0.0.0 | T3           |  |
|    |       |    | 7.0.0.4    | 10.20.1.1    | 0.0.0.1 | T1 <<<< Best |  |
| 5  | 70005 | +R | 1.0.22.2/  |              |         |              |  |
| 32 | T1    |    | 10.20.1.2  | 0.0.0.1 [NnP | ]       | 2 R(R)       |  |
|    |       |    | 7.16.0.11  | 10.20.1.2    | 0.0.0.0 | T3           |  |
|    |       |    | 7.0.0.6    | 10.20.1.2    | 0.0.0.1 | T1 <<<< Best |  |
| 6  | 70006 | +R | 10.20.1.3/ |              |         |              |  |
| 32 | T1    |    | 10.20.1.3  | 0.0.0.1 [NnP | ]       | 2 R(R)       |  |
|    |       |    | 7.16.0.10  | 10.20.1.2    | 0.0.0.0 | T3           |  |
|    |       |    | 7.0.0.5    | 10.20.1.3    | 0.0.0.1 | T1 <<<< Best |  |
| 7  | 70007 | +R | 10.20.1.4/ |              |         |              |  |
| 32 | T1    |    | 10.20.1.4  | 0.0.0.0 [NnP | ]       | 2 L(R)       |  |
|    |       |    | 7.16.0.6   | 10.20.1.2    | 0.0.0.1 | T3           |  |

```

 7.0.0.5 10.20.1.4 0.0.0.0 T1 <<<< Best
8 70008 +R 10.20.1.5/
32 T1 10.20.1.5 0.0.0.0 [NnP] 2 R(R)
 7.16.0.4 10.20.1.2 0.0.0.1 T3
 7.0.0.5 10.20.1.5 0.0.0.0 T1 <<<< Best
9 70009 +R 10.20.1.6/
32 LT1 10.20.1.6 0.0.0.0 [NnP] 0 -
 7.16.0.2 10.20.1.2 0.0.0.1 T3
 7.0.0.7 10.20.1.6 0.0.0.0 T1 <<<< Best
10 70010 +R 10.20.1.1/
32 T1 10.20.1.1 0.0.0.1 [NnP] 2 L(R)
 7.16.0.12 10.20.1.2 0.0.0.0 T3
 7.0.0.5 10.20.1.1 0.0.0.1 T1 <<<< Best
11 70011 +R 10.20.1.2/
32 T1 10.20.1.2 0.0.0.0 [NnP] 2 R(R)
 7.16.0.7 10.20.1.2 0.0.0.1 T3
 7.0.0.7 10.20.1.2 0.0.0.0 T1 <<<< Best
994 70994 +R 1.0.44.4/
32 T1 10.20.1.5 0.0.0.0 [NnPB] 0 -
 7.16.0.1 10.20.1.5 0.0.0.0 T1 <<<< Best
995 70995 +R 1.0.55.5/
32 T1 10.20.1.4 0.0.0.0 [NnPB] 0 -
 7.16.0.1 10.20.1.4 0.0.0.0 T1 <<<< Best
996 70996 +R 1.0.22.2/
32 LT1 10.20.1.6 0.0.0.1 [NnPB] 0 -
997 70997 - 1.0.66.6/
32 T1 10.20.1.2 0.0.0.1 [NnPB] - -
 7.16.0.1 10.20.1.2 0.0.0.1 T1 <<<< Best
998 70998 +R 1.0.33.3/
32 T1 10.20.1.1 0.0.0.1 [NnPB] 0 -
 7.16.0.1 10.20.1.1 0.0.0.1 T1 <<<< Best
999 70999 +R 1.0.11.1/
32 T1 10.20.1.3 0.0.0.1 [NnPB] 0 -
 7.16.0.1 10.20.1.3 0.0.0.1 T1 <<<< Best

```

```

Sts: R:reported I:incomplete W:wrong N:not reported F:failed +:SR-ack -
:no route

```

```
Type: L:local M: mapping Srv Tx: route type

```

```
FRR: L:Lfa R:RLfa (R):Reported (F):Failed

```

```
=====
*A:Dut-F#

```

---

## 4 IS-IS

### 4.1 Configuring IS-IS


Intermediate-system-to-intermediate-system (IS-IS) is a link-state interior gateway protocol (IGP) which uses the Shortest Path First (SPF) algorithm to determine routes. Routing decisions are made using the link-state information. IS-IS evaluates topology changes and, if necessary, performs SPF recalculations.

Entities within IS-IS include networks, intermediate systems, and end systems. In IS-IS, a network is an autonomous system (AS), or routing domain, with end systems and intermediate systems. A router is an intermediate system. End systems are network devices which send and receive protocol data units (PDUs), the OSI term for packets. Intermediate systems send, receive, and forward PDUs.

End system and intermediate system protocols allow routers and nodes to identify each other. IS-IS sends out link-state updates periodically throughout the network, so each router can maintain current network topology information.

IS-IS supports large ASs by using a two-level hierarchy. A large AS can be administratively divided into smaller, more manageable areas. A system logically belongs to one area. Level 1 routing is performed within an area. Level 2 routing is performed between areas. The routers can be configured as Level 1, Level 2, or both Level 1/2.

[Figure 14](#) displays an example of an IS-IS routing domain.

**Figure 14 IS-IS Routing Domain**

OSRG033

### 4.1.1 Routing

OSI IS-IS routing uses two-level hierarchical routing. A routing domain can be partitioned into areas. Level 1 routers know the topology in their area, including all routers and end systems in their area but do not know the identity of routers or destinations outside of their area. Level 1 routers forward traffic with destinations outside of their area to a Level 2 router in their area.

Level 2 routers know the Level 2 topology, and know which addresses are reachable by each Level 2 router. Level 2 routers do not need to know the topology within any Level 1 area, except to the extent that a Level 2 router can also be a Level 1 router within a single area. By default, only Level 2 routers can exchange PDUs or routing information directly with external routers located outside the routing domain.

In IS-IS, there are two types of routers:

- Level 1 intermediate systems — Routing is performed based on the area ID portion of the ISO address called the *network entity title* (NET). Level 1 systems route within an area. They recognize, based on the destination address, whether the destination is within the area. If so, they route toward the destination. If not, they route to the nearest Level 2 router.

- Level 2 intermediate systems — Routing is performed based on the area address. They route toward other areas, disregarding other area's internal structure. A Level 2 intermediate system can also be configured as a Level 1 intermediate system in the same area.

The Level 1 router's area address portion is manually configured (see [ISO Network Addressing](#)). A Level 1 router will not become a neighbor with a node that does not have a common area address. However, if a Level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, then the Level 1 router accepts the other node as a neighbor, as address B is common to both routers. Level 2 adjacencies are formed with other Level 2 nodes whose area addresses do not overlap. If the area addresses do not overlap, the link is considered by both routers to be Level 2 only and only Level 2 LSPDs flow on the link.

Within an area, Level 1 routers exchange LSPs which identify the IP addresses reachable by each router. Specifically, zero or more IP address, subnet mask, and metric combinations can be included in each LSP. Each Level 1 router is manually configured with the IP address, subnet mask, and metric combinations, which are reachable on each interface. A Level 1 router routes as follows:

- If a specified destination address matches an IP address, subnet mask, or metric reachable within the area, the PDU is routed via Level 1 routing.
- If a specified destination address does not match any IP address, subnet mask, or metric combinations listed as reachable within the area, the PDU is routed towards the nearest Level 2 router.

Level 2 routers include in their LSPs, a complete list of IP address, subnet mask, and metrics specifying all the IP addresses which reachable in their area. This information can be obtained from a combination of the Level 1 LSPs (by Level 1 routers in the same area). Level 2 routers can also report external reachability information, corresponding to addresses reachable by routers in other routing domains or autonomous systems.

## 4.1.2 IS-IS Frequently Used Terms

- Area — An area is a routing sub-domain which maintains detailed routing information about its own internal composition, and also maintains routing information which allows it to reach other routing sub-domains. Areas correspond to the Level 1 sub-domain.
- End system — End systems send NPDUs to other systems and receive NPDUs from other systems, but do not relay NPDUs. This International Standard does not specify any additional end system functions beyond those supplied by ISO 8473 and ISO 9542.

- Neighbor — A neighbor is an adjacent system reachable by traversing a single sub-network by a PDU.
- Adjacency — An adjacency is a portion of the local routing information which pertains to the reachability of a single neighboring end or intermediate system over a single circuit. Adjacencies are used as input to the decision process to form paths through the routing domain. A separate adjacency is created for each neighbor on a circuit and for each level of routing (Level 1 and Level 2) on a broadcast circuit.
- Circuit — The subset of the local routing information base pertinent to a single local Subnetwork Point of Attachments (SNPAs).
- Link — The communication path between two neighbors. A link is up when communication is possible between the two SNPAs.
- Designated IS — The intermediate system on a LAN which is designated to perform additional duties. In particular, the designated IS generates link-state PDUs on behalf of the LAN, treating the LAN as a pseudonode.
- Pseudonode — Where a broadcast sub-network has  $n$  connected intermediate systems, the broadcast sub-network itself is considered to be a pseudonode. The pseudonode has links to each of the  $n$  intermediate systems and each of the ISs has a single link to the pseudonode (rather than  $n-1$  links to each of the other intermediate systems). Link-state PDUs are generated on behalf of the pseudonode by the designated IS.
- Broadcast sub-network — A multi-access subnetwork that supports the capability of addressing a group of attached systems with a single PDU.
- General topology sub-network — A topology that is modeled as a set of point-to-point links, each of which connects two systems. There are several generic types of general topology subnetworks, multipoint links, permanent point-to-point links, dynamic and static point-to-point links.
- Routing sub-domain — A routing sub-domain consists of a set of intermediate systems and end systems located within the same routing domain.
- Level 2 sub-domain — Level 2 sub-domain is the set of all Level 2 intermediate systems in a routing domain.

### 4.1.3 ISO Network Addressing

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP).



---

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.


Each network entity has a special network address called a Network Entity Title (NET). Structurally, a NET is identical to an NSAP address but has an n-selector of 00. Most end systems have one NET. Intermediate systems can have up to three area IDs (area addresses).

NSAP addresses are divided into three parts. Only the area ID portion is configurable.

- Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.
- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

Of the total 20 bytes comprising the NET, only the first 13 bytes, the area ID portion, can be manually configured. As few as one byte can be entered or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

Routers with common area addresses form Level 1 adjacencies. Routers with no common NET addresses form Level 2 adjacencies, if they are capable ([Figure 15](#)).

**Figure 15 Using Area Addresses to Form Adjacencies**

### 4.1.3.1 IS-IS PDU Configuration

The following PDUs are used by IS-IS to exchange protocol information:

- IS-IS hello PDU — Routers with IS-IS enabled send hello PDUs to IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- Link-state PDUs — Contain information about the state of adjacencies to neighboring IS-IS systems. LSPs are flooded periodically throughout an area.
- Complete sequence number PDUs — In order for all routers to maintain the same information, CSNPs inform other routers that some LSPs can be outdated or missing from their database. CSNPs contain a complete list of all LSPs in the current IS-IS database.
- Partial sequence number PDUs (PSNPs) — PSNPs are used to request missing LSPs and acknowledge that an LSP was received.

### 4.1.3.2 IS-IS Operations

The routers perform IS-IS routing as follows:

- Hello PDUs are sent to the IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- IS-IS neighbor relationships are formed if the hello PDUs contain information that meets the criteria for forming an adjacency.
- The routers can build a link-state PDU based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- The routers flood LSPs to the adjacent neighbors except the neighbor from which they received the same LSP. The link-state database is constructed from these LSPs.
- A Shortest Path Tree (SPT) is calculated by each IS, and from this SPT the routing table is built.

## 4.1.4 IS-IS Route Summarization

IS-IS IPv4 route summarization allows users to create aggregate IPv4 addresses that include multiple groups of IPv4 addresses for a given IS-IS level. IPv4 Routes redistributed from other routing protocols also can be summarized. It is similar to the OSPF area-range command. IS-IS IPv4 route summarization helps to reduce the size of the LSDB and the IPv4 routing table, and it also helps to reduce the chance of route flapping.

IPv4 route summarization supports:

- Level 1, Level 1-2, and Level 2
- Route summarization for the IPv4 routes redistributed from other protocols
- Metric used to advertise the summary address is the smallest metric of all the more specific IPv4 routes.

### 4.1.4.1 Partial SPF Calculation

IS-IS supports partial SPF calculation, also referred to as partial route calculation. When an event does not change the topology of the network, IS-IS is not perform full SPF but instead performs an IP reach calculation for the impacted routes. Partial SPF is performed at the receipt of IS-IS LSPs with changes to IP reach TLVs and in general, for any IS-IS LSP TLV and sub-TLV change that does not impact the network topology.

---

## 4.1.5 IS-IS MT-Topology Support

Multi-Topology IS-IS (MT-ISIS) support within SR OS allows for the creation of different topologies within IS-IS that contribute routes to specific route tables for IPv4 unicast, IPv6 unicast, IPv4 multicast, and IPv6 multicast. This capability allows for non-congruent topologies between these different routing tables. As a result, networks are able to control which links or nodes are to be used for forwarding different types of traffic.

For example, MT-ISIS could allow all links to carry IPv4 traffic, while only a subset of links can also carry IPv6 traffic.

SR OS supports the following Multi-Topologies:

- IPv4 Unicast – MT-ID 0
- IPv6 Unicast – MT-ID 2
- IPv4 Multicast – MT-ID 3
- IPv6 Multicast – MT-ID 4

### 4.1.5.1 Native IPv6 Support

IS-IS IPv6 TLVs for IPV6 routing is supported in SR OS. This support is considered native IPv6 routing within IS-IS. However, it has limitations in that IPv4 and IPv6 topologies must be congruent, otherwise traffic may be blackholed. Service providers should ensure that the IPv4 topology and IPv6 topologies are the same if native IPv6 routing is used within IS-IS.

## 4.1.6 IS-IS Administrative Tags

IS-IS admin tags enable a network administrator to configure route tags to tag IS-IS route prefixes. These tags can subsequently be used to control Intermediate System-to-Intermediate System (IS-IS) route redistribution or route leaking.

The IS-IS support for route tags allows the tagging of IP addresses of an interface and use the tag to apply administrative policy with a route map. A network administrator can also tag a summary route and then use a route policy to match the tag and set one or more attributes for the route.

Using these administrative policies allow the operator to control how a router handles the routes it receives from and sends to its IS-IS neighboring routers. Administrative policies are also used to govern the installation of routes in the routing table.

Route tags allow:

- Policies to redistribute routes received from other protocols in the routing table to IS-IS.
- Policies to redistribute routes between levels in an IS-IS routing hierarchy.
- Policies to summarize routes redistributed into IS-IS or within IS-IS by creating aggregate (summary) addresses.

#### 4.1.6.1 Setting Route Tags

IS-IS route tags are configurable in the following ways:

- Setting a route tag for an IS-IS interface.
- Setting a route tag on an IS-IS passive interface.
- Setting a route tag for a route redistributed from another protocol to IS-IS.
- Setting a route tag for a route redistributed from one IS-IS level to another IS-IS level.
- Setting a route tag for an IS-IS default route.
- Setting a route tag for an IS-IS summary address.

#### 4.1.6.2 Using Route Tags

Although an operator on this or on a neighboring IS-IS router has configured setting of the IS-IS administrative tags, it will not have any effect unless policies are configure to instruct how to process the given tag value.

Policies can process tags where IS-IS is either the origin, destination or both origin and destination protocol.

```
config>router>policy-options>policy-statement>entry>from
config>router>policy-options>policy-statement>entry>action tag tag-value
config>router>policy-options>policy-statement# default-action tag tag-value
```

---

### 4.1.6.3 Unnumbered Interface Support

IS-IS supports unnumbered point-to-point interface with both Ethernet and PPP encapsulations.

Unnumbered interfaces borrow the address from other interfaces such as system or loopback interfaces and uses it as the source IP address for packets originated from the interface. This feature supports both dynamic and static ARP for unnumbered interfaces to allow interworking with unnumbered interfaces that may not support dynamic ARP.

An unnumbered interface is an IPv4 capability only used in cases where IPv4 is active (IPv4-only and mixed IPv4/IPv6 environments). When configuring an unnumbered interface, the interface specified for the unnumbered interface (system or other) must have an IPv4 address. Also, the interface type for the unnumbered interface will automatically be point-to-point. The unnumbered option can be used in IES and VPRN access interfaces, as well as in a network interface with MPLS support.

### 4.1.7 Segment Routing in Shortest Path Forwarding

Segment routing adds to IS-IS and OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface/next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as Segment ID (SID).

When segment routing is used together with MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels. This is the scope of the features described in this section.

Segment routing using MPLS labels can be used in both shortest path routing applications and in traffic engineering applications. This section focuses on the shortest path forwarding applications.

When a received IPv4 or IPv6 prefix SID is resolved, the Segment Routing module programs the ILM with a swap operation and also the LTN with a push operation both pointing to the primary/LFA NHLFE. An IPv4 or IPv6 SR tunnel to the prefix destination is also added to the TTM and is available for use by shortcut applications and L2/L3 services.

Segment routing introduces the remote LFA feature which expands the coverage of the LFA by computing and automatically programming SR tunnels which are used as backup next-hops. The SR shortcut tunnels terminate on a remote alternate node which provides loop-free forwarding for packets of the resolved prefixes. When the **loopfree-alternate** option is enabled in an IS-IS or OSPF instance, SR tunnels are protected with an LFA backup next-hop. If the prefix of a given SR tunnel is not protected by the base LFA, the remote LFA automatically computes a backup next-hop using an SR tunnel if the **remote-lfa** option is also enabled in the IGP instance.

#### 4.1.7.1 Configuring Segment Routing in Shortest Path

The user enables segment routing in an IGP routing instance using the following sequence of commands.

First, the user configures the global label block, referred to as Segment Routing Global Block (SRGB), which is reserved for assigning labels to segment routing prefix SIDs originated by this router. This range is carved from the system dynamic label range and is not instantiated by default:

**CLI Syntax:** `configure>router>mpls-labels>sr-labels start start-value  
end end-value`

Next, the user enables the context to configure segment routing parameters within a given IGP instance:

**CLI Syntax:** `configure>router>isis>segment-routing  
configure>router>ospf>segment-routing`


The key parameter is the configuration of the prefix SID index range and the offset label value which this IGP instance uses. Because each prefix SID represents a network global IP address, the SID index for a prefix must be unique network-wide. Thus, all routers in the network are expected to configure and advertise the same prefix SID index range for a given IGP instance. However, the label value used by each router to represent this prefix, that is, the label programmed in the ILM, can be local to that router by the use of an offset label, referred to as a start label:

*Local Label (Prefix SID) = start-label + {SID index}*

The label operation in the network becomes thus very similar to LDP when operating in the independent label distribution mode (RFC 5036) with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on advertised prefix SID index using the above formula.

Figure 16 shows an example of a router advertising its loopback address and the resulting packet label encapsulation throughout the network.

**Figure 16 Packet Label Encapsulation using Segment Routing Tunnel**



al\_0636

Router N-6 advertises loopback 10.10.10.1/32 with a prefix index of 5. Routers N-1 to N-6 are configured with the same SID index range of [1,100] and an offset label of 100 to 600 respectively. The following are the actual label values programmed by each router for the prefix of PE2:

- N-6 has a start label value of 600 and programs an ILM with label 605.
- N-3 has a start label of 300 and swaps incoming label 305 to label 605.
- N-2 has a start label of 200 and swaps incoming label 205 to label 305.

Similar operations are performed by N-4 and N-5 for the bottom path.

N-1 has an SR tunnel to N-6 with two ECMP paths. It pushes label 205 when forwarding an IP or service packet to N-6 via downstream next-hop N-2 and pushes label 405 when forwarding via downstream next-hop N-4.

The CLI for configuring the prefix SID index range and offset label value for a given IGP instance is as follows:

**CLI Syntax:**

```
configure>router>isis>segment-routing>prefix-sid-range
{global | start-label label-value max-index index-
value}
```



```
configure>router>ospf>segment-routing>prefix-sid-range
{global | start-label label-value max-index index-
value}
```

There are two mutually-exclusive modes of operation for the prefix SID range on the router. In the global mode of operation, the user configures the global value and this IGP instance assumes the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. Once one IGP instance selected the **global** option for the prefix SID range, all IGP instances on the system is restricted to do the same.

The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. Once the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change fails if an already allocated SID index/label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user thus configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values (start-label + index) must be within the SRGB or the configuration fails. Furthermore, the code checks for overlaps of the resulting net label value range across IGP instances and strictly enforces that these ranges do not overlap.

The user must shutdown the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. In addition, any range change fails if an already allocated SID index/label goes out of range.

The user can, however, change the SRGB on the fly as long as it does not reduce the current per-IGP instance SID index/label range defined with the **prefix-sid-range**. Otherwise, the user must shutdown the segment routing context of the IGP instance and delete and re-configure the **prefix-sid-range** command.

Finally, the user brings up segment routing on that IGP instances by un-shutting the context:

**CLI Syntax:**

```
configure>router>isis>segment-routing>no shutdown
configure>router>ospf>segment-routing>no shutdown
```

This command fails if the user has not previously enabled the **router-capability** option in the IGP instance. Segment routing is a new capability and needs to be advertised to all routers in a given domain so that routers which support the capability only programs the node SID in the data path towards neighbors which support it.

**CLI Syntax:**

```
configure>router>isis>advertise-router-capability {area
| as}
```

```
configure>router>ospf>advertise-router-capability {link
| area | as}
```

The IGP segment routing extensions are area-scoped. As a consequence, the user must configure the flooding scope to **area** in OSPF and to **area** or **as** in IS-IS, otherwise performing **no shutdown** of the segment-routing node fail.

The **segment-routing** command is mutually exclusive with the **igp-shortcut** and **advertise-tunnel-link** options under IGP, because an SR tunnel cannot resolve to an RSVP tunnel next-hop.

Next, the user assigns a node SID index or label to the prefix representing the primary address of an IPv4 or IPv6 network interface of type **loopback** using one of the following commands:

**CLI Syntax:**

```
config>router>isis>interface>ipv4-node-sid index value
config>router>ospf>interface>node-sid index value
config>router>isis>interface>ipv4-node-sid label value
config>router>ospf>interface>node-sid label value
config>router>isis>interface>ipv6-node-sid index value
config>router>isis>interface>ipv6-node-sid label value
```

Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address. The same applies to the non-primary IPv6 addresses of an interface.

Above commands should fail if the network interface is not of type loopback or if the interface is defined in an IES or a VPRN context. Also, assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.

Also, for OSPF the protocol version number and the instance number dictates if the node-SID index/label is for an IPv4 or IPv6 address of the interface. Specifically, the support of address families in OSPF is as follows:

- ospfv2: always IPv4 only

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required because the index, and thus the label ranges, of the various IGP instances are not allowed to overlap.

For an individual adjacency, values for the label may be provisioned for an IS-IS or OSPF interface. If they are not provisioned, then they will be dynamically allocated by the system from the dynamic label range. The following CLI commands are used:

```
config>router>isis>interface
[no] ipv4-adjacency-sid label <value>
[no] ipv6-adjacency-sid label <value>

config>router>ospf>area>interface
[no] adjacency-sid label <value>
```

The *value* must correspond to a label in a reserved label block in provisioned mode referred to by the **srlb** command (see [Segment Routing Local Block](#) for more details of SRLBs).

A static label *value* for an adjacency SID is persistent. Therefore, the P-bit of the Flags field in the Adjacency-SID TLV advertised in the IGP is set to 1.

By default, a dynamic adjacency SID is advertised for an interface. However, if a static adjacency SID value is configured, then the dynamic adjacency SID is deleted and only the static adjacency SID used. Changing an adjacency SID from dynamic (for example, **no adjacency-sid**) to static, or vice versa, may result in traffic being dropped as the ILM is reprogrammed.

For a provisioned adjacency SID for an interface, a backup is calculated similar to a regular adjacency SID when **sid-protection** is enabled for that interface.

Provisioned adjacency SIDs are only supported on point-to-point interfaces.

## 4.1.7.2 Segment Routing Operational Procedures

### 4.1.7.2.1 Prefix Advertisement and Resolution

Once segment routing is successfully enabled in the IS-IS or OSPF instance, the router performs the following operations. See [IS-IS Control Protocol Changes](#) for details of all TLVs and sub-TLVs for both IS-IS and OSPF protocols.

1. Advertise the Segment Routing Capability Sub-TLV to routers in all areas/levels of this IGP instance. However, only neighbors with which it established an adjacency interprets the SID/label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.

2. Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. Then the segment routing module programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Assign and advertise automatically an adjacency SID label for each formed adjacency over a network IP interface in the new Adjacency SID sub-TLV. The following points should be considered.
  - Adjacency SID is advertised for both numbered and unnumbered network IP interface.
  - Adjacency SID is not advertised for an IES interface because access interfaces do not support MPLS.
  - The adjacency SID must be unique per instance and per adjacency. Furthermore, ISIS MT=0 can establish an adjacency for both IPv4 and IPv6 address families over the same link and in such a case a different adjacency SID is assigned to each next-hop. However, the existing IS-IS implementation assigns a single Protect-Group ID (PG-ID) to the adjacency and as such when the state machine of a BFD session tracking the IPv4 or IPv6 next-hop times out, an action is triggered for the prefixes of both address families over that adjacency.

The segment routing module programs the incoming label map (ILM) with a swap to an implicit null label operation, for each advertised adjacency SID.

4. Resolve received prefixes and, if a prefix SID sub-TLV exists, the Segment Routing module programs the ILM with a swap operation and an LTN with a push operation, both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM. If a node SID resolves over an IES interface, the data path is not programmed and a trap is raised. Thus, only next-hops of an ECMP set corresponding to network IP interfaces are programmed in data path; next-hops corresponding to IES interfaces are not programmed. If however the user configures the interface as network on one side and IES on the other side, MPLS packets for the SR tunnel received on the access side is dropped.
5. LSA filtering causes SIDs not to be sent in one direction which means some node SIDs is resolved in parts of the network upstream of the advertisement suppression.

When the user enables segment routing in a given IGP instance, the main SPF and LFA SPF are computed normally and the primary next-hop and LFA backup next-hop for a received prefix are added to RTM without the label information advertised in the prefix SID sub-TLV. In all cases, the segment routing (SR) tunnel is not added into RTM.

#### 4.1.7.2.2 Error and Resource Exhaustion Handling


When the prefix corresponding to a node SID is being resolved, the following procedures are followed.

##### Procedure 1: Providing support of multiple topologies for the same destination prefix

The SR OS supports assigning different prefix-SID indexes and labels to the same prefix in different IGP instances. While other routers that receive these prefix SIDs programs a single route into RTM, based on the winning instance ID as per RTM route type preference, the SR OS adds two tunnels to this destination prefix in TTM. This provides for the support of multiple topologies for the same destination prefix.

For example: In two different instances (L2, IS-IS instance 1 and L1, IS-IS instance 2—see [Figure 17](#)), Router D has the same prefix destination, with different SIDs (SIDx and SIDy).

**Figure 17** Programming multiple tunnels to the same destination



Assume the following route type preference in RTM and tunnel type preference in TTM are configured:

- ROUTE\_PREF\_ISIS\_L1\_INTER (RTM) 15
- ROUTE\_PREF\_ISIS\_L2\_INTER (RTM) 18
- ROUTE\_PREF\_ISIS\_TTM 10



**Note:** the TTM tunnel type preference is not used by the SR module. It is put in the TTM and is used by other applications such as VPRN auto-bind and BGP shortcut to select a TTM tunnel.

**Step 1.** Router A performs the following resolution within the single IS-IS instance 1, level 2. All metrics are the same, and ECMP = 2.

- For prefix N, the RTM entry is:
  - prefix N
  - nhop1 = B
  - nhop2 = C
  - preference 18
- For prefix N, the SR tunnel TTM entry is:
  - tunnel-id 1: prefix N-SIDx
  - nhop1 = B
  - nhop2 = C
  - tunl-pref 10

**Step 2.** Add IS-IS instance 2 (Level 1) in the same setup, but in routers A, B, and C only.

- For prefix N, the RTM entry is:
  - prefix N
  - nhop1 = B
  - preference 15

RTM prefers L1 route over L2 route
- For prefix N, there are two SR tunnel entries in TTM:
 

SR entry for L2:

  - tunnel-id 1: prefix N-SIDx
  - nhop1 = B
  - nhop2= C
  - tunl-pref 10

SR entry for L1:

  - tunnel-id 2: prefix N-SIDy

---

**Procedure 2: Resolving received SID indexes or labels to different routes of the same prefix within the same IGP instance**

Two variations of this procedure can occur.

- a. When SR OS does not allow assigning the *same* SID index or label to different routes of the same prefix within the same IGP instance, it resolves only one of them if received from another SR implementation and based on the RTM active route selection.
- b. When SR OS does not allow assigning *different* SID indexes or labels to different routes of the same prefix within the same IGP instance, it resolves only one of them if received from another SR implementation and based on the RTM active route selection.

The selected SID will be used for ECMP resolution to all neighbors. If the route is inter-area and the conflicting SIDs are advertised by different ABRs, ECMP towards all ABRs uses the selected SID.

**Procedure 3: Checking for SID error prior to programming ILM and NHLFE**

If any of the following conditions are true, the router logs a trap and generates a syslog error message and will not program the ILM and NHLFE for the prefix SID.

- Received prefix SID index falls outside of the locally configured SID range.
- One or more resolved ECMP next-hops for a received prefix SID did not advertise SR Capability sub-TLV.
- Received prefix SID index falls outside the advertised SID range of one or more resolved ECMP next-hops.

**Procedure 4: Programming ILM/NHLFE for duplicate prefix-SID indexes/labels for different prefixes**

Two variations of this procedure can occur.

- a. For received duplicate prefix-SID indexes or labels for different prefixes *within the same* IGP instance, the router:
  - programs ILM/NHLFE for the first one
  - logs a trap and a syslog error message
  - does not program the subsequent one in data path
- b. For received duplicate prefix-SID index for different prefixes *across* IGP instances, there are two options.

- 
- In the global SID index range mode of operation, the resulting ILM label values is the same across the IGP instances. The router:
    - programs ILM/NHLFE for the prefix of the winning IGP instance based on the RTM route type preference
    - logs a trap and a syslog error message
    - does not program the subsequent prefix SIDs in data path
  - In per-instance SID index range mode of operation, the resulting ILM label will have different values across the IGP instances. The router programs ILM/NHLFE for each prefix as expected.

#### **Procedure 5: Programming ILM/NHLFE for the same prefix across IGP instances**


The behavior in the case of a global SID index range is illustrated by the IS-IS example in [Figure 18](#).

In global SID index range mode of operation, the resulting ILM label values is the same across the IGP instances. The router programs ILM/NHLFE for the prefix of the winning IGP instance based on the RTM route type preference. The router logs a trap and a syslog error message, and does not program the other prefix SIDs in data path.

In per-instance SID index range mode of operation, the resulting ILM label has different values across the IGP instances. The router programs ILM/NHLFE for each prefix as expected.



**Figure 18 Handling of Same Prefix and SID in different IS-IS Instances**



Assume the following route type preference in RTM and tunnel type preference in TTM are configured:

- ROUTE\_PREF\_ISIS\_L1\_INTER (RTM) 15
- ROUTE\_PREF\_ISIS\_L2\_INTER (RTM) 18
- ROUTE\_PREF\_ISIS\_TTM 10



**Note:** The TTM tunnel type preference is not used by the SR module. It is put in the TTM and is used by other applications such as VPRN auto-bind and BGP shortcut to select a TTM tunnel.

**Step 1.** Router A performs the following resolution within the single IS-IS instance 1, level 2. All metrics are the same, and ECMP = 2.

- For prefix N, the RTM entry is:
  - prefix N
  - nhop1 = B
  - nhop2 = C
  - preference 18
- For prefix N, the SR tunnel TTM entry is:

- tunnel-id 1: prefix N-SIDx
- nhop1 = B
- nhop2 = C
- tunl-pref 10

**Step 2.** Add IS-IS instance 2 (Level 1) in the same setup, but in routers A, B, and E only.

– For prefix N, the RTM entry is:

- prefix N
- nhop1 = B
- preference 15

RTM prefers L1 route over L2 route.

– For prefix N, there is one SR tunnel entry for L2 in TTM:

- tunnel-id 1: prefix N-SIDx
- nhop1 = B
- nhop2 = C
- tunl-pref 10

#### **Procedure 6: Handling ILM resource exhaustion while assigning an SID index/label**

If the system exhausted an ILM resource while assigning an SID index/label to a local loopback interface, then index allocation is failed and an error is returned in CLI. In addition, the router logs a trap and generates a syslog error message.

#### **Procedure 7: Handling ILM/NHLFE/other IOM or CPM resource exhaustion while resolving or programming an SID index/label**

If the system exhausted an ILM, NHLFE, or any other IOM or CPM resource while resolving and programming a received prefix SID or programming a local adjacency SID, then following occurs.

- The IGP instance goes into overload and a trap and syslog error message are generated.
- The segment routing module deletes the tunnel.

The user must manually clear the IGP overload condition after freeing resources. After the IGP is brought back up, it attempts to program at the next SPF all tunnels which previously failed the programming operation.

---

### 4.1.7.3 Segment Routing Tunnel Management

The segment routing module adds to TTM a shortest path SR tunnel entry for each resolved remote node SID prefix and programs the data path with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs. The LFA backup next-hop for a given prefix which was advertised with a node SID will only be computed if the **loopfree-alternate** option is enabled in the IS-IS or OSPF instance. The resulting SR tunnel which is populated in TTM is automatically protected with FRR when an LFA backup next-hop exists for the prefix of the node SID.

With ECMP, a maximum of 32 primary next-hops (NHLFEs) are programmed for the same tunnel destination per IGP instance. ECMP and LFA next-hops are mutually exclusive as per existing implementation.

The default preference for shortest path SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference of the various tunnel types. This includes the preference of both SR tunnels based on shortest path (referred to as SR-ISIS and SR-OSPF).

The global default TTM preference for the tunnel types is as follows:

- ROUTE\_PREF\_RSVP            7
- ROUTE\_PREF\_SR\_TE         8
- ROUTE\_PREF\_LDP            9
- ROUTE\_PREF\_OSPF\_TTM      10
- ROUTE\_PREF\_ISIS\_TTM      11
- ROUTE\_PREF\_BGP\_TTM      12
- ROUTE\_PREF\_GRE            255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of an SR tunnel in this case is based on lowest IGP instance-id.

The TTM preference is used in the case of BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the tunnel binding commands are configured to the **any** value which parses the TTM for tunnels in the protocol preference order. The user can choose to either go with the global TTM preference or list explicitly the tunnel types to be used. When the tunnel types are listed explicitly, the TTM preference is still used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a more preferred tunnel type is performed as soon as one is available. See [BGP Shortcut](#)

Using [Segment Routing Tunnel](#), [BGP Label Route Resolution Using Segment Routing Tunnel](#), and [Service Packet Forwarding with Segment Routing](#) for the detailed service and shortcut binding CLI.

For SR-ISIS and SR-OSPF, the user can configure the preference of each specific IGP instance away from the above default values.

**CLI Syntax:**

```
configure>router>isis>segment-routing>tunnel-table-pref
 preference <1..255>
configure>router>ospf>segment-routing>tunnel-table-pref
 preference <1..255>
```



**Note:** The preference of SR-TE LSP is not configurable and is the second most preferred tunnel type after RSVP-TE. This is independent if the SR-TE LSP was resolved in IS-IS or OSPF.

#### 4.1.7.3.1 Tunnel MTU Determination

The MTU of an SR tunnel populated into TTM is determined as in the case of an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Segment routing, however, supports remote LFA which programs an LFA backup next-hop adding another label to the tunnel for a total of two labels.

Based on the above, the user is provided with a CLI to configure the MTU of all SR tunnels within each IGP instance:

**CLI Syntax:**

```
configure>router>isis (ospf)>segment-routing>tunnel-mtu
 bytes
```

There is no default value for this new command. If the user does not configure an SR tunnel MTU, the MTU is fully determined by IGP as explained below.

The MTU of the SR tunnel, in bytes, is then determined as follows:

$$SR\_Tunnel\_MTU = MIN \{ Cfg\_SR\_MTU, IGP\_Tunnel\_MTU - (1 + frr-overhead) * 4 \}$$

Where:

- *Cfg\_SR\_MTU* is the MTU configured by the user for all SR tunnels within a given IGP instance using the above CLI. If no value was configured by the user, the SR tunnel MTU is fully determined by the IGP interface calculation explained next.

- *IGP\_Tunnel\_MTU* is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- *frr-overhead* is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated anytime any of the above parameters used in its calculation changes. This includes when the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.



**Note:** For the purpose of fragmentation of IP packets forwarded in GRT or in a VPRN over an SR shortest path tunnel, the IOM always deducts the worst case MTU (5 labels or 6 labels if hash label feature is enabled) from the outgoing interface MTU for the decision to fragment or not the packet. In this case, the above formula is not used.

#### 4.1.7.4 Segment Routing Local Block

Some labels that are provisioned through CLI or a management interface must be allocated from the Segment Routing Local Block (SRLB). The SRLB references a reserved label block configured under **configure>router>mpls-labels**. Refer to the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for further information on reserved label blocks.

The label block to use is specified by the **srlb** command under IS-IS or OSPF:

```
config>router>isis>segment-routing
[no] srlb <reserved-label-block-name>

config>router>ospf> segment-routing
[no] srlb <reserved-label-block-name>
```

Provisioned labels for adjacency SIDs and adjacency SID sets must be allocated from the configured SRLB. If no SRLB is specified, or the requested label does not fall within the SRLB, or the label is already allocated, then the request will be rejected.

##### 4.1.7.4.1 Bundling Adjacencies in Adjacency Sets

An adjacency set is a bundle of adjacencies, represented by a common adjacency SID for the bundled set. It enables, for example, a path for an SR-TE LSP through a network to be specified while allowing the local node to spray packets across the set of links identified by a single adjacency SID.

SR OS supports both parallel adjacency sets (for example, those where adjacencies originating on one node terminate on a second, common node), and the ability to associate multiple interfaces on a specified node, irrespective of whether the far end of those interfaces' respective links terminate on the same node.

An adjacency set is created under IS-IS or OSPF using the following CLI:

```
config
 router
 isis | ospf
 segment-routing
 [no] adjacency-set <id>
 family [ipv4 | ipv6]
 parallel [no-advertise]
 no parallel
 exit
 ...
 exit
 exit
config
 router
 ospf
 segment-routing
 [no] adjacency-set <id>
 parallel [no-advertise]
 no parallel
 exit
 ...
 exit
```

The **adjacency-set** *id* command specifies an adjacency set, where *id* is an unsigned integer in the range 0 to 4294967295.

In IS-IS, each adjacency set is assigned an address family, IPv4 or IPv6. The family command for IS-IS indicates the address family of the adjacency set. For OSPF, the address family of the adjacency set is implied by the OSPF version and the instance.

The **parallel** command indicates that all members of the adjacency set must terminate on the same neighboring node. The system raises a trap, when the **parallel** command is configured, if a user attempts to add an adjacency terminating on a neighboring node that differs from the existing members of the adjacency set. See [Associating an Interface with an Adjacency Set](#) for details about how to add interfaces to an adjacency set. In addition, the system stops advertising the adjacency set and deprograms it from TTM. The **parallel** command is enabled by default.

By default, parallel adjacency sets are advertised in the IGP. The **no-advertise** option prevents a parallel adjacency set from being advertised in the IGP; it is only advertised if the **parallel** command is configured. To prevent issues in the case of ECMP if a non-parallel adjacency set is used, coordination may be required by an external controller of the label sets for SIDs at all downstream nodes. As a result, non-parallel adjacency sets are not advertised in the IGP. The label stack below the adjacency set label must be valid at any downstream node that exposes it, even though it is sprayed over multiple downstream next-hops.

Parallel adjacency sets are programmed in TTM (unless there is an erroneous configuration of a non-parallel adjacency, as described). Non-parallel adjacency sets are not added to TTM or RTM, meaning that they cannot be used as a hop at the originating node. Parallel adjacency sets that are advertised are included in the link state database and TE database, but non-parallel adjacency sets are not included because they are not advertised.

An adjacency set with only one next hop is also advertised as an individual adjacency SID with the S flag set. However, the system does not calculate a backup for an adjacency set even if it has only one next hop.

## Associating an Interface with an Adjacency Set

IS-IS or OSPF interfaces are associated with one or more adjacency sets using the following CLI. Both numbered and unnumbered interfaces can be assigned to the same adjacency set.

```
config
 router
 isis
 interface
 [no] adjacency-set <id>
 [no] adjacency-set <id>
 [no] adjacency-set <id>
config
 router
 ospf
 area
 interface
 [no] adjacency-set <id>
 [no] adjacency-set <id>
 [no] adjacency-set <id>
```

If an interface is assigned to an adjacency set, then a common adjacency SID value is advertised for every interface in the set, in addition to the adjacency SID corresponding to the IPv4 and or IPv6 adjacency for the interface. Each IS-IS or OSPF advertisement therefore contains two adjacency SID TLVs for an address family:

- an adjacency SID for the interface (a locally-unique value).
- an adjacency SID TLV for the adjacency set.

This TLV is distinguished by having the S-bit (ISIS) or G-bit (OSPF) in the flags field set to 1. Its value is the same as other adjacency SIDs in the set at that node.

By default, both the adjacency SID for an interface and the adjacency SID for a set are dynamically allocated by the system. However, it is possible for the user to configure an alternative, static value for the SID (see [Provisioning Adjacency SID Values for an Adjacency Set](#) for more information).

A maximum of 32 interfaces can be bound to a common adjacency set. Configuration of more than 32 interfaces is blocked by the system and a CLI error will be generated.

Only point-to-point interfaces can be assigned to an adjacency set.

If a user attempts to assign an IES interface to an adjacency set, the system generates a CLI warning and segment routing does not program the association.

The IGP blocks the configuration of an adjacency set under an interface when the adjacency set has not yet been created under segment-routing.

In IS-IS, it is possible to add L1, L2, or a mix of L1 and L2 adjacencies to the same adjacency set.

### Provisioning Adjacency SID Values for an Adjacency Set

For an adjacency set, static values are configured using the **sid** CLI command, as follows:

```
config>router>isis>segment-routing
 [no] adjacency-set <id>
 family [ipv4 | ipv6]
 [no] sid label <value>
 parallel [no-advertise]
 no parallel
 exit
 [no] adjacency-set <id>
 family [ipv4 | ipv6]
 [no] sid label <value>
 parallel [no-advertise]
 no parallel
 exit
 ...

config>router>ospf>segment-routing
 [no] adjacency-set <id>
 [no] sid label <value>
 parallel [no-advertise]
```



```

 no parallel
 exit
 [no] adjacency-set <id>
 [no] sid label <value>
 parallel [no-advertise]
 no parallel
 exit
 ...

```

If **no sid** is configured, a dynamic value is allocated to the adjacency set. A user may change the dynamic value to specify a static SID value. Changing an adjacency set value from dynamic to a static, or vice versa, may result in traffic being dropped as the ILM is reprogrammed.

The *value* must correspond to a label in the reserved label block in provisioned mode referred to by the **srld** command. A CLI error is generated if a user attempts to configure an invalid *value*. If a label is not configured, then the label *value* is dynamically allocated by the system from the dynamic labels range. If a static adjacency set label is configured, then the system does not advertise a dynamic adjacency set label.

A static label value for an adjacency set SID is persistent. Therefore, the P-bit of the Flags field in the Adjacency-SID TLV, referring to the adjacency set should be set to 1.

#### 4.1.7.5 Entropy Label for IS-IS Segment Routing

The router supports the MPLS entropy label, as specified in RFC 6790, on IS-IS segment-routed tunnels. LSR nodes in a network can load-balance labeled packets in a more granular way than by hashing on the standard label stack. Refer to the *MPLS Guide* for more information.

#### 4.1.7.6 Remote LFA with Segment Routing

The user enables the remote LFA next-hop calculation by the IGP LFA SPF by appending the following new option in the existing command which enables LFA calculation:

**CLI Syntax:**

```

configure>router>isis>loopfree-alternate remote-lfa
configure>router>ospf>loopfree-alternate remote-lfa

```


SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when both of the following conditions are met:

- The remote-lfa option is enabled in an IGP instance.
- The LFA next hop calculation did not result in protection for one or more prefixes resolved to a given interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing/tearing-down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node which puts the packets back into the shortest without looping them back to the node which forwarded them over the repair tunnel. A repair tunnel can in theory be an RSVP LSP, an LDP-in-LDP tunnel, or an SR tunnel. In SR OS, this feature is restricted to use an SR repair tunnel to the remote LFA node.

The remote LFA algorithm for link protection is described in RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*. Unlike the regular LFA calculation, which is calculated per prefix, the LFA algorithm for link protection is a per-link LFA SPF calculation. As such, it provides protection for all destination prefixes which share the protected link by using the neighbor on the other side of the protected link as a proxy for all these destinations. Assume the topology in [Figure 19](#).

**Figure 19 Remote LFA Algorithm**



When the LFA SPF in node C computes the per-prefix LFA next-hop, prefixes which use link C-B as the primary next-hop has no LFA next-hop due to the ring topology. If node C used node link C-D as a back-up next-hop, node D would loop a packet back to node C. The remote LFA then runs the following algorithm, referred to as the “PQ Algorithm” in RFC 7490:

1. Compute the extended P space of Node C with respect to link C-B: set of nodes reachable from node C without any path transiting the protected link (link C-B). This yields nodes D, E, and F.

The determination of the extended P space by node C uses the same computation as the regular LFA by running SPF on behalf of each of the neighbors of C.



**Note:** RFC 7490 initially introduced the concept of P space, which would have excluded node F because, from the node C perspective, node C has a couple of ECMP paths, one of which goes via link C-B. However, because the remote LFA next-hop is activated when link C-B fails, this rule can be relaxed and node F can be included, which then yields the extended P space.

The user can limit the search for candidate P nodes to reduce the amount of SPF calculations in topologies where many eligible P nodes can exist. A CLI command is provided to configure the maximum IGP cost from node C for a P node to be eligible:

- **configure>router>isis>loopfree-alternate remote-lfa max-pq-cost**  
*value*
- **configure>router>ospf>loopfree-alternate remote-lfa max-pq-cost**  
*value*

2. Compute the Q space of node B with respect to link C-B: set of nodes from which the destination proxy (node B) can be reached without any path transiting the protected link (link C-B).


The Q space calculation is effectively a reverse SPF on node B. In general, one reverse SPF is run on behalf of each of C neighbors to protect all destinations resolving over the link to the neighbor. This yields nodes F and A in the example of [Figure 19](#).

The user can limit the search for candidate Q nodes to reduce the amount of SPF calculations in topologies where many eligible Q nodes can exist. The CLI above is also used to configure the maximum IGP cost from node C for a Q node to be eligible.

3. Select the best alternate node: this is the intersection of extended P and Q spaces. The best alternate node or PQ node is node F in the example of [Figure 19](#). From node F onwards, traffic follows the IGP shortest path.

If many PQ nodes exist, the lowest IGP cost from node C is used to narrow down the selection and if more than one PQ node remains, the node with lowest router-id is selected.

The details of the label stack encoding when the packet is forwarded over the remote LFA next-hop is shown in [Figure 20](#).

**Figure 20 Remote LFA Next-Hop in Segment Routing**

al\_0648

The label corresponding to the node SID of the PQ node is pushed on top of the original label of the SID of the resolved destination prefix. If node C has resolved multiple node SIDs corresponding to different prefixes of the selected PQ node, it pushes the lowest node SID label on the packet when forwarded over the remote LFA backup next-hop.

If the PQ node is also the advertising router for the resolved prefix, the label stack is compressed in some cases depending on the IGP:

- In IS-IS, the label stack is always reduced to a single label, which is the label of the resolved prefix owned by the PQ node.
- In OSPF, the label stack is reduced to the single label of the resolved prefix when the PQ node advertised a single node SID in this OSPF instance. If the PQ node advertised a node SID for multiple of its loopback interfaces within this same OSPF instance, the label stack is reduced to a single label only in the case where the SID of the resolved prefix is the lowest SID value.

---

The following rules and limitations apply to the remote LFA implementation:

- LFA policy is currently supported for IP next-hops only. It is not supported with tunnel next-hops when IGP shortcuts are used for LFA backup. Remote LFA is also a tunnel next-hop and, as such, a user configured LFA policy is not applied in the selection of a remote LFA backup next-hop when multiple candidates are available.
- As a result, if an LFA policy is applied and does not find an LFA IP next-hop for a set of prefixes, the remote LFA SPF runs to search for a remote LFA next-hop for the same prefixes. The selected remote LFA next-hops, if found, may not satisfy the LFA policy constraints.
- If the user excludes a network IP interface from being used as an LFA next-hop using the CLI command **loopfree-alternate-exclude** under the interface's IS-IS or OSPF context, the interface is also excluded from being used as the outgoing interface for a remote LFA tunnel next-hop.
- As with the regular LFA algorithm, the remote LFA algorithm computes a backup next-hop to the ABR advertising an inter-area prefix and not to the destination prefix itself.

#### 4.1.7.7 Topology Independent LFA

The Topology-Independent LFA (TI-LFA) feature improves the protection coverage of a network topology by computing and automatically instantiating a repair tunnel to a Q node which is not in the shortest path from the computing node. The repair tunnel uses the shortest path to the P node and a source-routed path from the P node to the Q node.

In addition, the TI-LFA algorithm selects the backup path that matches the post-convergence path. This helps capacity planning in the network since traffic always flows on the same path when transitioning to the FRR next-hop and then onto the new primary next-hop.

At a high level, the TI-LFA link protection algorithm searches for the closest Q node to the computing node and then selects the closest P node to this Q node, up to a maximum number of labels. This is performed on each of the post-convergence paths to each destination node or prefix D.

When the TI-LFA feature is enabled in IS-IS, it provides a TI-LFA link-protect backup path in IS-IS MT=0 for a SR-ISIS IPV4/IPV6 tunnel (node SID and adjacency SID), for a IPV4 SR-TE LSP, and for LDP IPV4 FEC when the LDP **fast-reroute backup-sr-tunnel** option is enabled.

#### 4.1.7.7.1 TI-LFA Configuration

Users can enable TI-LFA in an IS-IS instance using one of the following command:

```
config>router>isis>loopfree-alternate [remote-lfa [max-pq-cost value]] [ti-lfa
[max-sr-frr-labels value]]
```

When the **ti-lfa** option is enabled in IS-IS, it provides a TI-LFA link-protect backup path in IS-IS MT=0 for an SR-ISIS IPV4 and IPV6 tunnel (node SID and adjacency SID), for a IPv4 SR-TE LSP, and for an LDP IPv4 FEC when the LDP **fast-reroute backup-sr-tunnel** option is enabled. See more details of the applicability of the various LFA options in [LFA Protection Option Applicability](#).

The **max-sr-frr-labels** parameter limits the search for the LFA backup next-hop:

- 0 — The IGP LFA SPF restricts the search to TI-LFA backup next-hop which does not require a repair tunnel, meaning that P node and Q node are the same and match a neighbor. This is also the case when both P and Q node match the advertising router for a prefix.
- 1 to 3 — The IGP LFA SPF widens the search to include a repair tunnel to a P node which itself is connected to the Q nodes with a 0-to-2 hops for a total of maximum of three labels: one node SID to P node and two adjacency SIDs from P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed and meaning that up to three adjacency SIDs can separate the P and Q nodes.
- 2 (default) — Corresponds to a repair tunnel to a non-adjacent P which is adjacent to the Q node. If the P node is a neighbor of the computing node, then the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

When the user attempts to change the **max-sr-frr-labels** parameter to a value that results in a change to the computed FRR overhead, then IGP checks that all SR-TE LSPs can properly account for the overhead based on the configuration of the LSP **max-sr-labels** and **additional-frr-labels** parameter values or the change is rejected.

The FRR overhead is computed by IGP and its value is set as follows:

- 0 if **segment-routing** is disabled in the IGP instance
- 0 if **segment-routing** is enabled but **remote-lfa** is disabled and **ti-lfa** is disabled
- 1 if **segment routing** is enabled and **remote-lfa** is enabled but **ti-lfa** is disabled, or if **segment-routing** is enabled and **remote-lfa** is enabled and **ti-lfa** is enabled but **ti-lfa max-sr-frr-labels labels** is set to 0.
- The value of **ti-lfa max-sr-frr-labels labels**, if **segment-routing** is enabled and **ti-lfa** is enabled regardless if **remote-lfa** is enabled or disabled.

The LFA commands mentioned above enables the base LFA feature with the **loopfree-alternate** command and to optionally add remote LFA with the **remote-lfa** option and TI-LFA with the **ti-lfa** option. The behavior when one or more of these options is enabled is explained in [TI-LFA Link-Protect Operation](#).

#### 4.1.7.7.2 TI-LFA Link-Protect Operation

##### LFA Protection Option Applicability

Depending on the configured options of the **loopfree-alternate** command, the LFA SPF in a IGP instance runs the various algorithms in the following order:

1. First compute a regular LFA for each node and prefix. In this step, a computed backup next-hop satisfies any applied LFA policy. This backup next-hop protects that specific prefix or node in the context of IP FRR, LDP FRR, SR FRR, and SR-TE FRR.

This backup next-hop protects that specific prefix or node in the context of IP FRR, LDP FRR, SR FRR, and SR-TE FRR.

2. Next, always follow with the TI-LFA if the **ti-lfa** option enabled for all prefixes and nodes regardless of the outcome of the first step.

A prefix or node for which a TI-LFA backup next-hop is found overrides the result from the first step in the context of LDP FRR, if the LDP **fast-reroute backup-sr-tunnel** option is enabled, in SR FRR and in SR-TE FRR.

With SR FRR and SR-TE FRR, the TI-LFA next-hop protects the node-SID of that prefix and for protecting any adjacency-SID terminating on the node-SID of that prefix.

The prefix or node continues the use backup next-hop found in (a) in LDP FRR, if the LDP **fast-reroute backup-sr-tunnel** option disabled, and in IP FRR.

3. Finally, run remote LFA only for the next-hop of prefixes and nodes which remain unprotected after the first and second steps if the **remote-lfa** option is enabled. A prefix or node for which a remote LFA backup next-hop is found uses it in the context of LDP FRR, when the LDP **fast-reroute backup-sr-tunnel** option is enabled, in SR FRR and in SR-TE FRR.

When protecting an adjacency SID, a parallel ECMP adjacency takes precedence over any other type of LFA backup path. Applying the above algorithm applicability rules result in the following selection:


- adjacency SID of an alternate ECMP next-hop
- TI-LFA backup next-hop

- LFA backup next-hop
- remote LFA backup next-hop

### TI-LFA Algorithm

At a high level, the TI-LFA link protection algorithm searches for the closest Q node to the computing node and then selects the closest P node to this Q node, up to a number of labels corresponding to the value of `ti-lfa max-sr-frr-labels labels`, on each of the post-convergence paths to each destination node or prefix D. Consider the topology in [Figure 21](#) where router R3 computes a TI-LFA next-hop for protecting link R3-R4.

**Figure 21** Selecting Link-Protect TI-LFA Backup Path



sw0043

For each destination node D:

1. Compute the post-convergence SPF on the topology without the protected link.  
In [Figure 21](#), R3 finds a single post-convergence path to destination D via R1.  
Note that the post-convergence SPF does not include IGP shortcut tunnels, unless advertised as forwarding adjacencies.
2. Compute the extended P-Space of R3 with respect to protected link R3-R4 on the post-convergence paths.



This is the set of nodes  $Y_i$  in the post-convergence paths which are reachable from R3 neighbors without any path transiting the protected link R3-R4.

R3 computes an LFA SPF rooted at each of its neighbors within the post-convergence paths, that is, R1, using the following equation:

$$\text{Distance\_opt}(R1, Y_i) < \text{Distance\_opt}(R1, R3) + \text{Distance\_opt}(R3, Y_i)$$

Where,  $\text{Distance\_opt}(A,B)$  is the shortest distance between A and B. The extended P-space calculation yields only node R1.

3. Compute Q-space of R3 with respect to protected link R3-R4 in the post-convergence paths.

This is the set of nodes  $Z_i$  in the post-convergence paths from which the neighbor node R4 of the protected link, acting as a proxy for all destinations D, can be reached without any path transiting the protected link R3-R4.

$$\text{Distance\_opt}(Z_i, R4) < \text{Distance\_opt}(Z_i, R3) + \text{Distance\_opt}(R3, R4)$$

The Q-space calculation yields nodes R2 and R4.

This is the same computation of the Q-space performed by the remote LFA algorithm, except that the TI-LFA Q-space computation is performed only on the post-convergence.

4. For each post-convergence path, search for the closest Q-node and select the closest P-node to this Q-node, up to a number of labels corresponding to the value of **ti-lfa max-sr-frr-labels** *labels*.

In the topology in [Figure 21](#), there is a single post-convergence path, a single P-node (R1), and the closest of the two found Q-nodes to the P-Node is R2.

R3 installs the repair tunnel to the P-Q set and includes the node-SID of R1 and the adjacency SID of the adjacency over link R1-R2 in the label stack. Note that since the P-node R1 is a neighbor of the computing node R3, the node SID of R1 is not needed and the label stack of the repair tunnel is compressed to the adjacency SID over link R1-R2 as shown in [Figure 21](#).

When a P-Q set is found on multiple ECMP post-convergence paths, the following selection rules are applied, in ascending order, to select a set from a single path:

- a. the lowest number of labels
- b. the lowest next hop router ID
- c. the lowest interface index if same next-hop router-id

If multiple links with adjacency SID exist between the selected P node and the selected Q node, the following rules are used to select one of them:

- a. the adjacency SID with the lowest metric
- b. the adjacency SID with the lowest SID value if same lowest metric


---

## TI-LFA Feature Interaction and Limitations

The following are feature interactions and limitations of the TI-LFA link protection.


- Enabling the **ti-lfa** option in an IS-IS or OSPF instance overrides the user configuration of the **loopfree-alternate-exclude** command under the interface's context in that IGP instance. In other words, the TI-LFA SPF uses that interface as a backup next-hop if it matches the post-convergence next-hop.
- Any prefix excluded from LFA protection using the **loopfree-alternate-exclude prefix-policy prefix-policy** command under the IGP instance context is also excluded from TI-LFA.
- Since the post-convergence SPF does not use paths transiting on a node in IS-IS overload, the TI-LFA backup path automatically will not transit on such a node.
- As with remote LFA, a user-configured LFA policy is not applied in the selection of a TI-LFA backup next-hop when multiple candidates are available.
- IES interfaces are skipped in TI-LFA computation since they do not support Segment Routing with MPLS encapsulation. If the only found TI-LFA backup next-hop matches an IES interface, IGP will treat this as if there were no TI-LFA backup paths and will fall back to using either a remote LFA or regular LFA backup path as per the selection rules in [LFA Protection Option Applicability](#).
- The TI-LFA feature provides link-protection only. Thus, if the protected link is a broadcast interface, the TI-LFA algorithm only guarantees protection of that link and not of the Pseudo-Node (PN) corresponding to that shared subnet. In other words, if the PN is in the post-convergence path, the TI-LFA backup path may still traverse again the PN. For example, node E in [Figure 22](#) computes a TI-LFA backup path to destination D via E-C-PN-D because it is the post-convergence path when excluding link E-PN from the topology. This TI-LFA backup does not protect against the failure of the PN.

**Figure 22 TI-LFA Backup Path via a Pseudo-Node**



- When the computing router selects an adjacency SID among a set of parallel adjacencies between the P and Q nodes, the selection rules in step 4 of [TI-LFA Algorithm](#) are used. However, these rules may not yield the same interface the P node itself would have selected in its post-convergence SPF since the latter is based on the lowest value of the locally managed interface index.

For example, node A in [Figure 23](#) computes the link-protect TI-LFA backup path for destination node E as path A-C-E, where C is the P node and E is the Q node and destination. C has a pair of adjacency SIDs with the same metric to E. Node A selects the adjacency over the P2P link C-E because it has the lowest SID value but node C may select the interface C-PN in its post-convergence path calculation if that interface has a lower interface index than P2P link C-E.

**Figure 23 Parallel Adjacencies between P and Q Nodes**

- When a node SID is advertised by multiple routers (anycast SID), the TI-LFA algorithm on a router which resolves the prefix of this SID computes the backup next-hop toward a single node owner of the prefix based on the rules for prefix and SID ECMP next-hop selection.

#### 4.1.7.7.3 Data Path Support

The TI-LFA repair tunnel can have a maximum of three additional labels pushed in addition to the label of the destination node or prefix. The user can set a lower maximum value for the additional FRR labels by configuring the **ti-lfa max-sr-frr-labels labels** CLI option. The default value is 2.

The data path models the backup path like a SR-TE LSP and thus uses a super-NHLFE pointing to the NHLFE of the first hop in the repair tunnel. That first hop corresponds to either an adjacency SID or a node SID of the P node.

There is the special case where the P node is adjacent to the node computing the TI-LFA backup, and the Q node is the same as the P node or adjacent to the P node. In this case, the data path at the computing router pushes either zero labels or one label for the adjacency SID between P and Q nodes. The backup path uses a regular NHLFE in this case like in base LFA or remote LFA features. [Figure 21](#) shows an example of a single label in the backup NHLFE.

## 4.1.7.8 Node Protection Support in TI-LFA and Remote LFA

This feature extends the Remote LFA and TI-LFA features by adding support for node protection. The extensions are additions to the original link-protect LFA SPF algorithm.

When node protection is enabled, the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations.

### 4.1.7.8.1 Feature Configuration

The following are the CLI commands to configure the remote LFA and TI-LFA node protection feature.

```
CLI Syntax: configure
 router
 [no] isis
 [no] loopfree-alternates
 [no] remote-lfa [max-pq-cost
 <0..4294967295,
 default=4261412864>]
 [no] node-protect [max-pq-nodes
 <1..32, default=16>]
 [no] ti-lfa [max-sr-frr-labels <0..3,
 default=2>]
 [no] node-protect
 exclude
 [no] prefix-policy prefix-policy
 [prefix-policy...(up to 5 max)]
 exit
 exit
```

A command is added to enable node-protect calculation to both Remote LFA (**node-protect** [max-pq-nodes <1..32, default=16>]) and TI-LFA (**node-protect**).

When the **node-protect** command is enabled, the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.


The parameter **max-pq-nodes** in Remote LFA controls the maximum number of candidate PQ nodes found in the LFA SPFs for which the node protection check is performed. As explained in [Remote LFA Node-Protect Operation](#), the node-protect condition means the router must run the original link-protect Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel. The default value is 16.

#### 4.1.7.8.2 TI-LFA Node-Protect Operation

The SR OS supports the node-protect extensions to the TI-LFA algorithm as described in *draft-bashandy-rtgwg-segment-routing-ti-lfa-05*.

[Figure 24](#) shows a simple topology to illustrate the operation of the node-protect in the [TI-LFA Algorithm](#).

**Figure 24 Application of the TI-LFA Algorithm for Node Protection**



The first change is that the algorithm has to protect a node instead of a link.

The following topology computations pertain to [Figure 24](#):

- Step 1.** Compute post-convergence SPF on the topology without the protected node.

In [Figure 24](#), R1 computes TI-LFA on the topology without the protected node R2 and finds a single post-convergence path to destination D via R3 and R6.

Prefixes owned by all other nodes in the topology have a post-convergence path via R3 and R6 except for prefixes owned by node R2. The latter will use the link R3-R2 and they can only benefit from link protection.

- Step 2.** Compute extended P-Space of R1 with respect to protected node R2 on the post-convergence paths.

This is the set of nodes  $Y_i$  in the post-convergence paths that are reachable from R1 neighbors, other than protected node R2, without any path transiting the protected node R2.

R1 computes an LFA SPF rooted at each of its neighbors within the post-convergence paths, for example, R3, using the following equation:

$$\text{Distance\_opt}(R3, Y_i) < \text{Distance\_opt}(R3, R2) + \text{Distance\_opt}(R2, Y_i)$$

Where  $\text{Distance\_opt}(A,B)$  is the shortest distance between A and B.

The extended P-space calculation yields node R3 only.

- Step 3.** Compute Q-space of R1 with respect to protected link R1-R2 on the post-convergence paths.

This is the set of nodes  $Z_i$  in the post-convergence paths from which node R2 can be reached without any path transiting the protected link R1-R2.

$$\text{Distance\_opt}(Z_i, R2) < \text{Distance\_opt}(Z_i, R1) + \text{Distance\_opt}(R1, R2)$$

The reverse SPF for the Q-space calculation is the same as in the link-protect algorithm and uses the protected node R2 as the proxy for all destination prefixes. Note that if the Q-space were to be computed with respect to the protected node R2 instead of link R1-R2, a reverse SPF would have to be done to each destination D which is very costly and would not scale. This however means the path from the Q node to the destination D, or the path from P node to Q node, is not guaranteed to avoid the protected router R2. The intersection of the Q-space with post-convergence path is modified in the next step to mitigate this risk.

This step yields nodes R3, R4, R5, and R6.

- Step 4.** For each post-convergence path, search for the closest Q-node to destination D and select the closest P-node to this Q-node, up to a number of labels corresponding to the value of **ti-lfa max-sr-frr-labels labels**.

This step yields the following P-Q sets depending on the value of the parameter **max-sr-frr-labels**:

- **max-sr-frr-labels=0**, R3 is the closest Q node to the destination D and R3 is the only P node. This case is the one which results in link protection via PQ node R3.
- **max-sr-frr-labels=1**, R6 is the closest Q node to the destination D and R3 is the only P node. The repair tunnel for this case uses the SID of the adjacency over link R3-R6 and is illustrated in [Figure 24](#).
- **max-sr-frr-labels=2**, R5 is the closest Q node to the destination D and R3 is the only P node. The repair tunnel for this case uses the SIDs of the adjacencies over links R3-R6 and R6-R5.
- **max-sr-frr-labels=3**, R4 is the closest Q node to the destination D and R3 is the only P node. The repair tunnel for this case uses the SIDs of the adjacencies over links R3-R6, R6-R5, and R5-R4.

Note this step of the algorithm is modified from link protection which prefers Q nodes which are the closest to the computing router R1. This is to minimize the probability that the path from the Q node to the destination D, or the path from P node to Q node, goes via the protected node R2 as explained in step 2. There is however still a probability that the found P-Q set achieves link protection only.

**Step 5.** Select the P-Q Set

If a candidate P-Q set is found on each of the multiple ECMP post-convergence paths in step 4, the following selection rules are applied in ascending order to select a single set:

- a. lowest number of labels
- b. lowest next-hop router-id
- c. lowest interface index if same next-hop router-id

If multiple parallel links with adjacency SID exist between the P and Q nodes of the selected P-Q set, the following rules are used to select one of them:

- a. Adjacency SID with lowest metric
- b. Adjacency SID with the lowest SID value if same lowest metric

For each destination prefix D, R1 programs the TI-LFA repair tunnel (**max-sr-frr-labels=1**):

1. For prefixes other than those owned by node R2, R1 programs a node-protect repair tunnel to the P-Q pair R3-R6 by pushing the SID of adjacency R3-R6 on top of the SID for destination D and programming a next-hop of R3.
2. For prefixes owned by node R2, R1 runs the link-protect TI-LFA algorithm and programs a simple link-protect repair tunnel which consists of a backup next-hop of R3 and pushing no additional label on top of the SID for the destination prefix.



### 4.1.7.8.3 Remote LFA Node-Protect Operation


The SR OS supports the node-protect extensions to the Remote LFA algorithm as described in RFC 8102.

Remote LFA follows a similar algorithm as TI-LFA but does not limit the scope of the calculation of the extended P-Space and of the Q-Space to the post-convergence paths.

Remote LFA adds an extra forward SPF on behalf of the PQ node to make sure that for each destination the selected PQ node does not use a path via the protected node.

Figure 25 shows the same topology as in Figure 24 but with the metric for link R3-R6 modified to 100.

**Figure 25 Application of the Remote LFA Algorithm for Node Protection**



Applying the node protect remote LFA algorithm to this topology yields the following steps:

- Step 1.** Compute extended P-Space of R1 with respect to protected node R2. This is the set of nodes  $Y_i$  which are reachable from R1 neighbors, other than protected node R2, without any path transiting the protected node R2.

R1 computes a LFA SPF rooted at each of its neighbors, for example, R3, using the following equation:

$$\text{Distance\_opt}(\text{R3}, \text{Yi}) < \text{Distance\_opt}(\text{R3}, \text{R2}) + \text{Distance\_opt}(\text{R2}, \text{Yi})$$

Where  $\text{Distance\_opt}(\text{A}, \text{B})$  is the shortest distance between A and B.

Nodes R3, R5 and R6 satisfy this inequality.

**Step 2.** Compute Q-space of R1 with respect to protected link R1-R2.

This is the set of nodes  $Z_i$  from which node R2 can be reached without any path transiting the protected link R1-R2.

$$\text{Distance\_opt}(\text{Zi}, \text{R2}) < \text{Distance\_opt}(\text{Zi}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R2})$$

The reverse SPF for the Q-space calculation is the same as in the remote LFA link-protect algorithm and uses the protected node R2 as the proxy for all destination prefixes.

This step yields nodes R3, R4, R5, and R6.

**Step 3.** For each PQ node found, run a forward SPF to each destination D.

This step is required to select only the subset of PQ nodes which does not traverse protected node R2.

$$\text{Distance\_opt}(\text{PQi}, \text{D}) < \text{Distance\_opt}(\text{PQi}, \text{R2}) + \text{Distance\_opt}(\text{R2}, \text{D})$$

Of the candidates PQ nodes R3, R5, and R6, only PQ nodes R5 and R6 satisfy this inequality. PQ node R6 is closer in terms of IGP cost to computing router R1 and is thus selected. The detailed selection rules of a single PQ node from the candidate list is provided in Step 4.

Note this step of the algorithm is applied to the subset of candidate PQ nodes out of steps 1 and 2 and to which the parameter `max-pq-cost` was already applied. This subset is further reduced in this step by retaining the candidate PQ nodes which provide the highest coverage among all protected nodes in the topology and which number does not exceed the value of parameter **max-pq-nodes**.

**Step 4.** Select a PQ Node.

If multiple PQ nodes satisfy the criteria in all the above steps, then R1 further selects the PQ node as follows:

- a. selects the lowest IGP cost from R1
- b. If more than one remains, R1 selects the PQ node reachable via the neighbor with the lowest router-id (OSPF) or system-id (ISIS);

If more than one remains, R1 selects the PQ node with the lowest router-id (OSPF) or system-id (ISIS).

For each destination prefix D, R1 programs the remote LFA backup path:

1. For prefixes of R4 or downstream of R4, R1 programs a node-protect remote LFA repair tunnel to the PQ node R6 by pushing the SID of node R6 on top of the SID for destination D and programming a next-hop of R3.
2. For prefixes owned by node R2, R1 runs the link-protect remote LFA algorithm and programs a simple link-protect repair tunnel which consists of a backup next-hop of R3 and pushing no additional label on top of the SID for the destination prefix.
3. Prefixes owned by nodes R3, R6, and R5 are not impacted by the failure of R2 since their primary next-hop is R3

#### **4.1.7.8.4 TI-LFA and Remote LFA Node Protection Feature Interaction and Limitations**

[LFA Protection Option Applicability](#) describes the order of activation of the various LFA types on a per prefix basis: TI-LFA, followed by base LFA, followed by remote LFA.

Node protection is enabled for TI-LFA and remote LFA separately. The base LFA prefers node protection over link protection.

The order of activation of the LFA types supersedes the protection type (node versus link). Consequently, it is possible that a prefix will be programmed with a link-protect backup next-hop by the more preferred LFA type. For example, a prefix is programmed with the only link-protect backup next-hop found by the base LFA while there exists a node-protect remote LFA next-hop.

#### **4.1.7.9 IPv6 Segment Routing using MPLS Encapsulation**

This feature implements support for SR IPv6 tunnels in IS-IS MT=0. The user can configure a node SID for the primary IPv6 global address of a loopback interface, which then gets advertised in IS-IS. IS-IS automatically assigns and advertises an adjacency SID for each adjacency with an IPv6 neighbor. After the node SID is resolved, it is used to install an IPv6 SR-ISIS tunnel in the TTM for use by the services.

##### **4.1.7.9.1 IS-IS MT=0 Extensions**

The IS-IS MT=0 extensions consist of supporting the advertising and resolution of the prefix SID sub-TLV within the IP Reach TLV-236 (IPv6), which is defined in RFC 5308.

The adjacency SID is still advertised as a sub-TLV of the Extended IS Reachability TLV 22, as defined in RFC 5305, *IS-IS Extensions for Traffic Engineering*, as in the case of an IPv4 adjacency. The router sets the V-Flag and I-Flag in the SR-Capabilities Sub-TLV to indicate that it is capable of processing SR MPLS encapsulated IPv4 and IPv6 packets on its network interfaces.

#### 4.1.7.9.2 Service and Forwarding Contexts Supported

The service and forwarding contexts supported with the SR-ISIS IPv6 tunnels are:

- SDP of type **sr-isis** with **far-end** option using IPv6 address
- VLL, VPLS, IES/VP RN spoke-interface, R-VPLS
- Support of PW redundancy within Epipe/Ipipe VLL, Epipe spoke termination on VPLS and R-VPLS, and Epipe/Ipipe spoke termination on IES/VP RN
- IPv6 static route resolution to indirect next-hop using Segment Routing IPv6 tunnel
- Remote mirroring and L3 encap LI

#### 4.1.7.9.3 Services Using SDP with a SR IPv6 Tunnel

The MPLS SDP of type **sr-isis** with a **far-end** option using an IPv6 address is supported. Note the SDP must have the same IPv6 **far-end** address, used by the control plane for the T-LDP session, as the prefix of the node SID of the SR IPv6 tunnel.

**Example:**

```
configure
 service
 [no] sdp sdp-id mpls
 [no] far-end ipv6-address
 sr-isis
 no sr-isis
```

The **bgp-tunnel**, **lsp**, **sr-te lsp**, **sr-ospf**, and **mixed-lsp-mode** commands are blocked within the SDP configuration context when the far-end is an IPv6 address.

SDP admin groups are not supported with an SDP using an SR IPv6 tunnel, or with SR-OSPF for IPv6 tunnels, and the attempt to assign them is blocked in the CLI.

Services that use LDP control plane such as T-LDP VPLS and R-VPLS, VLL, and IES/VP RN spoke interface have the spoke SDP (PW) signaled with an IPv6 T-LDP session because the **far-end** option is configured to an IPv6 address. The spoke SDP for these services binds to an SDP that uses an SR IPv6 tunnel where the prefix matches the **far-end** address. SR OS also supports the following:

- the IPv6 PW control word with both data plane packets and VCCV OAM packets
- Hash Label and Entropy Label, with the above services
- network domains in VPLS

The PW switching feature is not supported with LDP IPv6 control planes. As a result, the CLI does not allow the user to enable the **vc-switching** option whenever one or both spoke SDPs uses an SDP that has the **far-end** configured as an IPv6 address.

L2 services that use BGP control plane such as dynamic MS-PW, BGP-AD VPLS, BGP-VPLS, BGP-VPWS, and EVPN MPLS cannot bind to an SR IPv6 tunnel because a BGP session to a BGP IPv6 peer does not support advertising an IPv6 next-hop for the L2 NLRI. As a result, these services will not auto-generate SDPs using an SR IPv6 tunnel. In addition, they skip any provisioned SDPs with **far-end** configured to an IPv6 address when the **use-provisioned-sdp** option is enabled.

SR OS also supports multi-homing with T-LDP active/standby FEC 128 spoke SDP using SR IPv6 tunnel to a VPLS/B-VPLS instance. BGP multi-homing is not supported because BGP IPv6 does not support signaling an IPv6 next-hop for the L2 NLRI.

The Shortest Path Bridging (SPB) feature works with spoke SDPs bound to an SDP that uses an SR IPv6 tunnel.

#### 4.1.7.10 Data Path Support

A packet received with a label matching either a node SID or an adjacency SID is forwarded according to the ILM type and operation, as described in [Table 32](#).

**Table 32** Data Path Support


| Label type                    | Operation                                                                                                                                                                           |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Top label is a local node SID | Label is popped and the packet is further processed.<br>If the popped node SID label is the bottom of stack label, the IP packet is looked up and forwarded in the appropriate FIB. |

**Table 32 Data Path Support (Continued)**

| Label type                             | Operation                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Top or next label is a remote node SID | Label is swapped to the calculated label value for the next-hop and forwarded according to the primary or backup NHLFE. With ECMP, a maximum of 32 primary next-hops (NHLFEs) are programmed for the same destination prefix and for each IGP instance. ECMP and LFA next-hops are mutually exclusive as per existing implementation.                 |
| Top or next label is an adjacency SID  | Label is popped and the packet is forwarded out on the interface to the next-hop associated with this adjacency SID label.<br>In effect, the data path operation is modeled like a swap to an implicit-null label instead of a pop.                                                                                                                   |
| Next label is BGP 3107 label           | The packet is further processed according to the ILM operation as in current implementation. <ul style="list-style-type: none"> <li>• The BGP label may be popped and the packet looked up in the appropriate FIB.</li> <li>• The BGP label may be swapped to another BGP label.</li> <li>• The BGP label may be stitched to an LDP label.</li> </ul> |
| Next label is a service label          | The packet is looked up and forwarded in the Layer 2 or VPRN FIB as in current implementation.                                                                                                                                                                                                                                                        |

A router forwarding an IP or a service packet over an SR tunnel pushes a maximum of two transport labels with a remote LFA next-hop. This is illustrated in [Figure 26](#).

**Figure 26 Transport Label Stack in Shortest Path Forwarding with Segment Routing**



al\_0648

Assume that a VPRN service in node B forwards a packet received on a SAP to a destination VPN-IPv4 prefix X advertised by a remote PE2 via ASBR/ABR node A. Router B is in a segment routing domain while PE2 is in an LDP domain. BGP label routes are used to distribute the PE /32 loopbacks between the two domains.

When node B forwards over the primary next-hop for prefix X, it pushes the node SID of the ASBR followed by the BGP 3107 label for PE2, followed by the service label for prefix X. When the remote LFA next-hop is activated, node B pushes one or more segment routing label: the node SID for the remote LFA backup node (node N).

When node N receives the packet while the remote LFA next-hop is activated, it pops the top segment routing label which corresponds to a local node SID. This results in popping this label and forwarding of the packet to the ASBR node over the shortest path (link N-Z).

When the ABR/ASBR node receives the packet from either node B or node Z, it pops the segment routing label which corresponds to a local node SID, then swaps the BGP label and pushes the LDP label of PE2 which is the next-hop of the BGP label route.

#### 4.1.7.10.1 Hash Label and Entropy Label Support

When the **hash-label** option is enabled in a service context, hash label is always inserted at the bottom of the stack as per RFC 6391.

The LSR adds the capability to check a maximum of 16 labels in a stack. The LSR is able to hash on the IP headers when the payload below the label stack of maximum size of 16 is IPv4 or IPv6, including when a MAC header precedes it (**eth-encap-ip** option).

The Entropy Label (EL) feature, as specified in RFC 6790, is supported on RSVP, LDP, segment-routed, and BGP transport tunnels. It uses the Entropy Label Indicator (ELI) to indicate the presence of the entropy label in the label stack. The ELI, followed by the actual entropy label, is inserted immediately below the transport label for which entropy label feature is enabled. If multiple transport tunnels have the entropy label feature enabled, the ELI/EL is inserted below the lowest transport label in the stack.

The LSR hashing operates as follows:

- If the **lbl-only** hashing option is enabled, or if one of the other LSR hashing options is enabled but a IPv4 or IPv6 header is not detected below the bottom of the label stack, the LSR hashes on the EL only.
- If the **lbl-ip** option is enabled, the LSR hashes on the EL and the IP headers.
- If the **ip-only** or **eth-encap-ip** is enabled, the LSR hashes on the IP headers only.

For more information about the Hash Label and Entropy Label features, see the "MPLS Entropy Label and Hash Label" section of the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide*.

#### 4.1.7.11 IS-IS Control Protocol Changes

This section describes the [IS-IS Control Protocol Changes](#). See [OSPF Control Protocol Changes](#) for details on OSPF control protocol changes.



---

#### 4.1.7.11.1 IS-IS Control Protocol Changes

New TLV/sub-TLVs are defined in *draft-ietf-isis-segment-routing-extensions* and are supported in the implementation of segment routing in IS-IS. Specifically:

- the prefix SID sub-TLV
- the adjacency SID sub-TLV
- the SID/Label Binding TLV
- SR-Capabilities Sub-TLV
- SR-Algorithm Sub-TLV

This section describes the behaviors and limitations of the IS-IS support of segment routing TLV and sub-TLVs.

SR OS supports advertising the IS router capability TLV (RFC 4971) only for topology MT=0. As a result, the segment routing capability Sub-TLV can only be advertised in MT=0 which restricts the segment routing feature to MT=0.

Similarly, if prefix SID sub-TLVs for the same prefix are received in different MT numbers of the same IS-IS instance, then only the one in MT=0 is resolved. When the prefix SID index is also duplicated, an error is logged and a trap is generated, as explained in [Error and Resource Exhaustion Handling](#).

I and V flags are both set to 1 when originating the SR capability sub-TLV to indicate support for processing both SR MPLS encapsulated IPv4 and IPv6 packets on its network interfaces. These flags are not checked when the sub-TLV is received. Only the SRGB range is processed.

The algorithm field is set to 0, meaning Shortest Path First (SPF) algorithm based on link metric, when originating the SR-Algorithm capability sub-TLV but is not checked when the sub-TLV is received.

Both IPv4 and IPv6 prefix and adjacency SID sub-TLVs originate within MT=0.

SR OS originates a single prefix SID sub-TLV per IS-IS IP reachability TLV and processes the first prefix SID sub-TLV only if multiple are received within the same IS-IS IP reachability TLV.

SR OS encodes the 32 bit index in the prefix SID sub-TLV. The 24 bit label is not supported.

SR OS originates a prefix SID sub-TLV with the following encoding of the flags and the following processing rules:

- The R-flag is set if the prefix SID sub-TLV, along with its corresponding IP reachability TLV, is propagated between levels. See below for more details about prefix propagation.
- The N-flag is always set because SR OS supports prefix SID of type node SID only.
- The P-Flag (no-PHP flag) is always set, meaning that the label for the prefix SID is pushed by the PHP router when forwarding to this router. The SR OS PHP router processes properly a received prefix SID with the P-flag set to zero and uses implicit-null for the outgoing label towards the router which advertised it as long as the P-Flag is also set to 1.
- The E-flag (Explicit-Null flag) is always set to zero. An SR OS PHP router, however, processes properly a received prefix SID with the E-flag set to 1 and, when the P-flag is also set to 1, it pushes explicit-null for the outgoing label towards the router which advertised it.
- The V-flag is always set to 0 to indicate an index value for the SID.
- The L-flag is always set to 0 to indicate that the SID index value is not locally significant.
- The algorithm field is always set to zero to indicate Shortest Path First (SPF) algorithm based on link metric and is not checked on a received prefix SID sub-TLV.
- The SR OS still resolves a prefix SID sub-TLV received without the N-flag set but with the prefix length equal to 32. A trap, however, is raised by IS-IS.
- The SR OS will not resolve a prefix SID sub-TLV received with the N flag set and a prefix length different than 32. A trap is raised by IS-IS.
- The SR OS resolves a prefix SID received within a IP reachability TLV based on the following route preference:
  - SID received via L1 in a prefix SID sub-TLV part of IP reachability TLV
  - SID received via L2 in a prefix SID sub-TLV part of IP reachability TLV
- A prefix received in an IP reachability TLV is propagated, along with the prefix SID sub-TLV, by default from L1 to L2 by an L1L2 router. A router in L2 sets up an SR tunnel to the L1 router via the L1L2 router, which acts as an LSR.
- A prefix received in an IP reachability TLV is not propagated, along with the prefix SID sub-TLV, by default from L2 to L1 by an L1L2 router. If the user adds a policy to propagate the received prefix, then a router in L1 sets up an SR tunnel to the L2 router via the L1L2 router, which acts as an LSR.
- If a prefix is summarized by an ABR, the prefix SID sub-TLV is not propagated with the summarized route between levels. To propagate the node SID for a /32 prefix, route summarization must be disabled.

- SR OS propagates the prefix SID sub-TLV when exporting the prefix to another IS-IS instance; however, it does not propagate it if the prefix is exported from a different protocol. Thus, when the corresponding prefix is redistributed from another protocol such as OSPF, the prefix SID is removed.

SR OS originates an adjacency SID sub-TLV with the following encoding of the flags:

- the F-flag is set to zero to indicate the IPv4 family and is set to 1 to indicate an IPv6 family for the adjacency encapsulation.
- the B-Flag is set to zero and is not processed on receipt
- the V-flag is always set to 1
- the L-flag is always set to 1
- the S-flag is set to zero as assigning adjacency SID to parallel links between neighbors is not supported. An adjacency received SID with S-Flag set is not processed.
- the weight octet is not supported and is set to all zeros.

SR OS can originate the SID/Label Binding TLV as part of the Mapping Server feature (see [4.1.8](#) for more information). It can process it properly if received. The following rules and limitations should be considered.

- Only the Mapping Server Prefix-SID Sub-TLV within the TLV is processed and the ILMs installed if the prefixes in the provided range are resolved.
- The range and FEC prefix fields are processed. Each FEC prefix is resolved normally, as for the prefix SID sub-TLV, meaning there must be an IP Reachability TLV received for the exact matching prefix.
- If the same prefix is advertised with both a prefix SID sub-TLV and a mapping server Prefix-SID sub-TLV. The resolution follows the following route preference:
  - SID received via L1 in a prefix SID sub-TLV part of IP reachability TLV
  - SID received via L2 in a prefix SID sub-TLV part of IP reachability TLV
  - SID received via L1 in a mapping server Prefix-SID sub-TLV
  - SID received via L2 in a mapping server Prefix-SID sub-TLV
- The entire TLV can be propagated between levels based on the settings of the S-flag. The TLV cannot be propagated between IS-IS instances (see [4.1.8](#) for more information). Finally, an L1L2 router is not propagated the prefix-SID sub-TLV from the SID/Label binding TLV (received from a mapping server) into the IP Reachability TLV if the latter is propagated between levels.
- The mapping server which advertised the SID/Label Binding TLV does not need to be in the shortest path for the FEC prefix.

- If the same FEC prefix is advertised in multiple binding TLVs by different routers, the SID in the binding TLV of the first router which is reachable is used. If that router becomes unreachable, the next reachable one is used.
- No check is performed if the content of the binding TLVs from different mapping servers are consistent or not.
- Any other sub-TLV, for example, the SID/Label Sub-TLV, ERO metric and unnumbered interface ID ERO, is ignored but the user can get a dump of the octets of the received but not-supported sub-TLVs using the existing IGP **show** command.

#### 4.1.7.12 BGP Shortcut Using Segment Routing Tunnel

The user enables the resolution of IPv4 prefixes using SR tunnels to BGP next-hops in TTM with the following command:

**CLI Syntax:**

```
configure>router>bgp>next-hop-resolution
 shortcut-tunnel
 [no] family {ipv4}
 resolution {any | disabled | filter}
 resolution-filter
 [no] sr-isis
 [no] sr-ospf
 [no] disallow-igp
 exit
 exit
exit
```

When **resolution** is set to **any**, any supported tunnel type in BGP shortcut context is selected following TTM preference. The following tunnel types are supported in a BGP shortcut context and in order of preference: RSVP, LDP, Segment Routing and BGP.

When the **sr-isis** or **sr-ospf** value is enabled, an SR tunnel to the BGP next-hop is selected in the TTM from the lowest preference IS-IS or OSPF instance. If many instances have the same lowest preference from the lowest numbered IS-IS or OSPF instance

See the BGP chapter for more details.

### 4.1.7.13 BGP Label Route Resolution Using Segment Routing Tunnel

The user enables the resolution of RFC 3107 BGP label route prefixes using SR tunnels to BGP next-hops in TTM with the following command:

**CLI Syntax:**

```
configure>router>bgp>next-hop-resolution
 labeled-routes
 transport-tunnel
 [no] family {label-ipv4 | label-ipv6 |
 vpn}
 resolution {any | disabled | filter}
 resolution-filter
 [no] sr-isis
 [no] sr-ospf
 exit
 exit
 exit
exit
```

When the **resolution** option is explicitly set to **disabled**, the default binding to LDP tunnel resumes. If **resolution** is set to **any**, any supported tunnel type in BGP label route context is selected following TTM preference.

The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing.

When the **sr-isis** or **sr-ospf** is specified using the **resolution-filter** option, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered IS-IS or OSPF instance.

See the BGP chapter for more details.

### 4.1.7.14 Service Packet Forwarding with Segment Routing

SDP subtypes of the MPLS type are available to allow service binding to an SR tunnel programmed in TTM by OSPF or IS-IS:

**\*A:7950 XRS-20# configure service sdp 100 mpls create**

**\*A:7950 XRS-20>config>service>sdp\$ sr-ospf**

**\*A:7950 XRS-20>config>service>sdp\$ sr-isis**

The SDP of type **sr-isis** or **sr-ospf** can be used with the **far-end** option. When the **sr-isis** or **sr-ospf** value is enabled, a tunnel to the far-end address is selected in the TTM from the lowest preference IS-IS or OSPF instance. If many instances have the same lowest preference from the lowest numbered IS-IS or OSPF instance. The SR-ISIS or SR-OSPF tunnel is selected at the time of the binding, following the tunnel selection rules. If a more preferred tunnel is subsequently added to the TTM, the SDP will not automatically switch to the new tunnel until the next time the SDP is being re-resolved.

The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-isis** and **sr-ospf** tunnel types.

The signaling protocol for the service labels for an SDP using an SR tunnel can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

SR tunnels can be used in VPRN and BGP EVPN with the **auto-bind-tunnel** command. See [Next-Hop Resolution](#) for more information.

Both VPN-IPv4 and VPN-IPv6 (6VPE) are supported in a VPRN or BGP EVPN service using segment routing transport tunnels with the **auto-bind-tunnel** command.

See [BGP](#) and refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information about the VPRN **auto-bind-tunnel** CLI command.

#### 4.1.7.15 Mirror Services and Lawful Intercept

The user can configure a spoke-SDP bound to an SR tunnel to forward mirrored packets from a mirror source to a remote mirror destination. In the configuration of the mirror destination service at the destination node, the **remote-source** command must use a spoke-sdp with VC-ID which matches the one the user configured in the mirror destination service at the mirror source node. The far-end option is not supported with an SR tunnel.

This also applies to the configuration of the mirror destination for an LI source.

Configuration at mirror source node:

**CLI Syntax:**

```
config mirror mirror-dest 10
 no spoke-sdp <sdp-id:vc-id>
 spoke-sdp <sdp-id:vc-id> [create]
 egress
 vc-label <egress-vc-label>
```



**Note:**

- *sdp-id* matches an SDP which uses an SR tunnel
- for vc-label, both static and t-ldp egress vc labels are supported

Configuration at mirror destination node:

**CLI Syntax:** \*A:7950 XRS-20# configure mirror mirror-dest 10 remote-source  
 spoke-sdp <SDP-ID>:<VC-ID> create <-- VC-ID  
*matching that of spoke-sdp configured in mirror destination context at mirror source node.*  
 ingress  
     vc-label <ingress-vc-label> <---  
         optional: both static and t-ldp ingress  
         vc label are supported.  
 exit  
     no shutdown  
 exit  
 exit



**Note:**

- the **far-end** command is not supported with SR tunnel at mirror destination node; user must reference a spoke-SDP using a segment routing SDP coming from mirror source node:
  - **far-end** *ip-address* [**vc-id** *vc-id*] [**ing-svc-label** *ingress-vc-label* | **tldp**] [**icb**]
  - **no far-end** *ip-address*
- for vc-label, both static and t-ldp ingress vc labels are supported

Mirroring and LI are also supported with the PW redundancy feature when the endpoint spoke-sdp, including the ICB, is using an SR tunnel. Routable Lawful Intercept Encapsulation (**config>mirror>mirror-dest>encap# layer-3-encap**) when the remote L3 destination is reachable over an SR tunnel is also supported.

## 4.1.8 Segment Routing Mapping Server Function for IPv4 Prefixes

### 4.1.8.1 Segment Routing Mapping Server

The mapping server feature allows the configuration and advertisement via IS-IS of the node SID index for prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server and using a prefix-SID sub-TLV within the SID/Label binding TLV in IS-IS.

The user configures the SR mapping database in IS-IS using the following CLI command:

**CLI Syntax:**

```
configure
router
 [no] isis
 segment-routing
 no segment-routing
 mapping-server
 sid-map node-sid {index value
 <0..4294967295> [range value
 <0..65535>]} prefix {{ip-
 address/mask} | {ip-
 address}{netmask}} [set-flags
 {s}] [level {1|2|1/2}]
 no sid-map node-sid index value}
```

The user can enter the node SID index for one prefix or a range of prefixes by specifying the first index value and, optionally, a range value. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. The user can enter the first prefix with a mask lower than 32 and the SID/Label Binding TLV is advertised but the routers will not resolve these prefix SIDs and instead originates a trap.

The **no** form of the **sid-map** command deletes the range of node SIDs beginning with the specified index value. The **no** form of the **mapping-server** command deletes all node SID entries in the IS-IS instance.

By setting the S-flag, the user can indicate to the IS-IS routers in the rest of the network that the flooding scope of the SID/Label binding TLV is the entire domain. In that case, a router receiving the TLV advertisement should leak it between IS-IS levels. If leaked from level 2 to level 1, the D-flag must be set and, once set, the TLV cannot be leaked back into level 2. Otherwise, the S-flag is clear by default and the TLV must not be leaked by routers receiving the mapping server advertisement.





**Note:** SR OS does not leak this TLV between IS-IS instances and does not support the multi-topology SID/Label Binding TLV format. In addition, the user can specify the mapping server's flooding scope for the generated SID/Label binding TLV using the **level** option. This option allows further narrowing of the flooding scope configured under the router IS-IS level-capability for a one or more SID/Label binding TLVs if required. The default flooding scope of the mapping server is L1/L2, which can be narrowed by what is configured under the router IS-IS level-capability.

The A-flag is used to indicate that a prefix for which the mapping server prefix SID is advertised is directly attached. The M-flag is used to advertise a SID for a mirroring context to provide protection against the failure of a service node. None of these flags are supported on the mapping server; the mapping client ignores them.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix, or range of prefixes, a prefix-SID sub-TLV within a IS-IS SID/label binding TLV in that instance. The flooding scope of the TLV from the mapping server is determined as explained above. No further check of the reachability of that prefix in the mapping server route table is performed and no check if the SID index is duplicate with some existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

## 4.1.8.2 Prefix SID Resolution for Segment Routing Mapping Server

### 1. IP Prefix Resolution

- a. SPF calculates the next-hops, up to **max-ecmp**, to reach a destination node.
- b. Each prefix inherits the next-hops of one or more destination nodes advertising it.
- c. A prefix advertised by multiple nodes, all reachable with same cost, will inherit up to **max-ecmp** next-hops from the advertising nodes.
- d. The next-hop selection, up to **max-ecmp**, value is deterministic and is based on sorting the next-hops by:
  - i. lowest next-hop router-id
  - ii. lowest interface index, for parallel links to same router-id

Each next-hop keeps a reference to the destination nodes of whom it was inherited.

### 2. Prefix SID Resolution

- a. For a specified prefix, IGP selects the SID value among multiple advertised values respecting the following preference order:

- i. the local intra-area SID owned by this router
- ii. the prefix SID sub-TLV advertised within an IP Reach TLV  
If multiple SIDs exist in this step, select the SID corresponding to the destination router or the ABR with the lowest system ID that is reachable using the first next-hop of the prefix.
- iii. the IS-IS SID and label binding TLV from the mapping server.  
If multiple SIDs exist in this step, select the following the preference rules in *draft-ietf-spring-conflict-resolution-05 [sid-conflict-resolution]* when applied to the SRMS entries of the conflicting SIDs. The order of these rules is as follows:
  - smallest range
  - smallest starting address
  - smallest algorithm
  - smallest starting SID



**Note:** If an L1L2 router acts as a mapping server and also re-advertises the mapping server prefix SID from other mapping servers, the redistributed mapping server prefix SID is preferred by other routers resolving the prefix, which may result in not selecting the mapping server respecting these rules.

- b. The selected SID is used with all ECMP next-hops step (1) towards all destination nodes or ABR nodes which advertised the prefix.
  - c. If duplicate prefix SIDs exist for different prefixes after the above steps, the first SID that is processed is programmed for its corresponding prefix. Subsequent SIDs cause a duplicate SID trap and are not programmed. The corresponding prefixes themselves are still resolved and programmed normally using IP next-next-hops.
3. SR Tunnel Programming
- a. If the prefix SID is resolved from a prefix SID sub-TLV advertised within an IP Reach TLV, the SR ILM is swapped to an SR NHLFE as with regular SR tunnel resolution.
  - b. If the prefix SID is resolved from a mapping server advertisement:
    - i. The SR ILM is stitched to an LDP FEC of the same prefix if one exists. The LDP FEC could have been resolved using a static route, a route within a IS-IS instance, or a route within a OSPF instance. The latter two can be the same or different as the IGP instance which advertised the mapping server prefix SID sub-TLV.

- 
- ii. Otherwise, the SR ILM label is swapped to a SR NHLFE label. This is only possible if a route is exported from another IGP instance into the local IGP instance without propagating the prefix SID sub-TLV with the route. Otherwise, when the prefix SID is propagated with the route, the resolution follows step [3a](#).

---

## 4.2 FIB Prioritization

The RIB processing of specific routes can be prioritized through the use of the **rib-priority command**. This command allows specific routes to be prioritized through the protocol processing so that updates are propagated to the FIB as quickly as possible.


The **rib-priority** command is configured within the global IS-IS routing context, and the administrator has the option to either specify a prefix-list or an IS-IS tag value. If a prefix list is specified then route prefixes matching any of the prefix list criteria will be considered high priority. If instead an IS-IS tag value is specified then any IS-IS route with that tag value will be considered high priority.

The routes that have been designated as high priority are the first routes processed and then passed to the FIB update process so that the forwarding engine can be updated. All known high priority routes should be processed before the IS-IS routing protocol moves on to other standard priority routes. This feature has the most impact when there are a large number of routes being learned through the IS-IS routing protocols.

## 4.3 IS-IS Configuration Process Overview

Figure 27 displays the process to provision basic IS-IS parameters.

**Figure 27** IS-IS Configuration and Implementation Flow



Router\_Config\_22

## 4.4 Configuration Notes

This section describes IS-IS configuration caveats.

### 4.4.1 General

- IS-IS must be enabled on each participating router.
- There are no default network entity titles.
- There are no default interfaces.
- By default, the routers are assigned a Level 1/Level 2 level capability.

---

## 4.5 Configuring IS-IS with CLI

This section provides information to configure IS-IS using the command line interface.

### 4.5.1 IS-IS Configuration Overview

#### 4.5.1.1 Router Levels

The router's level capability can be configured globally and on a per-interface basis. The interface-level parameters specify the interface's routing level. The neighbor capability and parameters define the adjacencies that are established.

IS-IS is not enabled by default. When IS-IS is enabled, the global default level capability is Level 1/2 which enables the router to operate as either a Level 1 and/or a Level 2 router with the associated databases. The router runs separate shortest path first (SPF) calculations for the Level 1 area routing and for the Level 2 multi-area routing to create the IS-IS routing table.

The level value can be modified on both or either of the global and interface levels to be only Level 1-capable, only Level 2-capable or Level 1 *and* Level 2-capable.

If the default value is not modified on any routers in the area, then the routers try to form both Level 1 and Level 2 adjacencies on all IS-IS interfaces. If the default values are modified to Level 1 or Level 2, then the number of adjacencies formed are limited to that level only.

#### 4.5.1.2 Area Address Attributes

The area-id command specifies the area address portion of the NET which is used to define the IS-IS area to which the router will belong. At least one area-id command should be configured on each router participating in IS-IS. A maximum of three area-id commands can be configured per router.

The area address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*. The routers in an area manage routing tables about destinations within the area. The Network Entity Title (NET) value is used to identify the IS-IS area to which the router belongs.

NSAP addresses are divided into three parts. Only the Area ID portion is configurable.

- Step 1.** Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- Step 2.** System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.
- Step 3.** Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The following example displays ISO addresses in IS-IS address format:

MAC address 00:a5:c7:6b:c4:9049.0011.00a5.c76b.c490.00  
IP address: 218.112.14.5 49.0011.2181.1201.4005.00

### 4.5.1.3 Interface Level Capacity

The level capability value configured on the interface level is compared to the level capability value configured on the global level to determine the type of adjacencies that can be established. The default level capability for routers and interfaces is Level 1/2.

[Table 33](#) displays configuration combinations and the potential adjacencies that can be formed.

**Table 33** Potential Adjacency

| Global Level | Interface Level | Potential Adjacency    |
|--------------|-----------------|------------------------|
| L 1/2        | L 1/2           | Level 1 and/or Level 2 |
| L 1/2        | L 1             | Level 1 only           |
| L 1/2        | L 2             | Level 2 only           |
| L 2          | L 1/2           | Level 2 only           |
| L 2          | L 2             | Level 2 only           |
| L 2          | L 1             | none                   |
| L 1          | L 1/2           | Level 1 only           |
| L 1          | L 2             | none                   |



**Table 33 Potential Adjacency (Continued)**

| Global Level | Interface Level | Potential Adjacency |
|--------------|-----------------|---------------------|
| L 1          | L 1             | Level 1 only        |

#### 4.5.1.4 Route Leaking

Nokia's implementation of IS-IS route leaking is performed in compliance with RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*. As previously stated, IS-IS is a routing domain (an autonomous system running IS-IS) which can be divided into Level 1 areas with a Level 2-connected subset (backbone) of the topology that interconnects all of the Level 1 areas. Within each Level 1 area, the routers exchange link state information. Level 2 routers also exchange Level 2 link state information to compute routes between areas.

Routers in a Level 1 area typically only exchange information within the Level 1 area. For IP destinations not found in the prefixes in the Level 1 database, the Level 1 router forwards PDUs to the nearest router that is in both Level 1/Level 2 with the *attached bit* set in its Level 1 link-state PDU.

There are many reasons to implement domain-wide prefix distribution. The goal of domain-wide prefix distribution is to increase the granularity of the routing information within the domain. The routing mechanisms specified in RFC 1195 are appropriate in many situations and account for excellent scalability properties. However, in certain circumstances, the amount of scalability can be adjusted which can distribute more specific information than described by RFC 1195.

Distributing more prefix information can improve the quality of the resulting routes. A well-known property of default routing is that loss of information can occur. This loss of information affects the computation of a route based upon less information which can result in sub-optimal routes.

### 4.5.2 Basic IS-IS Configuration

For IS-IS to operate on the routers, IS-IS must be explicitly enabled, and at least one area address and interface must be configured. If IS-IS is enabled but no area address or interface is defined, the protocol is enabled but no routes are exchanged. When at least one area address and interface are configured, then adjacencies can be formed and routes exchanged.

To configure IS-IS, perform the following tasks:

- Enable IS-IS (specifying the instance ID of multi-instance IS-IS is to be enabled).
- If necessary, modify the level capability on the global level (default is level-1/2).
- Define area addresses
- Configure IS-IS interfaces.

The following output displays IS-IS default values.

```
*A:Dut-D>config>router>isis# info detail

no system-id
no router-id
level-capability level-1/2
no graceful-restart
no auth-keychain
no authentication-key
no authentication-type
authentication-check
csnp-authentication
no ignore-lsp-errors
no ignore-narrow-metric
lsp-lifetime 1200
lsp-mtu-size 1492
lsp-refresh-interval 600
no export-limit
no export
no import
hello-authentication
psnp-authentication
no traffic-engineering
no reference-bandwidth
no default-route-tag
no disable-ldp-sync
no advertise-passive-only
no advertise-router-capability
no hello-padding
no ldp-over-rsvp
no advertise-tunnel-link
no ignore-attached-bit
no suppress-attached-bit
no iid-tlv-enable
no poi-tlv-enable
no prefix-limit
no loopfree-alternate
no loopfree-alternate-exclude
no rib-priority high
ipv4-routing
no ipv6-routing
ipv4-multicast-routing native
ipv6-multicast-routing native
no multi-topology
no unicast-import-disable both
no multicast-import both
no strict-adjacency-check
igp-shortcut
 shutdown
 tunnel-next-hop
```

```
 family ipv4
 resolution disabled
 resolution-filter
 no rsvp
 no sr-te
 exit
 family ipv6
 resolution disabled
 resolution-filter
 no rsvp
 no sr-te
 exit
 family srv4
 resolution disabled
 resolution-filter
 no rsvp
 no sr-te
 exit
 family srv6
 resolution disabled
 resolution-filter
 no rsvp
 no sr-te
 exit
 exit
exit
timers
 lsp-wait 5000 lsp-initial-wait 10 lsp-second-wait 1000
 sfp-wait 10000 sfp-initial-wait 1000 sfp-second-wait 1000
exit
level 1
 advertise-router-capability
 no hello-padding
 no lsp-mtu-size
 no auth-keychain
 no authentication-key
 no authentication-type
 csnf-authentication
 external-preference 160
 hello-authentication
 no loopfree-alternate-exclude
 preference 15
 psnf-authentication
 no wide-metrics-only
 default-metric 10
 default-ipv4-multicast-metric 10
 default-ipv6-unicast-metric 10
 default-ipv6-multicast-metric 10
exit
level 2
 advertise-router-capability
 no hello-padding
 no lsp-mtu-size
 no auth-keychain
 no authentication-key
 no authentication-type
 csnf-authentication
 external-preference 165
 hello-authentication
```

```
no loopfree-alternate-exclude
preference 18
psnp-authentication
no wide-metrics-only
default-metric 10
default-ipv4-multicast-metric 10
default-ipv6-unicast-metric 10
default-ipv6-multicast-metric 10
exit
segment-routing
shutdown
adj-sid-hold 15
no export-tunnel-table
no prefix-sid-range
no tunnel-table-pref
no tunnel-mtu
mapping-server
shutdown
exit
exit
no shutdown
```

### 4.5.3 Common Configuration Tasks

To implement IS-IS in your network, you must enable IS-IS on each participating router.

To assign different level to the routers and organize your network into areas, modify the level capability defaults on end systems from Level 1/2 to Level 1. Routers communicating to other areas can retain the Level 1/2 default.

On each router, at least one area ID also called the area address should be configured as well as at least one IS-IS interface.

- Enable IS-IS.
- Configure global IS-IS parameters.
  - Configure area addresses.
- Configure IS-IS interface-specific parameters.

### 4.5.4 Configuring IS-IS Components

Use the CLI syntax displayed in the following subsections to configure IS-IS components.

### 4.5.4.1 Enabling IS-IS

IS-IS must be enabled in order for the protocol to be active.



**Caution:** Careful planning is essential to implement commands that can affect the behavior of global and interface levels.

To configure IS-IS on a router, enter the following command:

**CLI Syntax:**     `isis`

**Example:**       `config>router# isis`

IS-IS also supports the concept of multi-instance IS-IS which allows separate instances of the IS-IS protocol to run independently of the SR-OS router.

Separate instances are created by adding a different instance ID as the optional parameter to the **config>router>isis** command.

### 4.5.4.2 Modifying Router-Level Parameters

When IS-IS is enabled, the default **level-capability** is Level 1/2. This means that the router operates with both Level 1 and Level 2 routing. To change the default value in order for the router to operate as a Level 1 router or a Level 2 router, you must explicitly modify the **level** value.

If the level is modified, the protocol shuts down and restarts. Doing this can affect adjacencies and routes.

The **level-capability** value can be configured on the global level and also on the interface level. The **level-capability** value determines which level values can be assigned on the router level or on an interface-basis.

In order for the router to operate as a Level 1 only router or as a Level 2 only router, you must explicitly specify the **level-number** value.

- Select **level-1** to route only within an area.
- Select **level-2** to route to destinations outside an area, toward other eligible Level 2 routers.

To configure the router level, enter the following commands:

**CLI Syntax:**    config>router# isis  
                  level-capability {level-1|level-2|level-1/2}  
                  level {1|2}

**Example:**        config>router# isis  
                  config>router>isis# level-capability 1/2  
                  config>router>isis# level 2

The following example displays the configuration:

```
A:ALA-A>config>router>isis# info
#-----
echo "ISIS"
#-----

level-capability level-1/2
level 2

A:ALA-A>config>router>isis#
```

#### 4.5.4.3 Configuring ISO Area Addresses

Use the following CLI syntax to configure an area ID also called an address. A maximum of 3 area-id can be configured.

**CLI Syntax:**    config>router# isis area-id *area-address*

The following example configures the router's area ID:

**Example:**        config>router>isis#  
                  config>router>isis# area-id 49.0180.0001  
                  config>router>isis# area-id 49.0180.0002  
                  config>router>isis# area-id 49.0180.0003

The following example displays the area ID configuration:

```
A:ALA-A>config>router>isis# info

area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003

A:ALA-A>config>router>isis#
```

#### 4.5.4.4 Configuring Global IS-IS Parameters

Commands and parameters configured on the global level are inherited to the interface levels. Parameters specified in the interface and interface-level configurations take precedence over global configurations.

The following example displays global-level IS-IS configuration command usage:

**Example:**

```
config>router# isis
config>router>isis# level-capability level-2
config>router>isis# authentication-check
config>router>isis# authentication-type password
config>router>isis# authentication-key test
config>router>isis# overload timeout 90
config>router>isis# traffic-engineering
```

The following example displays the modified global-level configuration.

```
A:ALA-A>config>router>isis# info

level-capability level-2
area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003
authentication-key "H5KBAWrAAQU" hash
authentication-type password
overload timeout 90
traffic-engineering

A:ALA-A>config>router>isis#
```

#### 4.5.4.5 Migration to IS-IS Multi-Topology

To migrate to IS-IS multi-topology for IPv6, perform the following tasks:

Enable the sending/receiving of IPv6 unicast reachability information in IS-IS MT TLVs on all the routers that support MT.

**CLI Syntax:** `config>router# isis multi-topology ipv6-unicast`

```
A:ALA-49>config>router>isis# info detail

...
ipv4-routing
ipv6-routing native
multi-topology
 ipv6-unicast
exit
...
```

```

A:ALA-49>config>router>isis#
```

Ensure that all MT routers have the IPv6 reachability information required by MT TLVs:

**CLI Syntax:** show>router# isis topology ipv6-unicast

```
A:ALA-49>config>router>isis# show router isis topology ipv6-unicast
=====
Topology Table
=====
Node Interface Nexthop

No Matching Entries
=====
A:ALA-49>config>router>isis#
```

**CLI Syntax:** show>router# isis database detail

```
A:ALA-49>config>router>isis# show router isis database detail
=====
Rtr Base ISIS Instance 0 Database (detail)
=====
Displaying Level 1 database

LSP ID : ALA-49.00-00 Level : L1
Sequence : 0x22b Checksum : 0x60e4 Lifetime : 1082
Version : 1 Pkt Type : 18 Pkt Ver : 1
Attributes: L1L2 Max Area : 3
SysID Len : 6 Used Len : 404 Alloc Len : 1492

TLVs :
Area Addresses :
 Area Address : (13) 47.4001.8000.00a7.0000.ffdd.0007
Supp Protocols :
 Protocols : IPv4 IPv6
IS-Hostname :
 Hostname : ALA-49
TE Router ID :
 Router ID : 10.10.10.104
Internal Reach :
 IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0 (I)
 IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10 (I)
MT IPv6 Reach. :
 MT ID : 2
 IPv6 Prefix : 3ffe::101:100/120
 Flags : Up Internal Metric : 10
 IPv6 Prefix : 10::/64
 Flags : Up Internal Metric : 10
I/f Addresses :
 IP Address : 10.10.10.104
```



```

IP Address : 10.10.4.3
IP Address : 10.10.5.3
IP Address : 10.10.7.3
IP Address : 10.10.0.16
IP Address : 10.0.0.104
I/f Addresses IPv6 :
 IPv6 Address : 3FFE::101:101
 IPv6 Address : 10::104
TE IP Reach. :
 IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0
 IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10
 IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10
 IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10
 IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10
 IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10
Authentication :
 Auth Type : Password(1) (116 bytes)

Level (1) LSP Count : 1

Displaying Level 2 database

LSP ID : ALA-49.00-00 Level : L2
Sequence : 0x22c Checksum : 0xb888 Lifetime : 1082
Version : 1 Pkt Type : 20 Pkt Ver : 1
Attributes: L1L2 Max Area : 3
SysID Len : 6 Used Len : 304 Alloc Len : 1492

TLVs :
Area Addresses :
 Area Address : (13) 47.4001.8000.00a7.0000.ffdd.0007
Supp Protocols :
 Protocols : IPv4 IPv6
IS-Hostname :
 Hostname : ALA-49
TE Router ID :
 Router ID : 10.10.10.104
Internal Reach :
 IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0 (I)
 IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10 (I)
 IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10 (I)
MT IPv6 Reach. :
 MT ID : 2
 IPv6 Prefix : 3ffe::101:100/120
 Flags : Up Internal Metric : 10
 IPv6 Prefix : 10::/64
 Flags : Up Internal Metric : 10
I/f Addresses :
 IP Address : 10.10.10.104
 IP Address : 10.10.4.3
 IP Address : 10.10.5.3
 IP Address : 10.10.7.3
 IP Address : 10.10.0.16
 IP Address : 10.0.0.104
I/f Addresses IPv6 :
 IPv6 Address : 3FFE::101:101

```

```

IPv6 Address : 10::104
TE IP Reach. :
IP Prefix : 10.10.10.104/32 (Dir. :Up) Metric : 0
IP Prefix : 10.10.4.0/24 (Dir. :Up) Metric : 10
IP Prefix : 10.10.5.0/24 (Dir. :Up) Metric : 10
IP Prefix : 10.10.7.0/24 (Dir. :Up) Metric : 10
IP Prefix : 10.10.0.0/24 (Dir. :Up) Metric : 10
IP Prefix : 10.0.0.0/24 (Dir. :Up) Metric : 10
Authentication :
Auth Type : MD5(54) (16 bytes)

```

```

Level (2) LSP Count : 1

```

```

Flags : D = Prefix Leaked Down
 : N = Node Flag
 : R = Re-advertisement Flag
 : S = Sub-TLVs Present
 : X = External Prefix Flag
=====

```

```

A:ALA-49>config>router>isis#

```

### Configure MT TLVs for IPv6 SPF:

**CLI Syntax:** config>router# isis ipv6-routing mt

```

A:ALA-49>config>router>isis# info detail

```

```

...
 ipv4-routing
 ipv6-routing mt
 multi-topology
 ipv6-unicast
 exit
...

```

```

A:ALA-49>config>router>isis#

```

### Verify IPv6 routes:

**CLI Syntax:** show>router# isis routes ipv6-unicast

```

A:ALA-49>config>router>isis# show router isis routes ipv6-unicast

```

```

=====
Rtr Base ISIS Instance 0 Route Table
=====

```

| Prefix[Flags] | Metric | Lvl/Typ | Ver. | SysID/Hostname  |
|---------------|--------|---------|------|-----------------|
| NextHop       |        |         | MT   | AdminTag/SID[F] |

```

No Matching Entries
=====

```

```

A:ALA-49>config>router>isis#

```

**CLI Syntax:** show>router# route-table ipv6

```

A:ALA-48>show>router# route-table ipv6

```

```

=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix Type Proto Age Pref
 Next Hop[Interface Name] Metric

10::/64 Local Local 05h35m28s 0
 to-104 0

No. of Routes: 1
=====
A:ALA-48>show>router#

```

#### 4.5.4.6 Configuring Interface Parameters

There are no interfaces associated with IS-IS by default. An interface belongs to all areas configured on a router. Interfaces cannot belong to separate areas. There are no default interfaces applied to the router's IS-IS instance. You must configure at least one IS-IS interface in order for IS-IS to work.

To enable IS-IS on an interface, first configure an IP interface in the **config>router>interface** context. Then, apply the interface in the **config>router>isis>interface** context.

You can configure both the Level 1 parameters and the Level 2 parameters on an interface. The **level-capability** value determines which level values are used.



**Note:** For point-to-point interfaces, only the values configured under Level 1 are used regardless of the operational level of the interface.

The following example displays the modified interface parameters:

**Example:**

```

config>router# isis
config>router>isis# level 1
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# level 2
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# interface ALA-1-2
config>router>isis>if# level-capability level-2
config>router>isis>if# exit
config>router>isis# interface ALA-1-3
config>router>isis>if# level-capability level-1
config>router>isis>if# interface-type point-to-point
config>router>isis>if# exit

```

```
config>router>isis# interface ALA-1-5
config>router>isis>if# level-capability level-1
config>router>isis>if# interface-type point-to-point
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# level-capability level-1/2
config>router>isis>if# exit
config>router>isis#
```

The following example displays the global and interface-level configurations.

```
A:ALA-A>config>router>isis# info


level-capability level-2
area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003
authentication-key "H5KBAWrAAQU" hash
authentication-type password
traffic-engineering
level 1
 wide-metrics-only
exit
level 2
 wide-metrics-only
exit
interface "system"
exit
interface "ALA-1-2"
 level-capability level-2
exit
interface "ALA-1-3"
 level-capability level-1
 interface-type point-to-point
exit
interface "ALA-1-5"
 level-capability level-1
 interface-type point-to-point
exit
interface "to-103"
exit

A:ALA-A>config>router>isis#
```

#### 4.5.4.6.1 Example: Configuring a Level 1 Area

Interfaces are configured in the **config>router>interface** context, as shown in [Figure 28](#).

**Figure 28 Configuring a Level 1 Area**



OSRG031

The following example displays the command usage to configure a Level 1 area.

```

A:ALA-A>config>router# isis
A:ALA-A>config>router>isis# area-id 47.0001
A:ALA-A>config>router>isis# level-capability level-1
A:ALA-A>config>router>isis# interface system
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis# interface A-B
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis# interface A-C
A:ALA-A>config>router>isis>if# exit
A:ALA-A>config>router>isis#

```

```

A:ALA-B>config>router# isis
A:ALA-B>config>router>isis# area-id 47.0001
A:ALA-B>config>router>isis# level-capability level-1
A:ALA-B>config>router>isis# interface system
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis# interface B-A
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis# interface B-C
A:ALA-B>config>router>isis>if# exit
A:ALA-B>config>router>isis#

```

```

A:ALA-C>config>router# isis
A:ALA-C>config>router>isis# area-id 47.0001
A:ALA-C>config>router>isis# level-capability level-1
A:ALA-C>config>router>isis# interface system
A:ALA-C>config>router>isis>if# exit
A:ALA-C>config>router>isis# interface "C-A"
A:ALA-C>config>router>isis>if# exit
A:ALA-C>config>router>isis# interface "C-B"
A:ALA-C>config>router>isis>if# exit

```

```

A:ALA-A>config>router>isis# info

level-capability level-1
area-id 49.0180.0001
interface "system"

```

```

 exit
 interface "A-B"
 exit
 interface "A-C"
 exit

A:ALA-A>config>router>isis#

A:ALA-B>config>router>isis# info

 level-capability level-1
 area-id 49.0180.0001
 interface "system"
 exit
 interface "B-A"
 exit
 interface "B-C"
 exit

A:ALA-B>config>router>isis#

A:ALA-C>config>router>isis# info
#-----
echo "ISIS"

 level-capability level-1
 area-id 49.0180.0001
 interface "system"
 exit
 interface "C-A"
 exit
 interface "C-B"
 exit


A:ALA-C>config>router>isis#

```

#### 4.5.4.6.2 Example: Modifying a Router's Level Capability

In the previous example, ALA-A, ALA-B, and ALA-C are configured as Level 1 systems. Level 1 systems communicate with other Level 1 systems in the same area. In this example, ALA-A is modified to set the level capability to Level 1/2, as shown in [Figure 29](#). Now, the Level 1 systems in the area with NET 47.0001 forward PDUs to ALA-A for destinations that are not in the local area.

**Figure 29 Configuring a Level 1/2 Area**



OSRG036

The following example displays the command usage to configure a Level 1/2 system.

```

A:ALA-A>config>router# isis
A:ALA-A>config>router>isis# level-capability level-1/2

```

#### 4.5.4.7 Configuring IS-IS Link Groups

IS-IS Link-Groups allows the creation of an administrative grouping of multiple IS-IS member interfaces that should be treated as a common group for ECMP purposes. If the number of operational links in the link-group drops below the operational-member value then all links associated with that IS-IS link group will have their interface metric increased by the configured offset amounts. As a result, IS-IS will then try to reroute traffic over lower cost paths.

Once triggered, the higher metric will not be reset to the originally configured IS-IS interface metric values until the number of active interfaces in the link bundle reaches the configured revertive threshold (**revert-members**).

Prerequisite are the following:

- 1 or more interface members.
- A configured operational-member (**oper-members**) value.
- A configured revertive-member (**revert-members**) value.
- Configured offset values for the appropriate address families.

---

## 4.6 IS-IS Configuration Management Tasks

This section discusses IS-IS configuration management tasks.

### 4.6.1 Disabling IS-IS

The **shutdown** command disables the IS-IS protocol instance on the router. The configuration settings are not changed, reset, or removed.

To disable IS-IS on a router, enter the following commands:

**CLI Syntax:**    `config>router# isis  
                  shutdown`

### 4.6.2 Removing IS-IS

The **no isis** command deletes the IS-IS protocol instance. The IS-IS configuration reverts to the default settings.

To remove the IS-IS configuration enter the following commands:

**CLI Syntax:**    `config>router#  
                  no isis`

### 4.6.3 Modifying Global IS-IS Parameters

You can modify, disable, or remove global IS-IS parameters without shutting down entities. Changes take effect immediately. Modifying the level capability on the global level causes the IS-IS protocol to restart.

The following example displays command usage to modify various parameters:

**Example:**        `config>router>isis# overload timeout 500  
config>router>isis# level-capability level-1/2  
config>router>isis# no authentication-check  
config>router>isis# authentication-key secretkey`

The following example displays the global modifications



```
A:ALA-A>config>router>isis# info

area-id 49.0180.0001
area-id 49.0180.0002
area-id 49.0180.0003
authentication-key "//oZrvtvFPn06S42lRIJsE" hash
authentication-type password
no authentication-check
overload timeout 500 on-boot
level 1
 wide-metrics-only
exit
level 2
 wide-metrics-only
exit
interface "system"
exit
interface "ALA-1-2"
 level-capability level-2
exit
interface "ALA-1-3"
 level-capability level-1
 interface-type point-to-point
exit
interface "ALA-1-5"
 level-capability level-1
 interface-type point-to-point
exit
interface "to-103"
exit
interface "A-B"
exit
interface "A-C"
exit

A:ALA-A>config>router>isis#
```

## 4.6.4 Modifying IS-IS Interface Parameters

You can modify, disable, or remove interface-level IS-IS parameters without shutting down entities. Changes take effect immediately. Modifying the level capability on the interface causes the IS-IS protocol on the interface to restart.

To remove an interface, issue the `no interface ip-int-name` command. To disable an interface, issue the `shutdown` command in the interface context.

The following example displays interface IS-IS modification command usage. For specific interface configuration and modification examples also see, [Example: Configuring a Level 1 Area](#) and [Example: Modifying a Router's Level Capability](#).

**Example:**

```
config>router# isis
config>router>isis# interface ALA-1-3
```

```
config>router>isis>if# passive
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# hello-authentication-type
 message-digest
config>router>isis>if# hello-authentication-key
 secretkey
config>router>isis>if# exit
```

The following example displays the modified interface parameters.

```
A:ALA-A>config>router>isis# info

 area-id 49.0180.0001
 area-id 49.0180.0002
 area-id 49.0180.0003
 authentication-key "//oZrvtvFPn06S42lRIJsE" hash
 authentication-type password
 no authentication-check
 overload timeout 500 on-boot
 level 1
 wide-metrics-only
 exit
 level 2
 wide-metrics-only
 exit
 interface "system"
 exit
 interface "ALA-1-2"
 level-capability level-2
 exit
 interface "ALA-1-3"
 level-capability level-1
 interface-type point-to-point
 passive
 exit
 interface "ALA-1-5"
 level-capability level-1
 interface-type point-to-point
 exit
 interface "to-103"
 hello-authentication-key "DvR3l264KQ6vXMTvbAZ1mE" hash
 hello-authentication-type message-digest
 exit
 interface "A-B"
 exit

A:ALA-A>config>router>isis#
```

---

## 4.6.5 Configuring Authentication using Keychains

The use of authentication mechanism is recommended to protect against malicious attack on the communications between routing protocol neighbors. These attacks could aim to either disrupt communications or to inject incorrect routing information into the systems routing table. The use of authentication keys can help to protect the routing protocols from these types of attacks. In addition, the use of authentication keychains provides the ability to configure authentication keys and make changes to them without affecting the state of the routing protocol adjacencies.

To configure the use of an authentication keychain within IS-IS, use the following steps:

1. Configure an authentication keychain within the **config>system>security** context. The configured keychain must include at least one valid key entry, using a valid authentication algorithm for the IS-IS protocol.
2. Associate the configured authentication keychain with IS-IS. Authentication keychains can be used to specify the authentication at the IS-IS global, and level context as well as for hello authentication at the interface and interface-level context.

The association of the authentication keychain is established through the **auth-keychain** *keychain-name* command at the global and level context. The hello authentication association is established through the **hello-auth-keychain** *keychain-name* command.

For a key entry to be valid, it must include a valid key, the current system clock value must be within the begin and end time of the key entry, and the algorithm specified in the key entry must be supported by the IS-IS protocol.

The IS-IS protocol supports the following algorithms:

- Clear text password (RFC 5304 and RFC 5310 formats)
- HMAC-MD5 (RFC 5304 and RFC 5310 formats)
- HMAC-SHA-1 (RFC 5310 format)
- HMAC-SHA-256 (RFC 5310 format)

The IS-IS key entry may also include the option parameter to determine how the IS-IS protocol encodes the authentication signature. The value of **basic** results in the use of RFC 5304 format. The default or a value of **isis-enhanced** results in using the RFC 5310 format.

Error handling:

- If a keychain exists but there are no active key entries with an authentication type that is valid for the associated protocol then inbound protocol packets will not be authenticated and discarded and no outbound protocol packets should be sent.
- If keychain exists, but the last key entry has expired, a log entry will be raised indicating that all keychain entries have expired. The IS-IS protocol requires that the protocol not revert to an unauthenticated state and requires that the old key is not to be used, therefore, once the last key has expired, all traffic will be discarded.

## 4.6.6 Configuring Leaking

IS-IS allows a two-level hierarchy to route PDUs. Level 1 areas can be interconnected by a contiguous Level 2 backbone. The Level 1 link-state database contains information only about that area. The Level 2 link-state database contains information about the Level 2 system and each of the Level 1 systems in the area. A Level 1/2 router contains information about both Level 1 and Level 2 databases. A Level 1/2 router advertises information about its Level 1 area toward the other Level 1/2 or Level 2 (only) routers.

Packets with destinations outside the Level 1 area are forwarded toward the closest Level 1/2 router which, in turn, forwards the packets to the destination area.

Sometimes, the shortest path to an outside destination is not through the closest Level 1/2 router, or, the only Level 1/2 system to forward packets out of an area is not operational. Route leaking provides a mechanism to leak Level 2 information to Level 1 systems to provide routing information regarding inter-area routes. Then, a Level 1 router has more options to forward packets.

Configure a route policy to leak routes from Level 2 into Level 1 areas in the **config>router>policy-options>policy-statement** context.

The following example shows the command usage to configure prefix list and policy statement parameters in the **config>router** context.

```
config>router>policy-options# prefix-list loops
..>policy-options>prefix-list# prefix 10.1.1.0/24 longer
..>policy-options>prefix-list# exit
..>policy-options# policy-statement leak
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry# from
..>policy-options>policy-statement>entry>from# prefix-list loops
..>policy-options>policy-statement>entry>from# level 2
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
..>policy-options>policy-statement>entry>to# level 1
```

```
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement#exit
..>policy-options# commit
..>policy-options#
```

```
A:ALA-A>config>router>policy-options# info

 prefix-list "loops"
 prefix 10.1.1.0/24 longer
 exit
 policy-statement "leak"
 entry 10
 from
 prefix-list "loop"
 level 2
 exit
 to
 level 1
 exit
 action accept
 exit
 exit
 exit
exit

```

```
A:ALA-A>config>router>policy-options#
```

Next, apply the policy to leak routes from Level 2 info Level 1 systems on ALA-A.

```
config>router#isis
config>router>isis# export leak

A:ALA-A>config>router>isis# info

 area-id 49.0180.0001
 area-id 49.0180.0002
 area-id 49.0180.0003
 authentication-key "//oZrvtvFPn06S42lRIJsE" hash
 authentication-type password
 no authentication-check
 export "leak"
 ...

A:ALA-A>config>router>isis#
```

After the policy is applied, create a policy to redistribute external IS-IS routes from Level 1 systems into the Level 2 backbone (see [Redistributing External IS-IS Routers](#)). In the **config>router** context, configure the following policy statement parameters:

```
config>router>policy-options# begin
..>policy-options# policy-statement "isis-ext"
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry$ from
```

```

..>policy-options>policy-statement>entry>from$ external
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
..>policy-options>policy-statement>entry>to$ level 2
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement# exit
..>policy-options# commit

```

```

A:ALA-A>config>router>policy-options# info

 prefix-list "loops"
 prefix 10.1.1.0/24 longer
 exit
 policy-statement "leak"
 entry 10
 from
 prefix-list "loop"
 level 2
 exit
 to
 level 1
 exit
 action accept
 exit
 exit
 exit
 policy-statement "isis-ext"
 entry 10
 from
 external
 exit
 to
 level 2
 exit
 action accept
 exit
 exit
 exit

A:ALA-A>config>router>policy-options#

```

## 4.6.7 Redistributing External IS-IS Routers

IS-IS does not redistribute Level 1 external routes into Level 2 by default. You must explicitly apply the policy to redistribute external IS-IS routes. Policies are created in the **config>router>policy-options** context. See [Route Policies](#) for more information.

The following example displays the policy statement configuration.

```
config>router>policy-options# info

 prefix-list "loops"
 prefix 10.1.1.0/24 longer
 exit
 policy-statement "leak"
 entry 10
 from
 prefix-list "loop"
 level 2
 exit
 to
 level 1
 exit
 action accept
 exit
 exit
 exit
 policy-statement "isis-ext"
 entry 10
 from
 external
 exit
 to
 level 2
 exit
 action accept
 exit
 exit
 exit
config>router>policy-options#

```

## 4.6.8 Specifying MAC Addresses for All IS-IS Routers

Specify the MAC address to use for all L1 or L2 IS-IS routers. The following example shows how to specify all L1 routers:

**Example:**      `all-l1isis 01-80-C2-00-00-14`

You can also specify the MAC address for all L2 IS-IS routers by using the **all-l2isis** command.





## 4.7 IS-IS Configuration Command Reference

### 4.7.1 Command Hierarchies

#### 4.7.1.1 Configuration Commands

```

config
 — router
 — [no] isis [isis-instance]
 — [no] advertise-passive-only
 — advertise-router-capability {area | as}
 — no advertise-router-capability
 — [no] advertise-tunnel-link
 — all-l1isis ieee-address
 — no all-l1isis
 — all-l2isis ieee-address
 — no all-l2isis
 — [no] area-id area-address
 — auth-keychain keychain-name
 — no auth-keychain
 — [no] authentication-check
 — authentication-key [authentication-key | hash-key] [hash | hash2]
 — no authentication-key
 — authentication-type {password | message-digest}
 — no authentication-type
 — [no] csnp-authentication
 — database-export [identifier id] [bgp-ls-identifier bgp-ls-id]
 — no database-export
 — default-route-tag tag
 — no default-route-tag
 — [no] disable-ldp-sync
 — entropy-label
 — [no] override-tunnel-elc
 — export policy-name [policy-name]
 — no export
 — export-limit number [log percentage]
 — no export-limit
 — [no] graceful-restart
 — [no] helper-disable
 — igp-shortcut
 — [no] shutdown
 — tunnel-next-hop
 — family {ipv4 | ipv6 | srv4 | srv6}
 — resolution {any | disabled | filter | match-family-ip}
 — resolution-filter
 — [no] rsvp
 — [no] sr-te

```

---

```

— [no] hello-authentication
— [no] hello-padding {none | adaptive | loose | strict}
— ignore-attached-bit
— [no] ignore-attached-bit
— [no] ignore-lsp-errors
— [no] ignore-narrow-metric
— [no] iid-tlv-enable
— import policy-name [policy-name]
— no import
— [no] interface ip-init-name
 — [no] adjacency-set id
 — [no] bfd-enable {ipv4 | ipv6} [include-bfd-tlv]
 — csnp-interval seconds
 — no csnp-interval
 — [no] default-instance
 — hello-auth-keychain keychain-name
 — no hello-auth-keychain
 — [no] hello-authentication
 — hello-authentication-key [authentication-key | hash-key] [hash |
 hash2]
 — no hello-authentication-key
 — hello-authentication-type {password | message-digest}
 — no hello-authentication-type
 — hello-padding {none | adaptive | loose | strict}
 — no hello-padding
 — interface-type {broadcast | point-to-point}
 — no interface-type
 — ipv4-adjacency-sid label value
 — no ipv4-adjacency-sid
 — ipv6-adjacency-sid label value
 — no ipv6-adjacency-sid
 — [no] ipv4-multicast-disable
 — ipv4-node-sid index value [clear-n-flag]
 — ipv4-node-sid label value [clear-n-flag]
 — no ipv4-node-sid [clear-n-flag]
 — [no] ipv6-multicast-disable
 — ipv6-node-sid index value [clear-n-flag]
 — ipv6-node-sid label value [clear-n-flag]
 — no ipv6-node-sid
 — [no] ipv6-unicast-disable
 — level {1 | 2}
 — hello-auth-keychain keychain-name
 — no hello-auth-keychain
 — hello-authentication-key [authentication-key | hash-key] [hash
 | hash2]
 — no hello-authentication-key
 — hello-authentication-type {password | message-digest}
 — no hello-authentication-type
 — hello-interval seconds
 — no hello-interval
 — hello-multiplier multiplier
 — no hello-multiplier
 — hello-padding {none | adaptive | loose | strict}
 — no hello-padding

```

---

```

— ipv4-multicast-metric metric
— no ipv4-multicast-metric
— ipv6-multicast-metric metric
— no ipv6-multicast-metric
— ipv6-unicast-metric metric
— no ipv6-unicast-metric
— metric metric
— no metric
— [no] passive
— priority number
— no priority
— sd-offset sd-offset
— no sd-offset
— sf-offset
— no sf-offset
— level-capability {level-1 | level-2 | level-1/2}
— lfa-policy-map route-nh-template template-name
— no lfa-policy-map
— load-balancing-weight [value]
— no load-balancing-weight
— [no] loopfree-alternate-exclude
— lsp-pacing-interval milliseconds
— no lsp-pacing-interval
— mesh-group [value | blocked]
— no mesh-group
— [no] passive
— retransmit-interval seconds
— no retransmit-interval
— [no] shutdown
— [no] sid-protection
— tag tag
— no tag
— ipv4-multicast-routing {native | mt}
— [no] ipv4-multicast-routing
— [no] ipv4-routing
— ipv6-multicast-routing {native | mt}
— [no] ipv6-multicast-routing
— [no] ipv6-routing {native | mt}
— [no] ldp-over-rsvp
— level {1 | 2}
— [no] advertise-router-capability
— auth-keychain name
— no auth-keychain
— authentication-key {authentication-key | hash-key} [{hash | hash2}]
— no authentication-key
— authentication-type {password | message-digest}
— no authentication-type
— [no] bier
— [no] shutdown
— template template-name
— no template
— [no] csnp-authentication
— [no] database-export-exclude
— default-ipv4-multicast-metric ipv4 metric

```

- 
- **no default-ipv4-multicast-metric**
  - **default-ipv6-multicast-metric** *ipv6 metric*
  - **no default-ipv6-multicast-metric**
  - **default-ipv6-unicast-metric** *ipv6 metric*
  - **no default-ipv6-unicast-metric**
  - **default-metric** *ipv4 metric*
  - **no default-metric**
  - **external-preference** *external-preference*
  - **no external-preference**
  - **[no] hello-authentication**
  - **hello-padding** {*none* | *adaptive* | *loose* | *strict*}
  - **no hello-padding**
  - **[no] loopfree-alternate-exclude**
  - **lsp-mtu-size** *size*
  - **no lsp-mtu-size**
  - **preference** *preference*
  - **no preference**
  - **[no] psnp-authentication**
  - **[no] wide-metrics-only**
  - **level-capability** {*level-1* | *level-2* | *level-1/2*}
  - **link-group** *link-group-name*
  - **no link-group**
    - **description** *string*
    - **no description**
    - **level** {*1* | *2*}
      - **ipv4-multicast-metric-offset** *offset-value*
      - **no ipv4-multicast-metric-offset**
      - **ipv6-multicast-metric-offset** *offset-value*
      - **no ipv6-multicast-metric-offset**
      - **ipv4-unicast-metric-offset** *offset-value*
      - **no ipv4-unicast-metric-offset**
      - **ipv6-unicast-metric-offset** *offset-value*
      - **no ipv6-unicast-metric-offset**
  - **[no] member** *interface-name*
    - **oper-members** [*value*]
    - **no oper-members**
    - **revert-members** [*value*]
    - **no revert-members**
  - **[no] loopfree-alternates**
    - **exclude**
      - **prefix-policy** *prefix-policy* [*prefix-policy*]
      - **no prefix-policy**
    - **remote-lfa** [*max-pq-cost value*]
    - **no remote-lfa**
      - **node-protect** [*max-pq-nodes value*]
      - **no node-protect**
    - **ti-lfa** [*max-sr-frr-labels value*]
    - **no ti-lfa**
      - **[no] node-protect**
  - **lsp-lifetime** *seconds*
  - **no lsp-lifetime**
  - **lsp-mtu-size** *size*
  - **no lsp-mtu-size**
  - **lsp-refresh-interval** [*seconds*] [*half-lifetime* {*enable* | *disable*}]

---

```

— no lsp-refresh-interval
— [no] multi-topology
 — [no] ipv4-multicast
 — [no] ipv6-multicast
 — [no] ipv6-unicast
— [no] multicast-import [{both | ipv4 | ipv6}]
— overload [timeout seconds] [max-metric]
— no overload
— [no] overload-export-external
— [no] overload-export-interlevel
— overload-on-boot [timeout seconds] [max-metric]
— no overload-on-boot
— [no] poi-tlv-enable
— [no] prefix-attributes-tlv
— prefix-limit limit [log-only] [threshold percent] [overload-timeout {seconds |
 forever}]
— no prefix-limit
— [no] psnp-authentication
— reference-bandwidth bandwidth-in-kbps
— reference-bandwidth [tbps Tera-bps] [gbps Giga-bps] [mbps Mega-bps]
 [kbps Kilo-bps]
— no reference-bandwidth
— rib-priority high {prefix-list-name | tag tag-value}
— no rib-priority
— router-id router-id
— no router-id
— segment-routing
— no segment-routing
 — adj-sid-hold seconds
 — no adj-sid-hold
 — [no] adjacency-set id
 — family {ipv4 | ipv6}
 — parallel [no-advertise]
 — no parallel
 — sid label value
 — no sid
 — entropy-label {force-disable | enable}
 — no entropy-label
 — export-tunnel-table ldp
 — no export-tunnel-table
 — mapping-server
 — no mapping-server
 — sid-map node-sid {index value [range value]} prefix {{ip-
 address/mask} | {ip-address} {netmask}} [set-flags {s}
 [level {1 | 2 | 1/2}] [clear-n-flag]
 — no sid-map node-sid index value
 — [no] shutdown
 — prefix-sid-range {global | start-label label-value max-index index-
 value}
 — no prefix-sid-range
 — [no] shutdown
 — srlb reserved-label-block-name
 — no srlb
 — tunnel-mtu bytes

```

- **no tunnel-mtu**
- **tunnel-table-pref** *preference*
- **no tunnel-table-pref**
- **[no] shutdown**
- **[no] standard-multi-instance**
- **[no] strict-adjacency-check**
- **summary-address** {*ip-prefix/mask* | *ip-prefix netmask*} [**level**] [**tag tag**]
- **no summary-address** {*ip-prefix/ip-prefix-length* | *ip-prefix netmask*}
- **suppress-attached-bit**
- **no suppress-attached-bit**
- **system-id** *isis-system-id*
- **no system-id**
- **[no] timers**
  - **lsp-wait** *lsp-wait* [**lsp-initial-wait initial-wait**] [**lsp-second-wait second-wait**]
  - **no lsp-wait**
  - **spf-wait** *spf-wait* [**spf-initial-wait initial-wait**] [**spf-second-wait second-wait**]
  - **no spf-wait**
- **[no] traffic-engineering**
- **[no] unicast-import-disable** [{**both** | **ipv4** | **ipv6**}]

## 4.7.2 Command Descriptions

### 4.7.2.1 Generic Commands

#### shutdown

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>isis<br>config>router>isis>igp-shortcut<br>config>router>isis>interface<br>config>router>isis>if>level<br>config>router>isis>level>bier<br>config>router>isis>segment-routing<br>config>router>isis>segm-rtng>mapping-server                                                                                                                                                                                       |
| <b>Description</b> | <p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>The <b>no</b> form of this command administratively enables an entity.</p> |

---

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | shutdown                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Special Cases</b> | <p><b>IS-IS IGP-Shortcut</b> — In the <b>config&gt;router&gt;isis&gt;igp-shortcut</b> context, the command enables or disables IGP shortcuts in the IGP instance.</p> <p><b>IS-IS Interface</b> — In the <b>config&gt;router&gt;isis&gt;interface</b> context, the command disables the IS-IS interface. By default, the IS-IS interface is enabled, <b>no shutdown</b>.</p> <p><b>IS-IS Interface and Level</b> — In the <b>config&gt;router&gt;isis&gt;interface <i>ip-int-name</i>&gt;level</b> context, the command disables the IS-IS interface for the level. By default, the IS-IS interface at the level is enabled, <b>no shutdown</b>.</p> |

### 4.7.2.2 IS-IS Commands

#### isis

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |               |         |                |   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|---------|----------------|---|
| <b>Syntax</b>      | <b>[no] isis</b> [ <i>isis-instance</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |               |         |                |   |
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |               |         |                |   |
| <b>Description</b> | <p>This command creates the context to configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance.</p> <p>The IS-IS protocol instance is enabled with the <b>no shutdown</b> command in the <b>config&gt;router&gt;isis</b> context. Alternatively, the IS-IS protocol instance is disabled with the <b>shutdown</b> command in the <b>config&gt;router&gt;isis</b> context.</p> <p>IS-IS instances are shutdown when created, so that all parameters can be configured prior to the instance being enabled.</p> <p>The <b>no</b> form of the command deletes the IS-IS protocol instance. Deleting the protocol instance removes all configuration parameters for this IS-IS instance.</p> |               |         |                |   |
| <b>Parameters</b>  | <p><i>isis-instance</i> — Specifies the instance ID for an IS-IS instance.</p> <table><tr><td><b>Values</b></td><td>1 to 31</td></tr><tr><td><b>Default</b></td><td>0</td></tr></table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>Values</b> | 1 to 31 | <b>Default</b> | 0 |
| <b>Values</b>      | 1 to 31                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |               |         |                |   |
| <b>Default</b>     | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |               |         |                |   |

#### advertise-passive-only

|                    |                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] advertise-passive-only</b>                                                                    |
| <b>Context</b>     | config>router>isis                                                                                    |
| <b>Description</b> | This command enables and disables IS-IS to advertise only prefixes that belong to passive interfaces. |

---

**Default**    no advertise-passive-only

## advertise-router-capability

**Syntax**    **advertise-router-capability {area | as}**  
              **no advertise-router-capability**

**Context**    config>router>isis

**Description**    This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability.

                  The parameters (area and as) control the scope of the capability advertisements.

                  The **no** form of this command disables this capability.

**Default**    no advertise-router-capability

**Parameters**    **area** — are only advertised within the area of origin.  
                  **as** — are only advertised throughout the entire autonomous system

## advertise-tunnel-link

**Syntax**    **[no] advertise-tunnel-link**

**Context**    config>router>isis

**Description**    This command enables the forwarding adjacency feature. With this feature, IS-IS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs per RFC 3906.

                  The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both **igp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win.

                  When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router-id of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA/LSP advertising the corresponding P2P unnumbered link is installed in the local routing database, then the node performs an SPF using it like any other link LSA/LSP. The link bi-directional check requires that a link, regular link or tunnel link, exists in the reverse direction for the tunnel to be used in SPF.



That the **igp-shortcut** option under the LSP name governs the use of the LSP with both the **igp-shortcut** and the **advertise-tunnel-link** options in IGP. In other words, the user can exclude a specific RSVP LSP from being used as a forwarding adjacency by entering the command **config>router>mpls>lsp>no igp-shortcut**.

The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.

The **no** form of this command disables forwarding adjacency and hence disables the advertisement of RSVP LSP into IGP.

**Default** no advertise-tunnel-link

## all-l1isis

|                    |                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>all-l1isis</b> <i>ieee-address</i><br><b>no all-l1isis</b>                                                                       |
| <b>Context</b>     | config>router>isis                                                                                                                  |
| <b>Description</b> | This command enables you to specify the MAC address to use for all L1 IS-IS routers. The MAC address should be a multicast address. |
| <b>Default</b>     | 01:80:c2:00:00:14                                                                                                                   |
| <b>Parameters</b>  | <i>ieee-address</i> — Specifies the destination MAC address for all L1 I-IS neighbors on the link for this IS-IS instance.          |

## all-l2isis

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>all-l2isis</b> <i>ieee-address</i><br><b>no all-l2isis</b>                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                       |
| <b>Description</b> | This command enables you to specify the MAC address to use for all Layer 2 IS-IS routers. The MAC address should be a multicast address. |
| <b>Default</b>     | 01:80:c2:00:00:15                                                                                                                        |
| <b>Parameters</b>  | <i>ieee-address</i> — Specifies the destination MAC address for all Layer 2 IS-IS neighbors on the link for this IS-IS instance.         |

## area-id

**Syntax** [no] **area-id** *area-address*

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command was previously named the <b>net</b> <i>network-entity-title</i> command. The <b>area-id</b> command allows you to configure the area ID portion of NSAP addresses which identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.</p> <p>A maximum of 3 <b>area addresses</b> can be configured.</p> <p>NSAP addresses are divided into three parts. Only the area ID portion is configurable.</p> <ul style="list-style-type: none"> <li>• Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.</li> <li>• System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.</li> <li>• Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.</li> </ul> <p>The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.</p> <p>For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.</p> <p>For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.</p> <p>For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.</p> <p>If multiple <b>area-id</b> commands are entered, the system ID of all subsequent entries must match the first <b>area</b> address.</p> <p>The <b>no</b> form of the command removes the area address.</p> |
| <b>Default</b>     | no area-id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>area-address</i> — The 1 — 13-byte address. Of the total 20 bytes comprising the NET, only the first 13 bytes can be manually configured. As few as one byte can be entered or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## auth-keychain

|                |                                  |
|----------------|----------------------------------|
| <b>Syntax</b>  | <b>auth-keychain</b> <i>name</i> |
| <b>Context</b> | config>router>isis               |

```
config>router>isis>level
config>service>vprn>isis>
config>service>vprn>isis>level
```

**Description** This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

**Default** no auth-keychain

**Parameters** *name* — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

## authentication-check

**Syntax** [no] authentication-check

**Context** config>router>isis

**Description** This command sets an authentication check to reject PDUs that do not match the type or key requirements.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generate a log event.

**Default** authentication-check

## authentication-key

**Syntax** authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]  
no authentication-key

**Context** config>router>isis  
config>router>isis>level

**Description** This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The [authentication-type](#) statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of the command removes the authentication key.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no authentication-key                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b> | <p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p><i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the <b>hash2</b> encrypted variable cannot be copied and pasted. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> |

## authentication-type

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-type {password   message-digest}</b><br><b>no authentication</b>                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>isis<br>config>router>isis>level                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.</p> <p>Both the authentication key and the authentication type on a segment must match. The <b>authentication-key</b> statement must also be included.</p> <p>Configure the authentication type on the global level in the <b>config&gt;router&gt;isis</b> context.</p> |

Configure or override the global setting by configuring the authentication type in the **config>router>isis>level** context.

The **no** form of the command disables authentication.

|                   |                                                                                                                                                                                                           |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no authentication-type                                                                                                                                                                                    |
| <b>Parameters</b> | <p><b>password</b> — Specifies that simple password (plain text) authentication is required.</p> <p><b>message-digest</b> — Specifies that MD5 authentication in accordance with RFC2104 is required.</p> |

## csnp-authentication

|                    |                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] csnp-authentication</b>                                                                                                                                                                          |
| <b>Context</b>     | config>router>isis<br>config>router>isis>level                                                                                                                                                           |
| <b>Description</b> | <p>This command enables authentication of individual IS-IS packets of complete sequence number PDUs (CSNP) type.</p> <p>The <b>no</b> form of the command suppresses authentication of CSNP packets.</p> |
| <b>Default</b>     | csnp-authentication                                                                                                                                                                                      |

## database-export

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database-export [identifier <i>id</i>] [bgp-ls-identifier <i>bgp-ls-id</i>]</b><br><b>no database-export</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables the population of the extended TE Database (TE-DB) with the link-state information from a given IGP instance.</p> <p>The extended TE-DB is used as a central point for importing all link-state information, link, node, and prefix, from IGP instances on the router or the vSROS controller of the NSP and to exporting it to BGP-LS on the router and to Java-VM proxy on the vSROS controller. This information includes the IGP, TE, and the SR information, prefix SID sub-TLV, adjacency SID sub-TLV, and router SR capability TLV.</p> <p>The <b>no</b> form of the command disables database exportation.</p> |
| <b>Default</b>     | no database-export                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Parameters**    **identifier** — This parameter is used to uniquely identify the IGP instance in the BGP-LS NLRI when a router has interfaces participating in multiple IGP instances. This parameter defaults to the IGP instance ID assigned by SR OS. However, given that the concept of instance ID defined in IS-IS (RFC 6822) is unique within a routing domain while the one specified for OSPF is local subnet significant (RFC 6549), the user can remove any overlap by configuring the new **identifier** value to be unique within a given IGP domain when this router sends the IGP link-state information using BGP-LS.

*id* — Specifies an entry ID to export.

**Values**        0 to 18446744073709551615

**bgp-ls-identifier** — This parameter is used, along with the Autonomous System Number (ASN) to correlate the BGP-LS NLRI advertisements of multiples BGP-LS speakers of the same IGP domain. If an NRC-P network domain has multiple IGP domains, BGP-LS speakers within each IGP domain must be configured with the same unique {bgp-ls-identifier, asn} tuple.

The BGP-LS identifier is optional and is only sent in a BGP-LS NLRI if configured in the IGP instance of an IGP domain.

Note that if this IGP instance participates in traffic engineering with RSVP-TE or SR-TE, the **traffic-engineering** option is not strictly required because enabling the extended TE-DB populates this information automatically. It is, however, recommended to enable it to make the configuration consistent with other routers in the network that do not require the enabling of the extended TE-DB.

*bgp-ls-id* — Specifies a BGP LS ID to export.

**Values**        0 to 4294967295

## default-route-tag

**Syntax**        **default-route-tag** *tag*  
                 **no default-route-tag**

**Context**       config>router>isis

**Description**   This command configures the route tag for default route.

**Default**        no default-route-tag

**Parameters**    *tag* — Assigns a default tag.

**Values**        1 to 4294967295

## disable-ldp-sync

**Syntax**        [no] **disable-ldp-sync**

**Context**       config>router>isis

**Description** This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF or IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different. It will then disable IGP-LDP synchronization for all interfaces. This command does not delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the `ldp-sync-timer` is configured.

**Default** no disable-ldp-sync

## entropy-label

**Syntax** **entropy-label**

**Context** config>router>isis

**Description** This command enables the context for the configuration of entropy label capabilities for the routing protocol.

## override-tunnel-elc

**Syntax** [**no**] **override-tunnel-elc**

**Context** config>router>isis>entropy-label

**Description** This command configures the ability to override any received entropy label capability advertisement. When enabled, the system assumes that all nodes for an IGP domain are capable of receiving and processing the entropy label on segment routed tunnels. This command can only be configured if **entropy-label** is enabled via the **config>router>isis>segment-routing>entropy-label** or **config>router>ospf>segment-routing>entropy-label** command.

The **no** version of this command disables the override. The system assumes entropy label capability for other nodes in the IGP instance if capability advertisements are received.

**Default** no override-tunnel-elc

## export

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

**Context** config>router>isis

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures export routing policies that determine the routes exported from the routing table to IS-IS.</p> <p>If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.</p> <p>If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.</p> <p>If an <b>aggregate</b> command is also configured in the <b>config&gt;router</b> context, then the aggregation is applied before the export policy is applied.</p> <p>Routing policies are created in the <b>config&gt;router&gt;policy-options</b> context.</p> <p>The <b>no</b> form of the command removes the specified <i>policy-name</i> or all policies from the configuration if no <i>policy-name</i> is specified.</p> |
| <b>Default</b>     | no export                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>policy-name</i> — The export policy name. Up to five <i>policy-name</i> arguments can be specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## export-limit

|                    |                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-limit</b> <i>number</i> [ <i>log percentage</i> ]<br><b>no export-limit</b>                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.</p> <p>The <b>no</b> form of the command removes the parameters from the configuration.</p>                              |
| <b>Default</b>     | no export-limit                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.</p> <p><b>Values</b> 1 to 4294967295</p> <p><b>log percentage</b> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p><b>Values</b> 1 to 100</p> |



## graceful-restart

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] graceful-restart</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command enables IS-IS graceful restart (GR) to minimize service interruption. When the control plane of a GR-capable router fails or restarts, the neighboring routers (GR helpers) temporarily preserve IS-IS forwarding information. Traffic continues to be forwarded to the restarting router using the last known forwarding tables. If the control plane of the restarting router becomes operationally and administratively up within the grace period, the restarting router resumes normal IS-IS operation. If the grace period expires, then the restarting router is presumed inactive and the IS-IS topology is recalculated to route traffic around the failure.</p> <p>The <b>no</b> form of the command disables graceful restart and removes the graceful restart configuration from the IS-IS instance.</p> |
| <b>Default</b>     | no graceful-restart                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## helper-disable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] helper-disable</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis>graceful-restart                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command disables helper support for IS-IS graceful restart (GR).</p> <p>When <b>graceful-restart</b> is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart), a restarting router, or both. The router only supports helper mode. It will not act as a restarting router, because the high availability feature set already preserves IS-IS forwarding information so that this functionality is not needed. This command is a historical command and should not be disabled. Configuring <b>helper-disable</b> has the effect of disabling graceful restart, because the router only supports helper mode.</p> <p>The <b>no helper-disable</b> command enables helper support and is the default when graceful restart is enabled.</p> |
| <b>Default</b>     | no helper-disable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## igp-shortcut

|                    |                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>igp-shortcut</b>                                                                                                              |
| <b>Context</b>     | config>router>isis                                                                                                               |
| <b>Description</b> | <p>This command enables the use of an RSVP-TE or SR-TE shortcut for resolving IGP routes by OSPF or IS-IS routing protocols.</p> |

This command instructs IGP to include RSVP LSPs and SR-TE LSPs originating on this node and terminating on the router ID of a remote node as direct links with a metric equal to the metric provided by MPLS.

During the IP reach calculation to determine the reachability of nodes and prefixes, LSPs are overlaid and the LSP metric is used to determine the subset of paths that are equal lowest cost to reach a node or a prefix. If the user enabled the **relative-metric** option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix that is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP-TE or SR-TE LSP, as well as the explicit-null IPv6 label at the bottom of the stack in the case of an IPv6 prefix. Any network event causing one or more IGP shortcuts to go down will trigger a full SPF computation, which may result in installing a new route over an updated set of tunnel next-hops and IP next-hops.

When **igp-shortcut** is enabled at the IGP instance level, all RSVP-TE and SR-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router ID of a remote node. LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IGP. The user can, however, exclude a specific RSVP-TE or SR-TE LSP from being used as a shortcut for resolving IGP routes by entering the **config>router>mpls>lsp>no igp-shortcut** command.

The SPF in IGP only uses RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If two or more options are enabled in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

The SPF in IGP uses SR-TE LSPs as IGP shortcuts only.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of tunnel and IP next-hops to program in the data path.

- Where a destination is a tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop), the tunnel with lowest tunnel-index is selected (the IP next-hop is never used in this case).
- Where a destination is not a tunnel-endpoint:
  - LSPs with metric higher than underlying IGP cost between the endpoint of the LSP are excluded
  - Tunnel next-hops are preferred over IP next-hops
  - Within tunnel next-hops:
    - i. The lowest endpoint-to-destination cost is selected
    - ii. If the endpoint-to-destination costs are the same, the lowest endpoint node router ID is selected
    - iii. If the router IDs are the same, the lowest tunnel index is selected

- Within IP next-hops:
  - i. The lowest downstream router ID is selected
  - ii. If the downstream router IDs are the same, the lowest interface-index is selected



**Note:** Although ECMP is not performed across both the IP and tunnel next-hops, the tunnel endpoint may lie in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

When both RSVP-TE and SR-TE IGP shortcuts are available, the IP reach calculation, in the unicast routing table, will first follow the above ECMP tunnel and IP next-hop selection rules when resolving a prefix over IGP shortcuts. After the set of ECMP tunnel and IP next-hops have been selected, the preference of tunnel type is then applied based on the user setting of the resolution of the family of the prefix. If the user enabled resolution of the prefix family to both RSVP-TE and SR-TE tunnel types, the TTM tunnel preference value is used to select one type for the prefix. In other words, an RSVP-TE LSP type is preferred to an SR-TE LSP type on a per-prefix basis.

The ingress IOM sprays the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the **igp-shortcut** and the **multicast-import** options are enabled in IGP. The unicast RTM can still make use of the tunnel next-hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

## tunnel-next-hop

|                    |                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-next-hop</b>                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis>igp-shortcut                                                                                                                                                                                                       |
| <b>Description</b> | This command enables the context to configure the resolution of IGP IPv4 prefix families, IGP IPv6 prefix families, SR-ISIS IPv4 tunnel families, SR-ISIS IPv6 tunnel families, and SR-OSPF IPv4 tunnel families using IGP shortcuts. |

The **resolution** node is introduced to provide flexibility in the selection of the tunnel types for each of the IP prefix and SR tunnel families.

The IPv4 **family** option causes the IS-IS or OSPF SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv4 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv4 prefixes owned by this node or for which this node is the IPv4 next hop.

The IPv6 **family** option causes the IS-IS or OSPFv3 SPF to include the IPv4 IGP shortcuts in the IP reach calculation of IPv6 nodes and prefixes. RSVP-TE or SR-TE LSPs terminating on a node identified by its router ID can be used to reach IPv6 prefixes owned by this node or for which this node is the IPv6 next-hop. The resolution of IPv6 prefixes is supported in OSPFv3 and in both IS-IS MT=0 and MT=2.

The IS-IS and OSPFv3 IPv6 routes resolved to IPv4 IGP shortcuts are used to:

- forward packets of IS-IS or OSPFv3 prefixes matching these routes
- forward CPM-originated IPv6 packets
- resolve the BGP next hop of BGP IPv6 prefixes
- resolve the indirect next hop of static IPv6 routes

In the data path, a packet for an IPv6 prefix has a label stack that consists of the IPv6 Explicit-Null label value of 2 at the bottom of the label stack followed by the label stack of the IPv4 RSVP-TE LSP.

There is no default behavior for IPv4 prefixes to automatically resolve to RSVP-TE or SR-TE LSPs used as IGP shortcuts by only enabling the **igp-shortcut** context. Instead, the user must enable the **ipv4 family** or **ipv6 family** and set the resolution to the value of **rsvp-te** to select the RSVP-TE tunnel type, or to the value of **sr-te** to select the SR-TE tunnel type.

Setting the **resolution** to the **any** value means that IGP selects the tunnels used as IGP shortcuts according to the TTM preference for the tunnel type. The RSVP-TE LSP type is of higher priority than the SR-TE LSP type.

An IP prefix of family=ipv4 or family= ipv6 always resolves to a single type of tunnel **rsvp-te** or **sr-te**. **Rsvp-te** type is preferred if both types are allowed by the prefix family resolution and both types exist in the set of tunnel next-hops of the prefix. The feature does not support mixing tunnel types per prefix.

If **resolution** for the IPv4 or IPv6 family is set to **disabled**, the corresponding prefixes are resolved to IP next-hops in the multicast routing table.

The **srv4 family** enables the resolution of SR-OSPF IPv4 tunnels and SR-ISIS IPv4 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-OSPF or an SR-ISIS IPv4 tunnel.

The **srv6 family** enables the resolution of SR-ISIS IPv6 tunnels in MT=0 over RSVP-TE IPv4 IGP shortcuts. A maximum of 32 ECMP tunnel next-hops can be programmed for an SR-ISIS IPv6 tunnel.

One or more RSVP-TE LSPs can be selected if **resolution=match-family-ip** and the corresponding IPv4 or IPv6 prefix resolves to RSVP-TE LSPs.



**Note:** An SR tunnel cannot resolve to SR-TE IGP shortcuts.

If **resolution** for the SRv4 or SRv6 tunnel family is set to **disabled**, the corresponding tunnels are resolved to IP next-hops in the multicast routing table.

To enable (disable) IGP shortcuts in the IGP instance, the user must perform a **shutdown** or **no shutdown** in the **igp-shortcut** context.

## family

|                    |                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>family {ipv4   ipv6   srv4   srv6}</b>                                                                                                                                                                                             |
| <b>Context</b>     | config>router>isis>igp-shortcut>tunnel-next-hop                                                                                                                                                                                       |
| <b>Description</b> | This command enables the context to configure the resolution of IGP IPv4 prefix family, IGP IPv6 prefix family, SR-ISIS IPv4 tunnel family, and SR-ISIS IPv6 tunnel family using IGP shortcuts.                                       |
| <b>Parameters</b>  | <p><b>ipv4</b> — Selects the IPv4 address family.</p> <p><b>ipv6</b> — Selects the IPv6 address family.</p> <p><b>srv4</b> — Selects the SR-ISIS IPv4 tunnel family.</p> <p><b>srv6</b> — Selects the SR-ISIS IPv6 tunnel family.</p> |

## resolution

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resolution {any   disabled   filter   match-family-ip}</b>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>isis>igp-shortcut>tunnel-next-hop>family                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command configures resolution mode in the resolution of the IP prefix or SR tunnel family using IGP shortcuts.                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><b>any</b> — Enables the binding to any supported tunnel type following TTM preference.</p> <p><b>disabled</b> — Disables the resolution of the IP prefix or SR tunnel family using IGP shortcuts.</p> <p><b>filter</b> — Enables the binding to the subset of tunnel types configured under <b>resolution-filter</b>.</p> <p><b>match-family-ip</b> — Enables the resolution of the SR tunnel family to match that of the corresponding IP prefix family.</p> |

## resolution-filter

|                |                                                        |
|----------------|--------------------------------------------------------|
| <b>Syntax</b>  | <b>resolution-filter</b>                               |
| <b>Context</b> | config>router>isis>igp-shortcut>tunnel-next-hop>family |

---

|                    |                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enables the context to configure the subset of tunnel types which can be used in the resolution of the IP prefix or SR tunnel family using IGP shortcuts. |
| <b>Parameters</b>  | <b>rsvp</b> — Selects the RSVP-TE tunnel type.<br><b>sr-te</b> — Selects the SR-TE tunnel type.                                                                        |

## rsvp

|                    |                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] rsvp</b>                                                                                                         |
| <b>Context</b>     | config>router>isis>igp-shortcut>tunnel-next-hop>family                                                                   |
| <b>Description</b> | This command selects the RSVP-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts. |

## sr-te

|                    |                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] sr-te</b>                                                                                                      |
| <b>Context</b>     | config>router>isis>igp-shortcut>tunnel-next-hop>family                                                                 |
| <b>Description</b> | This command selects the SR-TE tunnel type in the resolution of the IP prefix or SR tunnel family using IGP shortcuts. |

## hello-authentication

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] hello-authentication</b>                                                                                                                                    |
| <b>Context</b>     | config>router>isis<br>config>router>isis>interface<br>config>router>isis>level                                                                                      |
| <b>Description</b> | This command enables authentication of individual IS-IS packets of HELLO type.<br><br>The <b>no</b> form of the command suppresses authentication of HELLO packets. |
| <b>Default</b>     | hello-authentication                                                                                                                                                |

## hello-padding

|                |                                                                                |
|----------------|--------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>[no] hello-padding {none   adaptive   loose   strict}</b>                   |
| <b>Context</b> | config>router>isis<br>config>router>isis>level<br>config>router>isis>interface |

config>router>isis>interface>level

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables IS-IS Hello (IIH) message padding to ensure that IS-IS LSPs can traverse the link. When this option is enabled, IS-IS Hello messages are padded to the maximum LSP MTU value, which can be set with the <b>lsp-mtu-size</b> command.</p> <p>The <b>no</b> form of the command disables IS-IS hello padding at this level. However, the router may still perform hello padding if it was set at a higher level in the configuration. To ensure that Hello message padding is disabled, set all levels of configuration to <b>no hello-padding</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | no hello-padding                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><b>none</b> — Specifies that the Hello message padding is not enabled at this level, even if it is configured at one of the parent levels.</p> <p><b>adaptive</b> — Specifies the adaptive padding option; this option is able to detect LSP MTU asymmetry from one side of the connection but uses more overhead than loose padding.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in state up. If the implementation supports RFC 3373/5303, “Three-Way Handshake for IS-IS Point-to- Point Adjacencies” then this is when the three-way state is Up. If the implementation use the “classic” algorithm described in ISO 10589, this is when adjacency state is Up. If the neighbor does not support the adjacency state TLV, then padding continues.</li> <li>2. broadcast interface—Padding starts until at least one adjacency is up on the interface.</li> </ol> <p><b>loose</b> — Specifies the loose padding option; the loose padding may not be able to detect certain situations such as asymmetrical LSP MTUs between the routing devices.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state.</li> <li>2. broadcast interface—Padding starts until there is at least one adjacency (broadcast only has up/down) is up on the interface.</li> </ol> <p><b>strict</b> — Specifies the strict padding option; this option is the most overhead-intensive but detects LSP MTU issues on both sides of a link.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—Padding is done for all adjacency states, and is continuous.</li> <li>2. broadcast interface—Padding is done for all adjacency states, and is continuous.</li> </ol> |

## ignore-attached-bit

|                |                                                                          |
|----------------|--------------------------------------------------------------------------|
| <b>Syntax</b>  | <p><b>ignore-attached-bit</b></p> <p><b>[no] ignore-attached-bit</b></p> |
| <b>Context</b> | config>router>isis                                                       |

---

|                    |                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes. |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|

|                |                        |
|----------------|------------------------|
| <b>Default</b> | no ignore-attached-bit |
|----------------|------------------------|

## ignore-lsp-errors

|               |                               |
|---------------|-------------------------------|
| <b>Syntax</b> | <b>[no] ignore-lsp-errors</b> |
|---------------|-------------------------------|

|                |                                               |
|----------------|-----------------------------------------------|
| <b>Context</b> | config>router>isis<br>config>service>vpn>isis |
|----------------|-----------------------------------------------|

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command specifies that IS-IS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged. |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|

The **no** form of the command specifies that IS-IS will not ignore LSP errors.

|                |                      |
|----------------|----------------------|
| <b>Default</b> | no ignore-lsp-errors |
|----------------|----------------------|

## ignore-narrow-metric

|               |                                  |
|---------------|----------------------------------|
| <b>Syntax</b> | <b>[no] ignore-narrow-metric</b> |
|---------------|----------------------------------|

|                |                    |
|----------------|--------------------|
| <b>Context</b> | config>router>isis |
|----------------|--------------------|

|                    |                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command specifies that IS-IS will ignore links with narrow metrics when wide-metrics support has been enabled. |
|--------------------|---------------------------------------------------------------------------------------------------------------------|

The **no** form of the command specifies that IS-IS will not ignore these links.

|                |                         |
|----------------|-------------------------|
| <b>Default</b> | no ignore-narrow-metric |
|----------------|-------------------------|

## iid-tlv-enable

|               |                            |
|---------------|----------------------------|
| <b>Syntax</b> | <b>[no] iid-tlv-enable</b> |
|---------------|----------------------------|

|                |                    |
|----------------|--------------------|
| <b>Context</b> | config>router>isis |
|----------------|--------------------|

|                    |                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enables IS-IS multi-instance (MI) as described in <i>draft-ietf-isis-mi-02</i> . Multiple instances allows the formation of instance-specific adjacencies that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV that identifies the instance and the topology to which the PDU belongs. |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

The **iid-tlv-enable** (based on *draft-ietf-isis-mi-02*) and **standard-multi-instance** (based on *draft-ginsberg-isis-mi-bis-01*) commands cannot be configured in the same instance, because the MAC addresses and PDUs in each standard are incompatible.



The **no** form of the command disables IS-IS MI.

**Default** no iid-tlv-enable

## import

**Syntax** **import** *policy-name* [*policy-name*]  
**no import**

**Context** config>router>isis

**Description** This command specifies up to five route policies as IS-IS import policies.

When a prefix received in an IS-IS LSP is accepted by an entry in an IS-IS import policy, it is installed in the routing table, if it is the most preferred route to the destination.

When a prefix received in an IS-IS LSP is rejected by an entry in an IS-IS import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination.

The flooding of LSPs is unaffected by IS-IS import policy actions.

The **no** form of the command removes all policies from the configuration.

**Default** no import

**Parameters** *policy-name* — Identify the import route policy name. Allowed values are any string up to 32 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified names must already be defined.

## interface

**Syntax** [**no**] **interface** *ip-int-name*

**Context** config>router>isis

**Description** This command creates the context to configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted.

The **no** form of the command removes IS-IS from the interface.

The **shutdown** command in the **config>router>isis>interface** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

---

|                   |                                                                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no interface                                                                                                                                               |
| <b>Parameters</b> | <i>ip-int-name</i> — Identify the IP interface name created in the <b>config&gt;router&gt;interface</b> context. The IP interface name must already exist. |

## bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bfd-enable {ipv4   ipv6} [include-bfd-tlv]</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command enables the use of bi-directional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on an IPv4 or IPv6 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set by the BFD command under the IP interface. This command must be given separately to enable/disable BFD for both IPv4 and IPv6.</p> <p>The <b>no</b> form of this command removes BFD from the associated adjacency.</p> |
| <b>Default</b>     | no bfd-enable ipv4<br>no bfd-enable ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>include-bfd-tlv</i> — Enables support for the IS-IS BFD TLV options, specified in RFC 6213, which specifies that a BFD session must be established before an IS-IS adjacency can transition to the established state. This option should be enabled on all IS-IS neighbors on a shared interface.                                                                                                                                                                                                                                           |

## csnp-interval

|                    |                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>csnp-interval seconds</b><br><b>no csnp-interval</b>                                                                                                                                                                                     |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces.<br>csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.                                                                             |
| <b>Parameters</b>  | <i>seconds</i> — The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer.<br><b>Values</b> 1 to 65535                                                                             |

## default-instance

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] default-instance</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command enables a non-MI capable router to establish an adjacency and operate with a router in a non-zero instance. If the router does not receive IID-TLVs, it will establish an adjacency in a single instance. Instead of establishing an adjacency in the standard instance 0, the router will establish an adjacency in the configured non-zero instance. The router will then operate in the configured non-zero instance so that it appears to be in the standard instance 0 to its neighbor. This feature is supported on point-to-point interfaces, broadcast interfaces are not supported.</p> <p>This feature must be configured on the router connected to non-MI capable routers and on all other SR OS routers in the area, so that they receive non-MI LSPs in the correct instance and not in the base instance.</p> <p>The <b>no</b> form of this command disables the functionality so that the router can only establish adjacencies in the standard instance 0.</p> |
| <b>Default</b>     | no default-instance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## hello-auth-keychain

|                    |                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello-auth-keychain</b> <i>name</i>                                                                                                                                          |
| <b>Context</b>     | config>router>isis<br>config>router>isis>level<br>config>service>vprn>isis>interface<br>config>service>vprn>isis>interface>level                                                |
| <b>Description</b> | This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session. |
| <b>Default</b>     | no hello-auth-keychain                                                                                                                                                          |
| <b>Parameters</b>  | <i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.                                                   |

## hello-authentication

|                    |                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] hello-authentication</b>                                                                                                                            |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                |
| <b>Description</b> | <p>This command enables hello authentication on the interface.</p> <p>The <b>no</b> form of the command disables hello authentication on the interface.</p> |

---

**Default**    hello-authentication

## hello-authentication-key

**Syntax**    **hello-authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]  
**no hello-authentication-key**

**Context**    config>router>isis>interface  
              config>router>isis>if>level

**Description**    This command configures the authentication key (password) for hello PDUs. Neighboring routers use the password to verify the authenticity of hello PDUs sent from this interface. Both the hello authentication key and the hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the hello authentication key in the interface context use the **hello-authentication-key** in the **config>router>isis>interface** context.

To configure or override the hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>interface>level** context.

If both IS-IS and hello-authentication are configured, hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including hello) protocol PDUs.

When the hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

The **no** form of the command removes the authentication-key from the configuration.

**Default**    no hello-authentication-key

**Parameters**    *authentication-key* — The hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

*hash-key* — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

**hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

**hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

## hello-authentication-type

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello-authentication-type</b> {password   message-digest}<br><b>no hello-authentication-type</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis>interface<br>config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command enables hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello <b>authentication-key</b> statement must also be included.</p> <p>To configure the hello authentication type at the interface context, use <b>hello-authentication-type</b> in the <b>config&gt;router&gt;isis&gt;interface</b> context.</p> <p>To configure or override the hello authentication setting for a given level, configure the <b>hello-authentication-type</b> in the <b>config&gt;router&gt;isis&gt;interface&gt;level</b> context.</p> <p>The <b>no</b> form of the command disables hello authentication.</p> |
| <b>Default</b>     | no hello-authentication-type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><b>password</b> — Specifies simple password (plain text) authentication is required.</p> <p><b>message-digest</b> — Specifies MD5 authentication in accordance with RFC 2104 (HMAC: Keyed-Hashing for Message Authentication) is required.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## interface-type

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface-type</b> {broadcast   point-to-point}<br><b>no interface-type</b>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command configures the IS-IS interface type as either broadcast or point-to-point.</p> <p>Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.</p> <p>If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually.</p> |

The **no** form of the command reverts to the default value.

|                      |                                                                                                              |
|----------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | interface-type point-to-point — For IP interfaces on SONET channels.                                         |
|                      | interface-type broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.                |
| <b>Special Cases</b> | <b>SONET</b> — Interfaces on SONET channels default to the point-to-point type.                              |
|                      | <b>Ethernet or Unknown</b> — Physical interfaces that are Ethernet or unknown default to the broadcast type. |
| <b>Parameters</b>    | <b>broadcast</b> — Configures the interface to maintain this link as a broadcast network.                    |
|                      | <b>point-to-point</b> — Configures the interface to maintain this link as a point-to-point link.             |

## ipv4-adjacency-sid

|                    |                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv4-adjacency-sid</b> <i>label value</i><br><b>no ipv4-adjacency-sid</b>                         |
| <b>Context</b>     | config>router>isis>interface                                                                         |
| <b>Description</b> | This command allows a static value to be assigned to an IPv4 adjacency SID in IS-IS segment routing. |
|                    | The <b>label</b> option specifies that the value is assigned to an MPLS label.                       |
|                    | The <b>no</b> form of the command removes the adjacency SID.                                         |
| <b>Parameters</b>  | <b>label</b> <i>value</i> — Specifies the adjacency SID label.                                       |
|                    | <b>Values</b> 18432 to 5248   1048575 (FP4 only)                                                     |

## ipv6-adjacency-sid

|                    |                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-adjacency-sid</b> <i>label value</i><br><b>no ipv6-adjacency-sid</b>                         |
| <b>Context</b>     | config>router>isis>interface                                                                         |
| <b>Description</b> | This command allows a static value to be assigned to an IPv6 adjacency SID in IS-IS segment routing. |
|                    | The <b>label</b> option specifies that the value is assigned to an MPLS label.                       |
|                    | The <b>no</b> form of the command removes the adjacency SID.                                         |
| <b>Parameters</b>  | <b>label</b> <i>value</i> — Specifies the adjacency SID label.                                       |
|                    | <b>Values</b> 18432 to 5248   1048575 (FP4 only)                                                     |

## ipv4-multicast-disable

|                    |                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv4-multicast-disable</b>                                                                                                                                     |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                           |
| <b>Description</b> | This command disables IS-IS IPv4 multicast routing for the interface.<br><br>The <b>no</b> form of the command enables IS-IS IPv4 multicast routing for the interface. |
| <b>Default</b>     | no ipv4-multicast-disable                                                                                                                                              |

## ipv4-node-sid

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv4-node-sid index <i>index-value</i> [clear-n-flag]</b><br><b>ipv4-node-sid label <i>label-value</i> [clear-n-flag]</b><br><b>no ipv4-node-sid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of <b>type loopback</b>. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.</p> <p>The command fails if the network interface is not of <b>type loopback</b> or if the interface is defined in an IES or a VPRN context. Also, assigning the same SID index or label value to the same interface in two different IGP instances is not allowed within the same node.</p> <p>The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required since the index and thus label ranges of the various IGP instance are not allowed to overlap.</p> <p>The <b>clear-n-flag</b> option allows the user to clear the N-flag (node-sid flag) in an IS-IS prefix SID sub-TLV originated for the IPv4 prefix of a loopback interface on the system.</p> <p>By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID, meaning that it belongs to this node only. However, when the user wants to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, you must clear the N-flag to assure interoperability with third party implementations, which may perform a strict check on the receiving end and drop duplicate prefix SID sub-TLVs when the N-flag is set.</p> <p>The SR OS implementation is relaxed on the receiving end and accepts duplicate prefix SIDs with the N-flag set or cleared. SR OS will resolve to the closest owner, or owners if ECMP is configured, of the prefix SID according to its cost.</p> |
| <b>Default</b>     | no ipv4-node-sid                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

---

|                   |                                                              |
|-------------------|--------------------------------------------------------------|
| <b>Parameters</b> | <b>index</b> <i>index-value</i> — Specifies the index value. |
|                   | <b>Values</b> 0 to 4294967295                                |
|                   | <b>label</b> <i>label-value</i> — Specifies the label value. |
|                   | <b>Values</b> 0 to 4294967295                                |
|                   | <b>clear-n-flag</b> — Clears the node SID flag.              |
|                   | <b>Default</b> no clear-n-flag                               |

## ipv6-multicast-disable

|                    |                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6-multicast-disable</b>                                                                                                                                     |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                           |
| <b>Description</b> | This command disables IS-IS IPv6 multicast routing for the interface.<br><br>The <b>no</b> form of the command enables IS-IS IPv6 multicast routing for the interface. |
| <b>Default</b>     | no ipv6-multicast-disable                                                                                                                                              |

## ipv6-node-sid

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-node-sid index</b> <i>index-value</i> [ <b>clear-n-flag</b> ]<br><b>ipv6-node-sid label</b> <i>label-value</i> [ <b>clear-n-flag</b> ]<br><b>no ipv6-node-sid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command assigns a node SID index or label value to the prefix representing the primary address of an IPv6 network interface of type loopback. Only a single node SID can be assigned to an IPv6 interface. When an IPv6 interface has multiple global addresses, the primary address is always the first one in the list, as displayed by the <b>interface info</b> command.<br><br>The above command should fail if the network interface is not of loopback type or if the interface is defined in an IES or a VPRN context. Assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.<br><br>The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required since the index and thus label ranges of the various IGP instance are not allowed to overlap.<br><br>The <b>clear-n-flag</b> option allows the user to clear the N-flag (node-sid flag) in an IS-IS prefix SID sub-TLV originated for the IPv6 prefix of a loopback interface on the system. |



By default, the prefix SID sub-TLV for the prefix of a loopback interface is tagged as a node SID, meaning that it belongs to this node only. However, when the user wants to configure and advertise an anycast SID using the same loopback interface prefix on multiple nodes, you must clear the N-flag to assure interoperability with third-party implementations, which may perform a strict check on the receiving end and drop duplicate prefix SID sub-TLVs when the N-flag is set.

The SR OS implementation is relaxed on the receiving end and accepts duplicate prefix SIDs with the N-flag set or cleared. SR OS will resolve to the closest owner, or owners if ECMP is configured, of the prefix SID according to its cost.

|                   |                                                              |
|-------------------|--------------------------------------------------------------|
| <b>Default</b>    | no ipv6-node-sid                                             |
| <b>Parameters</b> | <b>index</b> <i>index-value</i> — Specifies the index value. |
|                   | <b>Values</b> 0 to 4294967295                                |
|                   | <b>label</b> <i>label-value</i> — Specifies the label value. |
|                   | <b>Values</b> 0 to 4294967295                                |
|                   | <b>clear-n-flag</b> — Clears the node SID flag.              |
|                   | <b>Default</b> no clear-n-flag                               |

## ipv6-unicast-disable

|                    |                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6-unicast-disable</b>                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command disables IS-IS IPv6 unicast routing for the interface.</p> <p>By default IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the <b>config&gt;router&gt;isis&gt;ipv6-routing mt</b> command is configured.</p> <p>The <b>no</b> form of the command enables IS-IS IPv6 unicast routing for the interface.</p> |
| <b>Default</b>     | no ipv6-unicast-disable                                                                                                                                                                                                                                                                                                                                                       |

## level

|                    |                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>level {1   2}</b>                                                                    |
| <b>Context</b>     | config>router>isis<br>config>router>isis>interface                                      |
| <b>Description</b> | This command creates the context to configure IS-IS Level 1 or Level 2 area attributes. |

A router can be configured as a Level 1, Level 2, or Level 1-2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies will not be established over this interface.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

**CLI Syntax:**

```
level>no hello-authentication-key
level>no hello-authentication-type
level>no hello-interval
level>no hello-multiplier
level>no metric
level>no passive
level>no priority
```

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | level 1 or level 2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Special Cases</b> | <p><b>Global IS-IS Level</b> — The <b>config&gt;router&gt;isis</b> context configures default global parameters for both Level 1 and Level 2 interfaces.</p> <p><b>IS-IS Interface Level</b> — The <b>config&gt;router&gt;isis&gt;interface</b> context configures IS-IS operational characteristics of the interface at Level 1 and/or Level 2. A logical interface can be configured on one Level 1 and one Level 2. In this case, each level can be configured independently and parameters must be removed independently. By default an interface operates in both Level 1 and Level 2 modes.</p> |
| <b>Parameters</b>    | <p>1 — Specifies the IS-IS operational characteristics of the interface at level 1.</p> <p>2 — Specifies the IS-IS operational characteristics of the interface at level 2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |

## advertise-router-capability

|                    |                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] advertise-router-capability</b>                                                                                                             |
| <b>Context</b>     | config>router>isis>level                                                                                                                            |
| <b>Description</b> | <p>This command enables router advertisement capabilities.</p> <p>The <b>no</b> form of the command disables router advertisement capabilities.</p> |
| <b>Default</b>     | advertise-router-capability                                                                                                                         |

## bier

|                    |                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bier</b>                                                                                             |
| <b>Context</b>     | config>router>isis>level                                                                                     |
| <b>Description</b> | This command enables BIER capabilities.<br><br>The <b>no</b> form of the command disables BIER capabilities. |

## template

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>template</b> <i>template-name</i><br><b>no template</b>                                                                               |
| <b>Context</b>     | config>router>isis>level>bier                                                                                                            |
| <b>Description</b> | This command assigns a BIER template to an IS-IS level.<br><br>The <b>no</b> form of the command removes templates from the IS-IS level. |
| <b>Parameters</b>  | <i>template-name</i> — Specifies the BIER template name.                                                                                 |

## database-export-exclude

|                    |                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] database-export-exclude</b>                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>isis>level                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command allows the user to prune the IGP link-state information of a specific IS-IS level from being exported into the extended TE-DB.<br><br>The no form of the command returns to the default behavior inherited from the database-export command at the IS-IS instance level. |
| <b>Default</b>     | no database-export-exclude                                                                                                                                                                                                                                                            |

## default-ipv4-multicast-metric

|                    |                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-ipv4-multicast-metric</b> <i>metric</i><br><b>no default-ipv4-multicast-metric</b>                       |
| <b>Context</b>     | config>router>isis>level                                                                                            |
| <b>Description</b> | This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3). |

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

|                   |                                                                                                  |
|-------------------|--------------------------------------------------------------------------------------------------|
| <b>Default</b>    | default-ipv4-multicast-metric 10                                                                 |
| <b>Parameters</b> | <i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT3) |
| <b>Values</b>     | 1 to 16777215                                                                                    |

## default-ipv6-multicast-metric

|                    |                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-ipv6-multicast-metric</b> <i>metric</i><br><b>no default-ipv6-multicast-metric</b>                                                                                                                                                     |
| <b>Context</b>     | config>router>isis>level                                                                                                                                                                                                                          |
| <b>Description</b> | This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4).<br><br>The <b>no</b> form of this command deletes the specified default metric and reverts to using the system default of 10. |
| <b>Default</b>     | default-ipv6-multicast-metric 10                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4).<br>1 to 16777215                                                                                                                                |

## default-ipv6-unicast-metric

|                    |                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>default-ipv6-unicast-metric</b> <i>ipv6 metric</i><br><b>no default-ipv6-unicast-metric</b>     |
| <b>Context</b>     | config>router>isis>level                                                                           |
| <b>Description</b> | This command specifies the default metric for IPv6 unicast.                                        |
| <b>Default</b>     | default-ipv6-unicast-metric 10                                                                     |
| <b>Parameters</b>  | <i>ipv6-metric</i> — Specifies the default metric for IPv6 unicast.<br><b>Values</b> 1 to 16777215 |

## default-metric

|               |                                                                      |
|---------------|----------------------------------------------------------------------|
| <b>Syntax</b> | <b>default-metric</b> <i>ipv4 metric</i><br><b>no default-metric</b> |
|---------------|----------------------------------------------------------------------|

---

|                    |                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>isis>level                                                                                                                                               |
| <b>Description</b> | This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface. |
| <b>Default</b>     | default-metric 10                                                                                                                                                      |
| <b>Parameters</b>  | <i>ipv4 metric</i> — Specifies the default metric for IPv4 unicast.                                                                                                    |
| <b>Values</b>      | 1 to 16777214                                                                                                                                                          |

## external-preference

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>external-preference</b> <i>preference</i><br><b>no external-preference</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis>level                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command configures the external route preference for the IS-IS level.</p> <p>The <b>external-preference</b> command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below.</p> <p>A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.</p> <p>Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the <b>ecmp</b> in the <b>config&gt;router</b> context.</p> |
| <b>Default</b>     | <p>external-preference (Level 1) — 160</p> <p>external-preference (Level 2) — 165</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>preference</i> — The preference for external routes at this level as expressed.</p> <p>Default preferences are listed in the <a href="#">Table 34</a>:</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 34**      **Default External Route Preferences**

| Route Type           | Preference | Configurable |
|----------------------|------------|--------------|
| Direct attached      | 0          | No           |
| Static-route         | 5          | Yes          |
| OSPF internal routes | 10         | No           |

**Table 34** Default External Route Preferences (Continued)

| Route Type             | Preference | Configurable     |
|------------------------|------------|------------------|
| IS-IS Level 1 internal | 15         | Yes <sup>1</sup> |
| IS-IS Level 2 internal | 18         | Yes <sup>1</sup> |
| OSPF external          | 150        | Yes              |
| IS-IS Level 1 external | 160        | Yes              |
| IS-IS Level 2 external | 165        | Yes              |
| BGP                    | 170        | Yes              |

Note:

1. Internal preferences are changed using the preference command in the config>router>isis>level *level-number* context.

**Values** 1 to 255

## hello-interval

**Syntax** **hello-interval** *seconds*  
**no hello-interval**

**Context** config>router>isis>if>level

**Description** This command configures the interval between IS-IS Hello PDUs issued on the interface at this level. The **hello-interval**, along with the **hello-multiplier**, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.



**Note:** The neighbor hold time is (hello multiplier × hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier × hello interval / 3) on designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be > 3 to reduce routing instability.

The **no** form of the command reverts to the default value.

**Default** 3 – for designated intermediate system interfaces  
9 – for non-designated intermediate system interfaces and point-to-point interfaces

**Parameters** *seconds* — The hello interval in seconds expressed as a decimal integer.

**Values** 1 to 20000

## hello-multiplier

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello-multiplier</b> <i>multiplier</i><br><b>no hello-multiplier</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures a hello multiplier. The <b>hello-multiplier</b>, along with the <b>hello-interval</b>, is used to calculate a hold time, which is communicated to a neighbor in a Hello PDU.</p> <p>The hold time is the time in which the neighbor expects to receive the next Hello PDU. If the neighbor receives a Hello within this time, the hold time is reset. If the neighbor does not receive a Hello within the hold time, it brings the adjacency down.</p> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p><b>Note:</b> The neighbor hold time is (hello multiplier × hello interval) on non-designated intermediate system broadcast interfaces and point-to-point interfaces and is (hello multiplier × hello interval / 3) on designated intermediate system broadcast interfaces. Hello values can be adjusted for faster convergence, but the hold time should always be &gt; 3 to reduce routing instability.</p> </div> </div> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | hello-multiplier 3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer.</p> <p><b>Values</b>      2 to 100</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## ipv4-multicast-metric

|                    |                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv4-multicast-metric</b> <i>metric</i><br><b>no ipv4-multicast-metric</b>                                                                                      |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                        |
| <b>Description</b> | <p>This command configures the IS-IS interface metric for IPv4 multicast.</p> <p>The <b>no</b> form of this command removes the metric from the configuration.</p> |
| <b>Default</b>     | no ipv4-multicast-metric                                                                                                                                           |
| <b>Parameters</b>  | <p><i>metric</i> — Specifies the IS-IS interface metric for IPv4 multicast.</p> <p><b>Values</b>      1 to 16777215</p>                                            |

---

## ipv6-multicast-metric

|                    |                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-multicast-metric</b> <i>metric</i><br><b>no ipv6-multicast-metric</b>                                                                               |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                 |
| <b>Description</b> | This command configures the IS-IS interface metric for IPv6 multicast.<br><br>The <b>no</b> form of this command removes the metric from the configuration. |
| <b>Default</b>     | no ipv6-multicast-metric                                                                                                                                    |
| <b>Parameters</b>  | <i>metric</i> — Specifies the IS-IS interface metric for IPv6 multicast.<br><b>Values</b> 1 to 16777215                                                     |

## ipv6-unicast-metric

|                    |                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-unicast-metric</b> <i>metric</i><br><b>no ipv6-unicast-metric</b>                                                                                 |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                               |
| <b>Description</b> | This command configures the IS-IS interface metric for IPv6 unicast.<br><br>The <b>no</b> form of this command removes the metric from the configuration. |
| <b>Default</b>     | no ipv6-unicast-metric                                                                                                                                    |
| <b>Parameters</b>  | <i>metric</i> — Specifies the IS-IS interface metric for IPv6 unicast.<br><b>Values</b> 1 to 16777215                                                     |

## ipv4-multicast-metric-offset

|                    |                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv4-multicast-metric-offset</b> <i>offset-value</i><br><b>no ipv4-multicast-metric-offset</b>                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>isis>link-group>level                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology.<br><br>The <b>no</b> form of the command reverts the offset value to 0. |
| <b>Default</b>     | no ipv4-multicast-metric-offset                                                                                                                                                                                                                                                                                        |



---

**Parameters**    *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

**Values**        0 to 6777215

## ipv6-multicast-metric-offset

**Syntax**        **ipv6-multicast-metric-offset** *offset-value*  
**no ipv6-multicast-metric-offset**

**Context**        config>router>isis>link-group>level

**Description**    This command sets the offset value for the IPv6 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv6 multicast topology.

The **no** form of the command reverts the offset value to 0.

**Default**        no ipv6-multicast-metric-offset

**Parameters**    *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

**Values**        0 to 6777215

## lsp-mtu-size

**Syntax**        **lsp-mtu-size** *size*  
**no lsp-mtu-size**

**Context**        config>router>isis  
config>router>isis>level

**Description**    This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then IS-IS must be restarted in order for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context.



**Note:** Using the **exec** command to execute a configuration file to change the LSP MTU-size from its default value will automatically restart IS-IS for the change to take effect.

The **no** form of the command reverts to the default value.

**Default**        lsp-mtu-size 1492

**Parameters** *size* — Specifies the LSP MTU size.

**Values** 490 to 9190

## preference

**Syntax** **preference** *preference*  
**no preference**

**Context** config>router>isis>level

**Description** This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

**Default** preference (Level 1) — 15

preference (Level 2) — 18

**Parameters** *preference* — The preference for external routes at this level expressed as a decimal integer.

Default preferences are listed in [Table 35](#):

**Table 35 Default Internal Route Preferences**

| Route Type             | Preference | Configurable     |
|------------------------|------------|------------------|
| Direct attached        | 0          | No               |
| Static-route           | 5          | Yes              |
| OSPF internal routes   | 10         | No               |
| IS-IS level 1 internal | 15         | Yes              |
| IS-IS level 2 internal | 18         | Yes              |
| OSPF external          | 150        | Yes              |
| IS-IS level 1 external | 160        | Yes <sup>1</sup> |
| IS-IS level 2 external | 165        | Yes <sup>1</sup> |

**Table 35 Default Internal Route Preferences (Continued)**

| Route Type | Preference | Configurable |
|------------|------------|--------------|
| BGP        | 170        | Yes          |

Note:

- External preferences are changed using the external-preference command in the `config>router>isis>level level-number` context.

**Values** 1 to 255

## metric

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>metric</b> <i>metric</i><br><b>no metric</b>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the metric used for the level on the interface.</p> <p>In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.</p> <p>If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | metric 10                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <p><i>metric</i> — The metric assigned for this level on this interface.</p> <p><b>Values</b> 1 to 16777215</p>                                                                                                                                                                                                                                                                                                                                                   |

## passive

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] passive</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>isis>if<br>config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.</p> <p>When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.</p> |

The **no** form of the command removes the passive attribute.

|                      |                                                                                                                                                                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | no passive                                                                                                                                                                                                                                |
| <b>Special Cases</b> | <p><b>Service Interfaces</b> — Service interfaces (defined using the service-prefix command in <b>config&gt;router</b>) are passive by default.</p> <p><b>All other Interfaces</b> — All other interfaces are not passive by default.</p> |

## priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>priority</b> <i>number</i></p> <p><b>no priority</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.</p> <p>This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | priority 64                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><i>number</i> — The priority for this interface at this level.</p> <p><b>Values</b>      0 to 127</p>                                                                                                                                                                                                                                                                                                                                                                                                              |

## psnp-authentication

|                    |                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] psnp-authentication</b>                                                                                                                                                                        |
| <b>Context</b>     | <p>config&gt;router&gt;isis</p> <p>config&gt;router&gt;isis&gt;level</p>                                                                                                                               |
| <b>Description</b> | <p>This command enables authentication of individual IS-IS packets of partial sequence number PDU (PSNP) type.</p> <p>The <b>no</b> form of the command suppresses authentication of PSNP packets.</p> |
| <b>Default</b>     | psnp-authentication                                                                                                                                                                                    |

## sd-offset

|                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sd-offset</b> <i>sd-offset</i><br><b>no sd-offset</b>                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>If the pre-FEC error rate of the associated DWDM port crosses the configured <b>sd-threshold</b>, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the <b>sd-threshold</b> value is configured under that port.</p> <p>The <b>no</b> form of the command reverts the offset value to 0.</p> |
| <b>Default</b>     | no sd-offset                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>sd-offset</i> — Specifies the amount the interface metric is increased by if the <b>sd-threshold</b> is crossed.</p> <p><b>Values</b> 0 to 16777215</p>                                                                                                                                                                                                                                |

## sf-offset

|                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sf-offset</b> <i>offset-value</i><br><b>no sf-offset</b>                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>isis>if>level                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>If the pre-FEC error rate of the associated DWDM port crosses the configured <b>sf-threshold</b>, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the <b>sf-threshold</b> value is configured under that port.</p> <p>The <b>no</b> form of the command reverts the offset value to 0.</p> |
| <b>Default</b>     | no sf-offset                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>offset-value</i> — Specifies the amount the interface metric is increased by if the <b>sf-threshold</b> is crossed.</p> <p><b>Values</b> 0 to 16777215</p>                                                                                                                                                                                                                             |

## wide-metrics-only

|                |                               |
|----------------|-------------------------------|
| <b>Syntax</b>  | <b>[no] wide-metrics-only</b> |
| <b>Context</b> | config>router>isis>level      |

**Description** This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When wide-metrics-only is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

The **no** form of the command reverts to the default value.

**Default** no wide-metrics-only

## level-capability

**Syntax** **level-capability {level-1 | level-2 | level-1/2}**  
**no level-capability**

**Context** config>router>isis  
config>router>isis>interface

**Description** This command configures the routing level for an instance of the IS-IS routing process.

An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 *and* 2.

[Table 36](#) displays configuration combinations and the potential adjacencies that can be formed.

**Table 36 Potential Adjacency**

| Global Level | Interface Level | Potential Adjacency    |
|--------------|-----------------|------------------------|
| L 1/2        | L 1/2           | Level 1 and/or Level 2 |
| L 1/2        | L 1             | Level 1 only           |
| L 1/2        | L 2             | Level 2 only           |
| L 2          | L 1/2           | Level 2 only           |
| L 2          | L 2             | Level 2 only           |
| L 2          | L 1             | none                   |
| L 1          | L 1/2           | Level 1 only           |
| L 1          | L 2             | none                   |
| L 1          | L 1             | Level 1 only           |

The **no** form of the command removes the level capability from the configuration.

---

|                      |                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | level-capability level-1/2                                                                                                                                                                                                                                                                                                                           |
| <b>Special Cases</b> | <p><b>IS-IS Router</b> — In the <b>config&gt;router&gt;isis</b> context, changing the <b>level-capability</b> performs a restart on the IS-IS protocol instance.</p> <p><b>IS-IS Interface</b> — In the <b>config&gt;router&gt;isis&gt;interface</b> context, changing the <b>level-capability</b> performs a restart of IS-IS on the interface.</p> |
| <b>Parameters</b>    | <p><b>level-1</b> — Specifies the router/interface can operate at Level 1 only.</p> <p><b>level-2</b> — Specifies the router/interface can operate at Level 2 only.</p> <p><b>level-1/2</b> — Specifies the router/interface can operate at both Level 1 and Level 2.</p>                                                                            |

## lfa-policy-map

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lfa-policy-map route-nh-template</b> <i>template-name</i><br><b>no lfa-policy-map</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command applies a route next-hop policy template to an OSPF or IS-IS interface.</p> <p>When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.</p> <p>If the user excluded the interface from LFA using the command <b>loopfree-alternate-exclude</b>, the LFA policy, if applied to the interface, has no effect.</p> <p>Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.</p> <p>The <b>no</b> form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.</p> |
| <b>Default</b>     | no lfa-policy-map                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>template-name</i> — Specifies the name of the template, up to 32 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## load-balancing-weight

|                |                                                                                  |
|----------------|----------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>load-balancing-weight</b> [ <i>value</i> ]<br><b>no load-balancing-weight</b> |
| <b>Context</b> | config>router>isis>interface                                                     |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures the weighted ECMP load-balancing weight for an IS-IS interface. If the interface becomes an ECMP next hop for an IPv4 or IPv6 route, and all the other ECMP next hops are interfaces with configured (non-zero) load-balancing weights, then the traffic distribution over the ECMP interfaces is proportional to the weights. In other words, the interface with the largest load-balancing weight should receive the most forwarded traffic if weighted ECMP is applicable.</p> <p>The <b>no</b> form of this command disables weighted ECMP for the interface and therefore effectively disables weighted ECMP for any IP prefix that has this interface as a next hop.</p> |
| <b>Default</b>     | no load-balancing-weight                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>value</i> — 0 to 4294967295                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## loopfree-alternate-exclude

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] loopfree-alternate-exclude</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>isis>level<br>config>router>isis>interface<br>config>service>vprn>isis>level<br>config>service>vprn>isis>interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.</p> <p>When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.</p> <p>The <b>no</b> form of this command re-instates the default value for this command.</p> |
| <b>Default</b>     | no loopfree-alternate-exclude                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## lsp-pacing-interval

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp-pacing-interval</b> <i>milliseconds</i><br><b>no lsp-pacing-interval</b> |
| <b>Context</b>     | config>router>isis>interface                                                    |
| <b>Description</b> | This command configures the interval at which LSPs are sent from the interface. |



To avoid overwhelming neighbors that have less CPU processing power with LSPs, the pacing interval can be configured to limit how many LSPs are sent at the interval. LSPs are sent in bursts at the interval up to the configured limit. If a value of 0 is configured, no LSPs are sent from the interface. The interval applies to all LSPs: LSPs generated by the router, and LSPs received from other routers.

The **no** form of the command reverts to the default value.



**Note:** The IS-IS pacing interval is 100 milliseconds for values < 100 milliseconds, and 1 second for values ≥ 100 milliseconds. For example, a pacing interval of 2 milliseconds means that a maximum of 50 LSPs are sent in a burst at 100 millisecond intervals. The default pacing interval of 100 milliseconds means that a maximum of 10 LSPs are sent in a burst at 1 second intervals.

**Default**      lsp-pacing-interval 100 — LSPs are sent in 100 millisecond intervals.

**Parameters**    *milli-seconds* — Specifies the interval in milliseconds during which IS-IS LSPs are sent from the interface expressed as a decimal integer.

**Values**        0 to 65535

## mesh-group

**Syntax**        **mesh-group {value | blocked}**  
                 **no mesh-group**

**Context**        config>router>isis>interface

**Description**    This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to create isolated islands that do not receive updates as (other) links fail.

The **no** form of the command removes the interface from the mesh group.

**Default**        no mesh-group

---

**Parameters**    *value* — The unique decimal integer value distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.

**Values**        1 to 2000000000

**blocked** — Prevents an interface from flooding LSPs.

## retransmit-interval

**Syntax**        **retransmit-interval** *seconds*  
**no retransmit-interval**

**Context**        config>router>isis>interface

**Description**    This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

The **no** form of the command reverts to the default value.

**Default**        retransmit-interval 5

**Parameters**    *seconds* — The interval in seconds that IS-IS LSPs can be sent on the interface.

**Values**        1 to 65535

## sid-protection

**Syntax**        [**no**] **sid-protection**

**Context**        config>router>isis>interface

**Description**    This command enables or disables adjacency SID protection by LFA and remote LFA.

While LFA and remote LFA Fast-Reroute (FRR) protection is enabled for all node SIDs and local adjacency SIDs when the user enables the **loopfree-alternate** option in IS-IS or OSPF at the LER and LSR, there are applications where the user wants traffic to never divert from the strict hop computed by CSPF for a SR-TE LSP. In that case, the user can disable protection for all adjacency SIDs formed over a given network IP interface using this command.

The protection state of an adjacency SID is advertised in the B-FLAG of the IS-IS or OSPF Adjacency SID sub-TLV.

**Default**        sid-protection

## tag

**Syntax**        **tag** *tag*

## no tag

|                    |                                                                                  |
|--------------------|----------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>isis>interface                                                     |
| <b>Description</b> | This command configures a route tag to the specified IP address of an interface. |
| <b>Default</b>     | no tag                                                                           |
| <b>Parameters</b>  | <i>tag</i> — Assigns a route tag.                                                |
| <b>Values</b>      | 1 to 4294967295                                                                  |

## ipv4-multicast-routing

|                    |                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv4-multicast-routing {native   mt}</b><br><b>[no] ipv4-multicast-routing</b>                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM.</p> <p>The <b>no</b> ipv4-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.</p> |
| <b>Default</b>     | ipv4-multicast-routing native                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><b>native</b> — Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.</p> <p><b>mt</b> — Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.</p>                                                                                                                                                                    |

## ipv4-routing

|                    |                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv4-routing</b>                                                                                                                               |
| <b>Context</b>     | config>router>isis                                                                                                                                     |
| <b>Description</b> | <p>This command specifies whether this IS-IS instance supports IPv4.</p> <p>The <b>no</b> form of the command disables IPv4 on the IS-IS instance.</p> |
| <b>Default</b>     | ipv4-routing                                                                                                                                           |

---

## ipv6-multicast-routing

|                    |                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-multicast-routing</b> { <b>native</b>   <b>mt</b> }<br><b>[no] ipv6-multicast-routing</b>                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv6 multicast RTM.</p> <p>The <b>no</b> ipv6-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.</p> |
| <b>Default</b>     | ipv6-multicast-routing native                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><b>native</b> — Causes IPv6 routes from the MT0 topology to be added to the multicast RTM for RPF checks.</p> <p><b>mt</b> — Causes IPv6 routes from the MT3 topology to be added to the multicast RTM for RPF checks.</p>                                                                                                                                                                    |

## ipv6-routing

|                    |                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6-routing</b> { <b>native</b>   <b>mt</b> }                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command enables IPv6 routing.</p> <p>The <b>no</b> form of the command disables support for IS-IS IPv6 TLVs for IPv6 routing.</p>                                                                                                                        |
| <b>Default</b>     | no ipv6-routing                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><b>native</b> — Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.</p> <p><b>mt</b> — Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.</p> |

## ldp-over-rsvp

|                    |                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ldp-over-rsvp</b>                                                                                                                 |
| <b>Context</b>     | config>router>isis                                                                                                                        |
| <b>Description</b> | <p>This command allows LDP over RSVP processing in IS-IS.</p> <p>The <b>no</b> form of the command disables LDP over RSVP processing.</p> |

---

**Default** no ldp-over-rsvp

## link-group

**Syntax** **link-group** *link-group-name*  
**no link-group**

**Context** config>router>isis

**Description** This command specifies the IS-IS link group associated with this particular level of the interface.

**Parameters** *link-group-name* — Specifies an IS-IS link group on the system up to 32 characters in length.

## description

**Syntax** **description** *string*  
**no description**

**Context** config>router>isis>link-group

**Description** This command adds a description string to the associated link-group. The string can be up to 256 characters long and can only contain printable characters. If the command is issued in the context of a link-group that already contains a description then the previous description string is replaced.

The **no** form of the command removes the description from the associated link-group.

**Default** no description

**Parameters** *string* — Character string to be associated with the associated link-group.

## member

**Syntax** [**no**] **member** *interface-name*

**Context** config>router>isis>link-group>level

**Description** This command adds or removes a link to the associated link-group. The interface name should already exist before it is added to a link-group.

The **no** form of the command removes the specified interface from the associated link-group.

**Parameters** *interface-name* — Name of the interface to be added or removed from the associated link-group.

---

## oper-members

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>oper-members</b> [ <i>value</i> ]<br><b>no oper-members</b>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis>link-group>level                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3.</p> <p>The <b>no</b> form of the command reverts the <b>oper-members</b> limit to 1.</p> |
| <b>Default</b>     | oper-members 1                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>value</i> — Specifies the threshold for operational members.<br><b>Values</b> 1 to 8                                                                                                                                                                                                                                                                                                                                                |

## revert-members

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>revert-members</b> [ <i>value</i> ]<br><b>no revert-members</b>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>isis>link-group>level                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command sets the threshold for the minimum number of operational links to return the associated link-group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured revert-member threshold then the configured offsets are removed.</p> <p>The <b>no</b> form of the command reverts the revert-members threshold back to the default which is equal to the oper-member threshold value.</p> |
| <b>Default</b>     | revert-members <i>oper-members</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>value</i> — Specifies the threshold for revertive members.<br><b>Values</b> 1 to 8                                                                                                                                                                                                                                                                                                                                                                                                                            |

## ipv4-unicast-metric-offset

|                |                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>ipv4-unicast-metric-offset</b> <i>offset-value</i><br><b>no ipv4-unicast-metric-offset</b> |
| <b>Context</b> | config>router>isis>link-group>level                                                           |

---

|                    |                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric.</p> <p>The <b>no</b> form of the command reverts the offset value to 0.</p> |
| <b>Default</b>     | no ipv4-unicast-metric-offset                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.</p> <p><b>Values</b>      0 to 6777215</p>                |

## ipv6-unicast-metric-offset

|                    |                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>ipv6-unicast-metric-offset</b> <i>offset-value</i></p> <p><b>no ipv6-unicast-metric-offset</b></p>                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>isis>link-group>level                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the <b>oper-members</b> threshold, the configured offset is applied to the interface metric for the IPv6 topology.</p> <p>The <b>no</b> form of the command reverts the offset value to 0.</p> |
| <b>Default</b>     | no ipv6-unicast-metric-offset                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.</p> <p><b>Values</b>      0 to 6777215</p>                                             |

## loopfree-alternates

|                    |                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] loopfree-alternates</b>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command enables Loop-Free Alternate (LFA) computation by SPF for the IS-IS routing protocol.</p> <p>When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary nexthop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.</p> |

The user enables the remote LFA next-hop calculation by the IGP LFA SPF by appending the `remote-lfa` option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when the latter resulted in no protection for one or more prefixes which are resolved to a given interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing/tearing-down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node which puts the packets back into the shortest without looping them back to the node which forwarded them over the repair tunnel. The remote LFA node is referred to as PQ node. A repair tunnel can in theory be an RSVP LSP, a LDP-in-LDP tunnel, or a SR tunnel. In this feature, it is restricted to use SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix like the regular LFA one. So, it provides protection for all destination prefixes which share the protected link by using the neighbor on the other side of the protected link as a proxy for all these destinations.

The Topology-Independent LFA (TI-LFA) further improves the protection coverage of a network topology by computing and automatically instantiating a repair tunnel to a Q node which is not in shortest path from the computing node. The repair tunnel uses shortest path to the P node and a source routed path from the P node to the Q node.

In addition, the TI-LFA algorithm selects the backup path which matches the post-convergence path. This helps the capacity planning in the network since traffic will always flow on the same path when transitioning to the FRR next-hop and then onto the new primary next-hop.

At a high level, the TI-LFA protection algorithm is searching for a candidate P-Q set separated with a number of hops such that the label stack size does not exceed the value of **ti-lfa max-sr-frr-labels**, on each of the post-convergence paths to each destination node or prefix D.

When the **ti-lfa** option is enabled in IS-IS, it provides TI-LFA node-protect or link-protect backup path in IS-IS MT=0 for an SR-ISIS IPV4/IPv6 tunnel (node SID and adjacency SID), for an IPv4 SR-TE LSP, and for LDP IPv4 FEC when the LDP **fast-reroute backup-sr-tunnel** option is enabled.

The **max-sr-frr-labels** parameter is used to limit the search for the TI-LFA backup next-hop:

- a. 0 — The IGP LFA SPF restricts the search to TI-LFA backup next-hop which does not require a repair tunnel, meaning that P node and Q node are the same and match a neighbor. This is also the case when both P and Q node match the advertising router for a prefix.
- b. 1 to 3 — The IGP LFA SPF will widen the search to include a repair tunnel to a P node which itself is connected to the Q nodes with a 0-to-2 hops for a total of maximum of three labels: one node SID to P node and two adjacency SIDs from P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed and meaning that up to three adjacency SIDs can separate the P and Q nodes.



- c. 2 (default) — Corresponds to a repair tunnel to a non-adjacent P which is adjacent to the Q node. If the P node is a neighbor of the computing node, then the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

When the **node-protect** command is enabled, the router will prefer a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.

The parameter **max-pq-nodes** in Remote LFA controls the maximum number of PQ nodes found in the LFA SPFs for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the **max-pq-cost** parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel.

The **no** form of this command disables the LFA computation by IGP SPF.

**Default** no loopfree-alternates

## exclude

**Syntax** **exclude**

**Context** config>router>isis>loopfree-alternates

**Description** This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF.



**Note:** Prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form of this command deletes the exclude prefix policy.

---

**Default** no exclude

## prefix-policy

**Syntax** **[no] prefix-policy** *prefix-policy* [*prefix-policy*]

**Context** config>router>isis>loopfree-alternates>exclude

**Description** This command specifies the name of the policy for the prefixes to exclude from the LFA SPF calculation in this ISIS instance.

The **no** form of this command deletes the exclude prefix policy.

**Default** no prefix-policy

**Parameters** **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. Up to five prefix policies can be specified. The specified name must have been already defined.

## remote-lfa

**Syntax** **remote-lfa** [**max-pq-cost** *value*]  
**no remote-lfa**

**Context** config>router>isis>loopfree-alternates

**Description** This command enables the use of the Remote LFA algorithm in the LFA SPF calculation for this ISIS instance.

The **no** form of this command disables the use of the Remote LFA algorithm in the LFA SPF calculation for this ISIS instance.

**Default** no remote-lfa

**Parameters** **max-pq-cost** *value* — Specifies the integer used to limit the search of candidate P and Q nodes in the remote LFA by setting the maximum IGP cost from the router performing the remote LFA calculation to the candidate P or Q node.

**Values** 0 to 4294967295

**Default** 4261412864

## node-protect

**Syntax** **node-protect** [**max-pq-nodes** *value*]  
**no node-protect**

**Context** config>router>isis>loopfree-alternates>remote-lfa

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.</p> <p>The <b>no</b> form of this command disables node-protect.</p>                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>     | no node-protect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><b>max-pq-nodes</b> <i>value</i> — Specifies the maximum number of PQ nodes found in the LFA SPF for which the node protection check is performed. The node-protect condition means the router must run the original Remote LFA algorithm plus one extra forward SPF on behalf of each PQ node found, potentially after applying the <b>max-pq-cost</b> parameter, to check if the path from the PQ node to the destination does not traverse the protected node. Setting this parameter to a lower value means the LFA SPFs will use less computation time and resources but may result in not finding a node-protect repair tunnel</p> <p><b>Values</b> 1 to 32</p> <p><b>Default</b> 16</p> |

## ti-lfa

|                    |                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>ti-lfa</b> [<b>max-sr-frr-labels</b> <i>value</i>]</p> <p><b>no ti-lfa</b></p>                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>isis>loopfree-alternates                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command enables the use of the Topology-Independent LFA algorithm in the LFA SPF calculation for this ISIS instance.</p> <p>The <b>no</b> form of this command disables the use of the Topology-Independent LFA algorithm in the LFA SPF calculation for this ISIS instance.</p>                   |
| <b>Default</b>     | no ti-lfa                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><b>max-sr-frr-labels</b> [<i>value</i>] — Specifies the maximum number of labels that the TI-LFA backup next-hop can use. The TI-LFA algorithm uses this value to limit the search for the Q node from the P node on the post-convergence path.</p> <p><b>Values</b> 0 to 3</p> <p><b>Default</b> 2</p> |

## node-protect

|                |                                               |
|----------------|-----------------------------------------------|
| <b>Syntax</b>  | <b>[no] node-protect</b>                      |
| <b>Context</b> | config>router>isis>loopfree-alternates>ti-lfa |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables node-protect in which the router prefers a node-protect over a link-protect repair tunnel for a given prefix if both are found in the Remote LFA or TI-LFA SPF computations. The SPF computations may only find a link-protect repair tunnel for prefixes owned by the protected node.</p> <p>The <b>no</b> form of this command disables node-protect.</p> |
| <b>Default</b>     | no node-protect                                                                                                                                                                                                                                                                                                                                                                     |

## lsp-lifetime

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp-lifetime</b> <i>seconds</i><br><b>no lsp-lifetime</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.</p> <p>Each LSP received is maintained in an LSP database until the <b>lsp-lifetime</b> expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.</p> <p>The LSP refresh timer is derived from this formula: <math>\text{lsp-lifetime}/2</math></p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | lsp-lifetime 1200                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>seconds</i> — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.<br><b>Values</b> 350 to 65535                                                                                                                                                                                                                                                                                                                                                                                                                  |

## lsp-refresh-interval

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp-refresh-interval</b> [ <i>seconds</i> ] [ <b>half-lifetime</b> {enable   disable}]<br><b>no lsp-refresh-interval</b>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the IS-IS LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for <b>lsp-lifetime</b> must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.</p> <p>The <b>no</b> form of this command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the no lsp-refresh-interval command will be rejected.</p> |

---

|                   |                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | lsp-refresh-interval 600 half-lifetime enable                                                                                                                                               |
| <b>Parameters</b> | <i>seconds</i> — Specifies the refresh interval.                                                                                                                                            |
| <b>Values</b>     | 150 to 65535                                                                                                                                                                                |
|                   | <b>half-lifetime</b> — Sets the refresh interval to always be half the <b>lsp-lifetime</b> value. When this parameter is set to <b>enable</b> , the configured refresh interval is ignored. |
| <b>Values</b>     | enable, disable                                                                                                                                                                             |

## multi-topology

|                    |                                                    |
|--------------------|----------------------------------------------------|
| <b>Syntax</b>      | [no] multi-topology                                |
| <b>Context</b>     | config>router>isis                                 |
| <b>Description</b> | This command enables IS-IS multi-topology support. |
| <b>Default</b>     | no multi-topology                                  |

## ipv4-multicast

|                    |                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] ipv4-multicast                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>is-is>multi-topology                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance.</p> <p>The <b>no</b> form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.</p> |
| <b>Default</b>     | no ipv4-multicast                                                                                                                                                                                                                 |

## ipv6-multicast

|                    |                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] ipv6-multicast                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>is-is>multi-topology                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables support for the IPv6 topology (MT4) within the associate IS-IS instance.</p> <p>The <b>no</b> form of this command disables support for the IPv6 topology (MT4) within the associated IS-IS instance.</p> |
| <b>Default</b>     | no ipv6-multicast                                                                                                                                                                                                                 |

---

## ipv6-unicast

|                    |                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ipv6-unicast</b>                                                                                                |
| <b>Context</b>     | config>router>isis>multi-topology                                                                                       |
| <b>Description</b> | <p>This command enables multi-topology TLVs.</p> <p>The <b>no</b> form of the command disables multi-topology TLVs.</p> |
| <b>Default</b>     | no ipv6-unicast                                                                                                         |

## multicast-import

|                    |                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] multicast-import [{both   ipv4   ipv6}]</b>                                                                                                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables the submission of routes into the multicast Route Table Manager (RTM) by IS-IS.</p> <p>The <b>no</b> form of the command disables the submission of routes into the multicast RTM.</p> |
| <b>Default</b>     | no multicast-import                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><b>both</b> — allows submission of both IPv4 and IPv6 routes</p> <p><b>ipv4</b> — allows submission of IPv4 routes only</p> <p><b>ipv6</b> — allows submission of IPv6 routes only</p>                      |

## overload

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>overload [timeout seconds] [max-metric]</b><br><b>no overload</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.</p> <p>During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic.</p> <p>If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.</p> |

The overload command is cleared from the configuration after a reboot if [overload-on-boot](#) is configured with or without a timeout value. To keep the IS-IS router in the overload state indefinitely after rebooting, configure overload-on-boot with no timeout value or configure the overload command with no overload-on-boot command.

The **overload** command can be useful in circumstances where the router is overloaded or used prior to executing a **shutdown** command to divert traffic around the router.

The **max-metric** parameter can be set to advertise transit links with the maximum metric of 0xffffffff (wide metrics) or 0x3f (regular metrics), instead of setting the overload bit when placing the router in overload.

The **no** form of the command causes the router to exit the overload state.

|                   |                                                                                         |
|-------------------|-----------------------------------------------------------------------------------------|
| <b>Default</b>    | no overload                                                                             |
| <b>Parameters</b> | <i>seconds</i> — The time, in seconds, that this router must operate in overload state. |
| <b>Default</b>    | infinity (overload state maintained indefinitely)                                       |
| <b>Values</b>     | 60 to 1800                                                                              |
| <b>max-metric</b> | — Set the maximum metric instead of overload.                                           |

## overload-export-external

|                    |                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] overload-export-external</b>                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                              |
| <b>Description</b> | This command enables external routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.<br><br>The <b>no</b> form of the command causes external routes to be withdrawn when the router is in overload. |
| <b>Default</b>     | no overload-export-external                                                                                                                                                                                                                                     |

## overload-export-interlevel

|                    |                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] overload-export-interlevel</b>                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command enables inter-level routes that are exported with an IS-IS export policy to continue to be advertised when the router is in overload.<br><br>The <b>no</b> form of the command causes inter-level routes to be withdrawn when the router is in overload. |

**Default** no overload-export-interlevel

## overload-on-boot

**Syntax** **overload-on-boot** [*timeout seconds*] [*max-metric*]  
**no overload-on-boot**

**Context** config>router>isis

**Description** When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

1. The timeout timer expires.
2. A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload : Manual on boot (Indefinitely in overload)
- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Use the show router isis status command to display the administrative and operational state as well as all timers.

**Default** no overload-on-boot

**Parameters** **timeout seconds** — Configure the timeout timer for overload-on-boot in seconds.

**Values** 60 to 1800

**max-metric** — Sets the maximum metric instead of overload.



## poi-tlv-enable

|                    |                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>poi-tlv-enable</b><br><b>no poi-tlv-enable</b>                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge.</p> <p>The <b>no</b> form of the command removes the POI functionality from the configuration.</p> |
| <b>Default</b>     | no poi-tlv-enable                                                                                                                                                                                                                                                                                                                                            |

## prefix-attributes-tlv

|                    |                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] prefix-attributes-tlv</b>                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables IS-IS Prefix Attributes TLV support to exchange extended IPv4 and IPv6 reachability information. Extended reachability information is required for traffic engineering features using path computation element (PCE) or optimal route reflection.</p> <p>The <b>no</b> form of the command removes the <b>prefix-attributes-tlv</b> configuration.</p> |
| <b>Default</b>     | no prefix-attributes-tlv                                                                                                                                                                                                                                                                                                                                                       |

## prefix-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefix-limit <i>limit</i> [log-only] [threshold <i>percent</i>] [overload-timeout {<i>seconds</i>   forever}]</b><br><b>no prefix-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the maximum number of prefixes that IS-IS can learn, and use to protect the system from a router that has accidentally advertised a large number of prefixes. If the number of prefixes reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, IS-IS will go into overload.</p> <p>The <b>overload-timeout</b> option controls the length of time that IS-IS is in the overload state when the <b>prefix-limit</b> is reached. The system automatically attempts to restart IS-IS at the end of this duration. If the <b>overload-timeout forever</b> option is used, IS-IS is not restarted automatically and stays in overload until the condition is manually cleared by the administrator. This is also the default behavior when the <b>overload-timeout</b> option is not configured.</p> |

The **no** form of the command removes the **prefix-limit**.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no prefix-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b> | <p><b>log-only</b> — Enables a warning message to be sent at the specified threshold percentage and also when the limit is exceeded. However, overload is not set when this parameter is configured.</p> <p><b>limit</b> — The number of prefixes that can be learned expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4294967296</p> <p><b>percent</b> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p><b>Values</b> 0 to 100</p> <p><b>seconds</b> — The time in minutes before IS-IS is restarted.</p> <p><b>Values</b> 1 to 1800</p> <p><b>forever</b> — Specifies that IS-IS should be restarted only after the execution of the <b>clear router isis overload prefix-limit</b> command.</p> |

## reference-bandwidth

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>reference-bandwidth</b> <i>bandwidth-in-kbps</i></p> <p><b>reference-bandwidth</b> [<b>tbps</b> <i>Tera-bps</i>] [<b>gbps</b> <i>Giga-bps</i>] [<b>mbps</b> <i>Mega-bps</i>] [<b>kbps</b> <i>Kilo-bps</i>]</p> <p><b>no reference-bandwidth</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures the reference bandwidth that provides the basis of bandwidth relative costing.</p> <p>In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula:</p> $\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$ <p>If the reference bandwidth is configured as 10 Gb (10,000,000,000), a 100 Mb/s interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed. (See wide-metrics-only in the <b>config&gt;router&gt;isis</b> context.)</p> <p>If the reference bandwidth is not configured, then all interfaces have a default metric of 10.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | no reference-bandwidth                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

---

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <p><i>bandwidth-in-kbps</i> — The reference bandwidth in kilobits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4000000000</p> <p><i>Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 4</p> <p><i>Giga-bps</i> — The reference bandwidth in gigabits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 999</p> <p><i>Mega-bps</i> — The reference bandwidth in megabits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 999</p> <p><i>Kilo-bps</i> — reference bandwidth in kilobits per second expressed as a decimal integer.</p> <p><b>Values</b> 1 to 999</p> |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## rib-priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>rib-priority high</b> {<i>prefix-list-name</i>   <b>tag</b> <i>tag-value</i>}</p> <p><b>no rib-priority</b></p>                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.</p> <p>The no rib-priority form of command disables RIB prioritization.</p>                                                       |
| <b>Default</b>     | no rib-priority high                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><i>prefix-list-name</i> — specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.</p> <p><i>tag tag-value</i> — specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.</p> <p><b>Values</b> 1 to 4294967295</p> |

## router-id

|                |                                                                     |
|----------------|---------------------------------------------------------------------|
| <b>Syntax</b>  | <p><b>router-id</b> <i>router-id</i></p> <p><b>no router-id</b></p> |
| <b>Context</b> | config>router>isis                                                  |

---

|                    |                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures the router ID.<br><br>The <b>no</b> form of the command deletes the router ID. |
| <b>Parameters</b>  | <i>router-id</i> — the IP address of the router                                                        |

## segment-routing

|                    |                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>segment-routing</b><br><b>no segment-routing</b>                                                   |
| <b>Context</b>     | config>router>isis                                                                                    |
| <b>Description</b> | This command enables the context to configure segment routing parameters within a given IGP instance. |

Segment routing adds to IS-IS and OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface/next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as Segment ID (SID).

When segment routing is used together with MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels.

Segment routing using MPLS labels can be used in both shortest path routing applications and in traffic engineering applications. This feature implements the shortest path forwarding application.

After segment routing is successfully enabled in the IS-IS or OSPF instance, the router will perform the following operations:

1. Advertise the Segment Routing Capability Sub-TLV to routers in all areas/levels of this IGP instance. However, only neighbors with which it established an adjacency will interpret the SID/label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.
2. Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. Then the segment routing module programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Assign and advertise automatically an adjacency SID label for each formed adjacency over a network IP interface in the new Adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, in effect with a swap to an implicit null label operation, for each advertised adjacency SID.
4. Resolve received prefixes and if a prefix SID sub-TLV exists, the Segment Routing module programs the ILM with a swap operation and also an LTN with a push operation both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in a given IGP instance, the main SPF and LFA SPF are computed normally and the primary next-hop and LFA backup next-hop for a received prefix are added to RTM without the label information advertised in the prefix SID sub-TLV.

## family

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>family {ipv4   ipv6}</b><br><b>no family</b>                                                                                          |
| <b>Context</b>     | config>router>isis>segment-routing>adjacency-set                                                                                         |
| <b>Description</b> | This command specifies the address family of an adjacency set in IS-IS.<br><br>The <b>no</b> form of the command reverts to the default. |
| <b>Default</b>     | family ipv4                                                                                                                              |
| <b>Parameters</b>  | <b>ipv4</b> — Specifies a family of IPv4.<br><b>ipv6</b> — Specifies a family of IPv6.                                                   |

## adj-sid-hold

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |               |          |                |    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------|----------------|----|
| <b>Syntax</b>      | <b>adj-sid-hold <i>seconds</i></b><br><b>no adj-sid-hold</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |               |          |                |    |
| <b>Context</b>     | config>router>isis>segment-routing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |          |                |    |
| <b>Description</b> | This command configures a timer to hold the ILM or LTM of an adjacency SID following a failure of the adjacency.<br><br>When an adjacency to a neighbor fails, IGP will withdraw the advertisement of the link TLV information as well as its adjacency SID sub-TLV. However, the LTN or ILM record of the adjacency SID must be kept in data path to maintain forwarding using the LFA or remote LFA backup for a period of time sufficient to allow the ingress LER and other routers which use this adjacency SID to activate a new path after IGP converges.<br><br>If the adjacency is restored before the timer expires, the timer is aborted as soon as the new ILM or LTN records are updated with the new primary and backup NHLFE information.<br><br>The <b>no</b> form of the command removes adjacency SID hold time. |               |          |                |    |
| <b>Default</b>     | adj-sid-hold 15                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |               |          |                |    |
| <b>Parameters</b>  | <b><i>seconds</i></b> — the adjacency SID hold time, in seconds<br><br><table> <tr> <td><b>Values</b></td><td>1 to 300</td></tr> <tr> <td><b>Default</b></td><td>15</td></tr> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>Values</b> | 1 to 300 | <b>Default</b> | 15 |
| <b>Values</b>      | 1 to 300                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |               |          |                |    |
| <b>Default</b>     | 15                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |          |                |    |

---

## export-tunnel-table

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-tunnel-table ldp</b><br><b>no export-tunnel-table</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>isis>segment-routing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command enables exporting to an IGP instance the LDP tunnels for the purpose of stitching a SR tunnel to a LDP FEC for the same destination IPv4 /32 prefix.</p> <p>In the SR-to-LDP data path direction, the SR mapping server provides a global policy for the prefixes corresponding to the LDP FECs the SR stitches to.</p> <p>When this command is enabled in the segment-routing context of an IGP instance, IGP listens to LDP tunnel entries in the TTM. Whenever a LDP tunnel destination matches a prefix for which IGP received a prefix-SID sub-TLV from a mapping server, it instructs the SR module to program the SR ILM and to stitch it to the LDP tunnel endpoint. The LDP FEC can be resolved via a static route, a IS-IS instance, or an OSPF instance.</p> <p>When an SR tunnel is stitched to a LDP FEC, packets forwarded will benefit from the protection of the LFA backup next-hop of the LDP FEC.</p> <p>When resolving a node SID, IGP will prefer resolution of prefix SID received in a IP Reach TLV over a prefix SID received via the mapping server. In other words the swapping of the SR ILM to a SR NHLFE is preferred over stitching it to a LDP tunnel endpoint.</p> <p>It is recommended to enable the bfd-enable option on the interfaces in both LDP and IGP instance contexts to speed up the failure detection and the activation of the LFA/remote-LFA backup next-hop in either direction of the stitching.</p> <p>This feature is limited to IPv4 /32 prefixes in both LDP and SR.</p> <p>The <b>no</b> form of the command disables the exporting of LDP tunnels to the IGP instance.</p> |
| <b>Default</b>     | no export-tunnel-table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <b>ldp</b> — exports LDP tunnels from the tunnel table into an IGP instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## mapping-server

|                    |                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mapping-server</b>                                                                               |
| <b>Context</b>     | config>router>isis>segment-routing                                                                       |
| <b>Description</b> | This command configures the context for the Segment Routing mapping server feature in an IS-IS instance. |

The mapping server feature allows the configuration and advertisement via IS-IS of the node SID index for IS-IS prefixes of routers which are in the LDP domain. This is performed in the router acting as a mapping server and using a prefix-SID sub-TLV within the SID/Label binding TLV in IS-IS.

The **no** form of the command deletes all node SID entries in the IS-IS instance.

## sid-map

**Syntax** **sid-map node-sid** {**index** *value* [**range** *value*]} **prefix** {{*ip-address/mask*} | {*ip-address*} {*netmask*}} [**set-flags** {**s**}] [**level** {**1** | **2** | **1/2**}] [**clear-n-flag**]  
**no sid-map node-sid index** *value*

**Context** config>router>isis>segm-rtn>mapping-server

**Description** This command configures the Segment Routing mapping server database in IS-IS.

The user can enter the node SID index for one or a range of prefixes by specifying the first index value and optionally a range value. The default value for the range option is 1. Only the first prefix in a consecutive range of prefixes must be entered. The user can enter the first prefix with a mask lower than 32 and the SID/Label Binding TLV is advertised but the routers will not resolve these prefix SIDs and will instead originate a trap.

The user can indicate to the IS-IS routers in the rest of the network that the flooding scope of the SID/Label binding TLV is the entire domain by setting the S-flag. In that case, a router receiving the TLV advertisement should leak it between ISIS levels. If leaked from level 2 to level 1, the D-flag must be set and once set the TLV cannot be leaked back into level 2. Otherwise, the S-flag is clear by default and the TLV must not be leaked by routers receiving the mapping server advertisement.

Note that the SR OS does not leak this TLV between IS-IS instances and does not support the multi-topology SID/Label Binding TLV format.

In addition, the user can specify the mapping server own flooding scope for the generated SID/Label binding TLV using the level option. This option allows further narrowing of the flooding scope configured under the router IS-IS level-capability for a one or more SID/Label binding TLVs if required. The default flooding scope of the mapping server is L1/L2 which can be narrowed by what is configured under the router IS-IS level-capability.

The A-flag and M-flag are not supported by the mapping server feature. The mapping client ignores them.

Each time a prefix or a range of prefixes is configured in the SR mapping database in any routing instance, the router issues for this prefix, or range of prefixes, a prefix-SID sub-TLV within a ISIS SID/label binding TLV in that instance. The flooding scope of the TLV from the mapping server is determined as explained above. No further check of the reachability of that prefix in the mapping server route table is performed and no check if the SID index is duplicate with some existing prefix in the local IGP instance database or if the SID index is out of range with the local SRGB.

The **no** form of the `sid-map` command deletes the range of node SIDs beginning with the specified index value.

|                   |                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>index</i> — Specifies the node SID index for the IS-IS prefix which will be advertised in a SID/Label Binding TLV   |
|                   | <b>Values</b> 0 to 4294967295                                                                                          |
|                   | <b>Default</b> none                                                                                                    |
|                   | <i>value</i> — Specifies the node SID range for the IS-IS prefix which will be advertised in a SID/Label Binding TLV   |
|                   | <b>Values</b> 0 to 65535                                                                                               |
|                   | <b>Default</b> none                                                                                                    |
|                   | <i>ip-address/mask</i> — Specifies the IP address and mask                                                             |
|                   | <b>Values</b> <i>ip-address</i> : <b>a.b.c.d.</b> (host bits must be 0)<br><i>mask</i> : <b>0 to 32</b>                |
|                   | <i>ip-address netmask</i> — Specifies the IP address netmask                                                           |
|                   | <b>Values</b> <b>a.b.c.d.</b> (network bits all 1 and host bits all 0)                                                 |
|                   | <b>set-flags</b> — Specifies the flooding scope of the SID/Label binding TLV                                           |
|                   | <b>Default</b> <b>S-flag clear</b><br>The TLV is not leaked by routers receiving the mapping server advertisement      |
|                   | <b>level {1   2   1/2}</b> — Configures the mapping server own flooding scope for the generated SID/Label binding TLV. |
|                   | <b>Default</b> 1/2                                                                                                     |
|                   | <b>clear-n-flag</b> — Specifies whether the node-sid flag (N-flag) should be cleared in a SID Label Binding TLV.       |

## prefix-sid-range

|                    |                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefix-sid-range {global   start-label label-value max-index index-value}</b><br><b>no prefix-sid-range</b> |
| <b>Context</b>     | config>router>isis>segment-routing                                                                             |
| <b>Description</b> | This command configures the prefix SID index range and offset label value for a given IGP instance.            |



The key parameter is the configuration of the prefix SID index range and the offset label value which this IGP instance will use. Since each prefix SID represents a network global IP address, the SID index for a prefix must be network-wide unique. Thus, all routers in the network are expected to configure and advertise the same prefix SID index range for a given IGP instance. However, the label value used by each router to represent this prefix; that is, the label programmed in the ILM can be local to that router by the use of an offset label, referred to as a start label:

$$\text{Local Label (Prefix SID)} = \text{start-label} + \{\text{SID index}\}$$

The label operation in the network becomes thus very similar to LDP when operating in the independent label distribution mode (RFC 5036, *LDP Specification*) with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the **global** mode of operation, the user configures the global value and this IGP instance will assume the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. Once one IGP instance selected the global option for the prefix SID range, all IGP instances on the system will be restricted to do the same. The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. Once the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change will be failed if an already allocated SID index/label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user thus configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. All resulting net label values ( $\text{start-label} + \text{index}$ ) must be within the SRGB or the configuration will be failed. Furthermore, the code checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce that these ranges do not overlap. The user must shutdown the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. In addition, any range change will be failed if an already allocated SID index/label goes out of range. The user can however change the SRGB on the fly as long as it does not reduce the current per IGP instance SID index/label range defined with the **prefix-sid-range**. Otherwise, the user must shutdown the segment routing context of the IGP instance and delete and re-configure the **prefix-sid-range** command.

|                   |                                                                                                                                                                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no prefix-sid-range                                                                                                                                                                                                                                                                |
| <b>Parameters</b> | <p><i>label-value</i> — Specifies the label offset for the SR label range of this IGP instance.</p> <p><b>Values</b> 0 to 524287</p> <p><i>index-value</i> — Specifies the maximum value of the prefix SID index range for this IGP instance.</p> <p><b>Values</b> 1 to 524287</p> |

---

tunnel-mtu

|                    |                                                                             |
|--------------------|-----------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-mtu</b> <i>bytes</i><br><b>no tunnel-mtu</b>                      |
| <b>Context</b>     | config>router>isis>segment-routing                                          |
| <b>Description</b> | This command configures the MTU of all SR tunnels within each IGP instance. |

The MTU of a SR tunnel populated into TTM is determined like in the case of an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote LFA can add, at most, one more label to the tunnel for a total of two labels. There is no default value for this command. If the user does not configure an SR tunnel MTU, the MTU is determined by IGP as explained below.

The MTU of the SR tunnel in bytes is then determined as follows:

$$SR\_Tunnel\_MTU = MIN \{ Cfg\_SR\_MTU, IGP\_Tunnel\_MTU - (1 + frr-overhead) * 4 \}$$

Where:

*Cfg\_SR\_MTU* is the MTU configured by the user for all SR tunnels within a given IGP instance using the above CLI. If no value was configured by the user, the SR tunnel MTU will be determined by the IGP interface calculation explained next.

*IGP\_Tunnel\_MTU* is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.

*frr-overhead* is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGMP instance. Otherwise, it is set to 0.

The SR tunnel MTU is dynamically updated anytime any of the above parameters used in its calculation changes. This includes when the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.

|                   |                                                                                    |
|-------------------|------------------------------------------------------------------------------------|
| <b>Default</b>    | no tunnel-mtu                                                                      |
| <b>Parameters</b> | <i>bytes</i> — Specifies the size of the Maximum Transmission Unit (MTU) in bytes. |
| <b>Values</b>     | 512 to 9198                                                                        |

## tunnel-table-pref

|                |                                                                           |
|----------------|---------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>tunnel-table-pref</b> <i>preference</i><br><b>no tunnel-table-pref</b> |
| <b>Context</b> | config>router>isis>segment-routing                                        |

**Description** This command configures the TTM preference of SR tunnels created by the IGP instance. This is used in the case of BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the new tunnel binding commands are configured to the **any** value which parses the TTM for tunnels in the protocol preference order. The user can choose to either go with the global TTM preference or list explicitly the tunnel types they want to use. When they list the tunnel types explicitly, the TTM preference will still be used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds to TTM a SR tunnel entry for each resolved remote node SID prefix and programs the data path with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs.

The default preference for SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference of the various tunnel types. This includes the preference of SR tunnels based on shortest path (referred to as **SR-ISIS** and **SR-OSPF**).

The global default TTM preference for the tunnel types is as follows:

- ROUTE\_PREF\_RSVP 7
- ROUTE\_PREF\_SR\_TE 8
- ROUTE\_PREF\_LDP 9
- ROUTE\_PREF\_OSPF\_TTM 10
- ROUTE\_PREF\_ISIS\_TTM 11
- ROUTE\_PREF\_BGP\_TTM 12
- ROUTE\_PREF\_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of a SR tunnel in this case will be based on lowest IGP instance-id.

|                   |                                                                                                                 |
|-------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | tunnel-table-pref 11                                                                                            |
| <b>Parameters</b> | <i>preference</i> — Specifies the integer value to represent the preference of IS-IS or OSPF SR tunnels in TTM. |
| <b>Values</b>     | 1 to 255                                                                                                        |
| <b>Default</b>    | 11                                                                                                              |

standard-multi-instance

|                |                              |
|----------------|------------------------------|
| <b>Syntax</b>  | [no] standard-multi-instance |
| <b>Context</b> | config>router>isis           |

**Description** This command enables IS-IS multi-instance (MI) as described in *draft-ginsberg-isis-mi-bis-01*. Multiple instances allow instance-specific adjacencies to be formed that support multiple network topologies on the same physical interfaces. Each instance has an LSDB, and each PDU contains a TLV identifying the instance and the topology to which the PDU belongs. A single topology is supported in each instance, so the instance-specific topology identifier (ITID) is set to 0 and cannot be changed.

The **standard-multi-instance** (based on *draft-ginsberg-isis-mi-bis-01*) and **iid-tlv-enable** (based on *draft-ietf-isis-mi-02*) commands cannot be configured in the same instance, because the MAC addresses and PDUs from the two standards are incompatible.

The **no** form of the command removes the **standard-multi-instance** configuration.

**Default** no standard-multi-instance

## strict-adjacency-check

**Syntax** [no] **strict-adjacency-check**

**Context** config>router>isis

**Description** This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well.

When disabled (**no strict-adjacency-check**), both routers only need to have one common address family to establish the adjacency.

**Default** no strict-adjacency-check

## summary-address

**Syntax** **summary-address** {*ip-prefix/mask* | *ip-prefix netmask*} [*level*] [**tag** *tag*]  
**no summary-address** {*ip-prefix/ip-prefix-length* | *ip-prefix netmask*}

**Context** config>router>isis

**Description** This command creates summary-addresses.

**Parameters** *ip-prefix/mask* — Specifies information for the specified IP prefix and mask length.

**Values** ipv4-prefix:  
• a.b.c.d (host bits must be 0)  
ipv4-prefix-length: [0 to 32]

ipv6-prefix:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ipv6-prefix-length: [0 to 128]

*netmask* — The subnet mask in dotted decimal notation.

**Values** 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

*level* — Specifies IS-IS level area attributes. If no level parameter is specified, the default is level-1/2.

**Values** level-1, level-2, level-1/2

*tag* — Assigns a route tag to the summary address.

**Values** 1 to 4294967295

suppress-attached-bit

|                    |                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>suppress-attached-bit</b><br><b>no suppress-attached-bit</b>                                                                                                            |
| <b>Context</b>     | config>router>isis                                                                                                                                                         |
| <b>Description</b> | This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it. |
| <b>Default</b>     | no suppress-attached-bit                                                                                                                                                   |

system-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>system-id</b> <i>isis-system-id</i><br><b>no system-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command configures the IS-IS system ID. The system ID has a fixed length of 6 octets; it is determined using the following preference:<br><br><ol style="list-style-type: none"><li>1. <b>configure&gt;router&gt;isis&gt;system-id</b></li><li>2. <b>configure&gt;router&gt;isis&gt;router-id</b></li><li>3. <b>configure&gt;router&gt;router-id</b></li><li>4. <b>configure&gt;router&gt;interface&gt;system&gt;address</b></li><li>5. The default system ID 2550.0000.0000, based on the default router ID 255.0.0.0</li></ol> |

The system ID is integral to IS-IS; therefore, for the **system-id** command to take effect, the IS-IS instance must be **shutdown** and then **no shutdown**. This will ensure that the configured and operational system ID are always the same.

The **no** form of the command removes the system ID from the configuration. The router ID is used when no system ID is specified.

|                   |                                                                                      |
|-------------------|--------------------------------------------------------------------------------------|
| <b>Default</b>    | no system-id                                                                         |
| <b>Parameters</b> | <i>isis-system-id</i> — Specifies 12 hexadecimal characters in dotted-quad notation. |
| <b>Values</b>     | aaaa.bbbb.cccc, where aaaa, bbbb, and cccc are hexadecimal numbers                   |

## timers

|                    |                                                 |
|--------------------|-------------------------------------------------|
| <b>Syntax</b>      | [no] timers                                     |
| <b>Context</b>     | config>router>isis                              |
| <b>Description</b> | This command configures the IS-IS timer values. |

## lsp-wait

|                    |                                                                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp-wait</b> <i>lsp-wait</i> [ <b>lsp-initial-wait</b> <i>initial-wait</i> ] [ <b>lsp-second-wait</b> <i>second-wait</i> ]                                                                                                                                                                              |
| <b>Context</b>     | config>router>isis>timers                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command is used to customize LSP generation throttling. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second <b>lsp-wait</b> timer until a maximum value is reached. |



**Note:** The timer granularity is 10 ms if the value is less than 500 ms, and 100 ms if the value is greater than or equal to 500 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

|                   |                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | lsp-initial-wait 10 lsp-second-wait 1000                                                                                            |
| <b>Parameters</b> | <i>lsp-max-wait</i> — Specifies the maximum interval in milliseconds between two consecutive occurrences of an LSP being generated. |
| <b>Values</b>     | 10 to 120000                                                                                                                        |
| <b>Default</b>    | 5000                                                                                                                                |

*initial-wait* — Specifies the initial LSP generation delay in milliseconds. Values less than 100 ms are internally rounded down to 0, so that there is no added initial LSP generation delay.

**Values** 10 to 100000

**Default** 10

*second-wait* — Specifies the hold time in milliseconds between the first and second LSP generation.

**Values** 10 to 100000

**Default** 1000

## spf-wait

**Syntax** [no] **spf-wait** *spf-wait* [**spf-initial-wait** *initial-wait*] [**spf-second-wait** *second-wait*]

**Context** config>router>isis>timers

**Description** This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command.

Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.



**Note:** The timer granularity is 100 ms. Timer values are rounded down to the nearest granularity, for example a configured value of 550 ms is internally rounded down to 500 ms.

**Default** spf-initial-wait 1000 spf-second-wait 1000

**Parameters** *spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

**Values** 10 to 120000

**Default** 10000

*initial-wait* — Specifies the initial SPF calculation delay in milliseconds after a topology change.

**Values** 10 to 100000

**Default** 1000

---

*second-wait* — Specifies the hold time in milliseconds between the first and second SPF calculation.

**Values** 10 to 100000

**Default** 1000

## traffic-engineering

**Syntax** [no] traffic-engineering

**Context** config>router>isis

**Description** This command configures traffic-engineering and determines if IGP shortcuts are required by BGP.

**Default** no traffic-engineering

## unicast-import-disable

**Syntax** [no] unicast-import-disable [{both | ipv4 | ipv6}]

**Context** config>router>isis

**Description** This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM; as such, an import policy must be explicitly configured.

**Default** no unicast-import-disable both

**Parameters** **both** — allows importation of both IPv4 and IPv6 routes

**ipv4** — allows importation of IPv4 routes only

**ipv6** — allows importation of IPv6 routes only



## 4.8 Show, Clear, and Debug Command Reference

### 4.8.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

#### 4.8.1.1 Show Commands

```
show
 — router
 — isis [isis-instance]
 — adjacency [ip-address | ip-int-name | nbr-system-id] [detail] [mt mt-id-number]
 — capabilities [system-id | lsp-id] [level level]
 — database [system-id | lsp-id] [detail] [level level]
 — hostname
 — interface [ip-int-name | ip-address] [detail]
 — interface [detail] mt mt-id-number
 — lfa-coverage
 — link-group-member-status name [level level]
 — link-group-status name [level level]
 — mapping-server [prefix ip-address[/mask]] [index index] [level level] [flags {s}]
 — prefix-sids [ipv4-unicast|ipv6-unicast|ipv4-multicast | ipv6-multicast | mt mt-id-number] [ip-prefix[/prefix-length]] [sid sid] [adv-router system-id | hostname] [srms | no-srms]
 — routes [ipv4-unicast | ipv6-unicast | ipv4-multicast | ipv6-multicast | mt mt-id-number] [ip-prefix[/prefix-length]] [alternative] [exclude-shortcut] [detail]
 — spf-log [detail]
 — sr-adj-sets [adjacency-set-id] [detail]
 — statistics
 — status
 — summary-address [ip-address [/prefix-length]]
 — topology [ipv4-unicast | ipv6-unicast | ipv4-multicast | ipv6-multicast | mt mt-id-number] [lfa] [detail]
```

#### 4.8.1.2 Clear Commands

```
clear
 — router
```

- **isis** [*isis-instance*]
  - **adjacency** [*system-id*]
  - **database** [*system-id*]
  - **export**
  - **overload** {*rtm* | *fib* | *prefix-limit*}
  - **spf-log**
  - **statistics**

### 4.8.1.3 Debug Commands

- debug
  - router
    - **isis** [*isis-instance*]
      - [no] **adjacency** [*ip-int-name* | *ip-address* | *nbr-system-id*]
      - [no] **cspf**
      - [no] **graceful-restart**
      - **interface** [*ip-int-name* | *ip-address*]
      - no **interface**
      - **leak** [*ip-address*]
      - no **leak**
      - [no] **lsdb** [*level-number*] [*system-id* | *lsp-id*]
      - [no] **misc**
      - **packet** [*packet-type*] [*ip-int-name* | *ip-address*] [**detail**]
      - **rtm** [*ip-address*]
      - no **rtm**
      - [no] **spf** [*level-number*] [*system-id*]
      - **summary** [*ip-address*]
      - no **summary**
      - **tunnel-endpoint** [**tunnel-spf**] [**tunnel-leak** *ip-address*]
      - no **tunnel-endpoint**

### 4.8.1.4 Tools Commands

- tools
  - dump
    - router
      - **isis**
        - **sr-adjacencies** [*ip-int-name* | *ip-address*] [**detail**]
        - **sr-database** [*sid* *sid*] [**prefix** *ip-address*] [**nh-type** {*sr* | *bgp* | *rsvp*}]  
[*ipv4-unicast* | *ipv6-unicast*] [**detail**]

## 4.8.2 Command Descriptions

### 4.8.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

#### isis

|                    |                                                                         |
|--------------------|-------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>isis</b> [ <i>isis-instance</i> ]                                    |
| <b>Context</b>     | show>router                                                             |
| <b>Description</b> | This command displays information for a specified IS-IS instance.       |
| <b>Parameters</b>  | <i>isis-instance</i> — Specifies the instance ID for an IS-IS instance. |
| <b>Values</b>      | 1 to 31                                                                 |
| <b>Default</b>     | 0                                                                       |

#### adjacency

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adjacency</b> [ <i>ip-address</i>   <i>ip-int-name</i>   <i>nbr-system-id</i> ] [ <b>detail</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command displays information regarding IS-IS neighbors. When no <i>ip-address</i> , <i>ip-int-name</i> , or <i>nbr-system-id</i> is specified, then all adjacencies are displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><i>ip-address</i> — When specified, only adjacencies with that interface is displayed.</p> <p><b>Values</b></p> <p>ipv4-address:</p> <ul style="list-style-type: none"> <li>• a.b.c.d (host bits must be 0)</li> </ul> <p>ipv6-address:</p> <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> <li>• d: [0 to 255]D</li> </ul> <p><i>ip-int-name</i> — When specified, only adjacencies with that interface is displayed.</p> <p><i>nbr-system-id</i> — When specified, only the adjacency with that ID is displayed.</p> <p><b>detail</b> — All output is displayed in the detailed format.</p> |

*mt-id-number* — Specifies the multi topology ID number for this route.

### Output Standard and Detailed IS-IS Adjacency Output

The following table describes the standard and detailed command output fields for an IS-IS adjacency.

**Table 37 Standard and Detailed Adjacency Output Fields**

| Label               | Description                                                                                                                                                                                                                                                 |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface           | Interface name associated with the neighbor.                                                                                                                                                                                                                |
| System-id           | Neighbor's system ID.                                                                                                                                                                                                                                       |
| Level               | 1-L1 only, 2-L2 only, 3-L1 and L2.                                                                                                                                                                                                                          |
| State               | Up, down, new, one-way, initializing, or rejected.                                                                                                                                                                                                          |
| Hold                | Hold time remaining for the adjacency.                                                                                                                                                                                                                      |
| SNPA                | Subnetwork point of attachment, MAC address of the next hop.                                                                                                                                                                                                |
| Circuit type        | Level on the interface L1, L2, or both.                                                                                                                                                                                                                     |
| Expires In          | Number of seconds until adjacency expires.                                                                                                                                                                                                                  |
| Priority            | Priority to become designated router.                                                                                                                                                                                                                       |
| Up/down transitions | Number of times neighbor state has changed.                                                                                                                                                                                                                 |
| Event               | Event causing last transition.                                                                                                                                                                                                                              |
| Last transition     | Time since last transition change.                                                                                                                                                                                                                          |
| Speaks              | Supported protocols (only IP).                                                                                                                                                                                                                              |
| IP address          | IP address of neighbor.                                                                                                                                                                                                                                     |
| MT enab             | Yes — The neighbor is advertising at least 1 non MTID#0.                                                                                                                                                                                                    |
| Topology            | Derived from the MT TLV in the IIH <ul style="list-style-type: none"> <li>• MT#0, MT#2 =&gt; "Topology : Unicast, IPv6-Unicast"</li> <li>• Native IPv4 or native IPv6 =&gt; "Topology : Unicast"</li> </ul> Not supported MTIDs => Topology line suppressed |

### Sample Output

```
*A:Dut-C# show router isis adjacency
```

```
=====
Rtr Base ISIS Instance 0 Adjacency
=====
System ID Usage State Hold Interface MT-ID
```

```

Dut-B L1L2 Up 23 to_Dut-B 0
Dut-D L1L2 Up 23 to_Dut-D1 0

Adjacencies : 2
=====
*A:Dut-C# show router isis adjacency Dut-D detail
=====
Rtr Base ISIS Instance 0 Adjacency (detail)
=====
SystemID : Dut-D SNPA : 00:00:00:00:00:04
Interface : to_Dut-D1 Up Time : 0d 00:05:23
State : Up Priority : 0
Nbr Sys Typ : L1L2 L. Circ Typ : L1L2
Hold Time : 19 Max Hold : 27
Adj Level : L1L2 MT Enabled : No
Topology : Unicast

IPv6 Neighbor : fe80::200:ff:fe00:4
IPv4 Neighbor : 1.3.4.4
IPv4 Adj SID : Label 262139
Restart Support : Disabled
Restart Status : Not currently being helped
Restart Suppressed : Disabled
Number of Restarts: 0
Last Restart at : Never

=====
*A:Dut-C#

*A:ALA-A# show router isis adjacency 180.0.7.12
=====
Rtr Base ISIS Instance 0 Adjacency
=====
System ID Usage State Hold Interface

asbr_east L2 Up 25 if2/5

Adjacencies : 1
=====
*A:ALA-A#

*A:ALA-A# show router isis adjacency if2/5
=====
Rtr Base ISIS Instance 0 Adjacency
=====
System ID Usage State Hold Interface

asbr_east L2 Up 20 if2/5

Adjacencies : 1
=====
*A:ALA-A#

*A:Dut-A# show router isis adjacency detail

```

```

=====
Rtr Base ISIS Instance 0 Adjacency (detail)
=====
SystemID : Dut-B SNPA : 20:81:01:01:00:01
Interface : ip-3FFE::A0A:101 Up Time : 0d 00:56:10
State : Up Priority : 64
Nbr Sys Typ : L1 L. Circ Typ : L1
Hold Time : 2 Max Hold : 2
Adj Level : L1 MT Enabled : Yes
Topology : Unicast, IPv6-Unicast

IPv6 Neighbor : FE80::2281:1FF:FE01:1
IPv4 Neighbor : 10.10.1.2
Restart Support : Disabled
Restart Status : Not currently being helped
Restart Supressed : Disabled
Number of Restarts: 0
Last Restart at : Never

SystemID : Dut-B SNPA : 20:81:01:01:00:01
Interface : ip-3FFE::A0A:101 Up Time : 0d 00:56:10
State : Up Priority : 64
Nbr Sys Typ : L2 L. Circ Typ : L2
Hold Time : 2 Max Hold : 2
Adj Level : L2 MT Enabled : Yes
Topology : Unicast, IPv6-Unicast

IPv6 Neighbor : FE80::2281:1FF:FE01:1
IPv4 Neighbor : 10.10.1.2
Restart Support : Disabled
Restart Status : Not currently being helped
Restart Supressed : Disabled
Number of Restarts: 0
Last Restart at : Never

SystemID : Dut-F SNPA : 00:00:00:00:00:00
Interface : ies-1-3FFE::A0A:1501 Up Time : 0d 01:18:34
State : Up Priority : 0
Nbr Sys Typ : L1L2 L. Circ Typ : L1L2
Hold Time : 5 Max Hold : 6
Adj Level : L1L2 MT Enabled : Yes
Topology : Unicast, IPv6-Unicast

IPv6 Neighbor : FE80::2285:FFFF:FE00:0
IPv4 Neighbor : 10.10.21.6
Restart Support : Disabled
Restart Status : Not currently being helped
Restart Supressed : Disabled
Number of Restarts: 0
Last Restart at : Never
=====
*A:Dut-A#

```

## capabilities

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>capabilities</b> [ <i>system-id</i>   <i>lsp-id</i> ] [ <i>level level</i> ] |
| <b>Context</b>     | show>router>isis                                                                |
| <b>Description</b> | This command displays the IS-IS capability information.                         |
| <b>Output</b>      |                                                                                 |

### Sample Output

```
*A:Dut-C# show router isis capabilities
=====
Rtr Base ISIS Instance 0 Capabilities
=====
Displaying Level 1 capabilities

LSP ID : Dut-A.00-00
 Router Cap : 10.20.1.1, D:0, S:0
 TE Node Cap : B E M P
 SR Cap: IPv4 MPLS-IPv6
 SRGB Base:20000, Range:20000
 SR Alg: metric based SPF
LSP ID : Dut-A.02-00
LSP ID : Dut-A.03-00
LSP ID : Dut-B.00-00
 Router Cap : 10.20.1.2, D:0, S:0
 TE Node Cap : B E M P
 SR Cap: IPv4 MPLS-IPv6
 SRGB Base:20000, Range:20000
 SR Alg: metric based SPF
LSP ID : Dut-C.00-00
 Router Cap : 10.20.1.3, D:0, S:0
 TE Node Cap : B E M P
 SR Cap: IPv4 MPLS-IPv6
 SRGB Base:20000, Range:20000
 SR Alg: metric based SPF
LSP ID : Dut-C.02-00
LSP ID : Dut-D.00-00
 Router Cap : 10.20.1.4, D:0, S:0
 TE Node Cap : B E M P
 SR Cap: IPv4 MPLS-IPv6
 SRGB Base:20000, Range:20000
 SR Alg: metric based SPF
LSP ID : Dut-D.01-00
LSP ID : Dut-E.00-00
 Router Cap : 10.20.1.5, D:0, S:0
 TE Node Cap : B E M P
 SR Cap: IPv4 MPLS-IPv6
 SRGB Base:20000, Range:20000
 SR Alg: metric based SPF
LSP ID : Dut-E.01-00
LSP ID : Dut-E.02-00
LSP ID : Dut-F.00-00
 Router Cap : 10.20.1.6, D:0, S:0
 TE Node Cap : B E M P
```

```
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-F.01-00
LSP ID : Dut-F.03-00
Level (1) Capability Count : 14
Displaying Level 2 capabilities
```

```

LSP ID : Dut-A.00-00
Router Cap : 10.20.1.1, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-A.02-00
LSP ID : Dut-A.03-00
```

```
LSP ID : Dut-B.00-00
Router Cap : 10.20.1.2, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
```

```
LSP ID : Dut-C.00-00
Router Cap : 10.20.1.3, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
```

```
LSP ID : Dut-C.02-00
LSP ID : Dut-D.00-00
Router Cap : 10.20.1.4, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
```

```
LSP ID : Dut-D.01-00
LSP ID : Dut-E.00-00
Router Cap : 10.20.1.5, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
```

```
LSP ID : Dut-E.01-00
LSP ID : Dut-E.02-00
LSP ID : Dut-F.00-00
```

```
Router Cap : 10.20.1.6, D:0, S:0
TE Node Cap : B E M P
SR Cap: IPv4 MPLS-IPv6
SRGB Base:20000, Range:20000
SR Alg: metric based SPF
LSP ID : Dut-F.01-00
LSP ID : Dut-F.03-00
Level (2) Capability Count : 14
```

```
=====
```



## database

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database</b> [ <i>system-id</i>   <i>lsp-id</i> ] [ <b>detail</b> ] [ <b>level</b> <i>level</i> ]                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command displays the entries in the IS-IS link state database.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><i>system-id</i> — Only the LSPs related to the specified <i>system-id</i> are listed. If no <i>system-id</i> or <i>lsp-id</i> are specified, all database entries are listed.</p> <p><i>lsp-id</i> — Only the specified LSP (hostname) is listed. If no <i>system-id</i> or <i>lsp-id</i> are specified, all database entries are listed.</p> <p><i>level</i> — Specifies the interface level (1, 2, or 1 and 2).</p> <p><b>detail</b> — All output is displayed in the detailed format.</p> |

### Output

#### Sample Output

```
*A:ALA-A# show router isis database
=====
Rtr Base ISIS Instance 0 Database
=====
LSP ID Sequence Checksum Lifetime Attributes

Displaying Level 1 database

abr_dfw.00-00 0x50 0x164f 603 L1L2
Level (1) LSP Count : 1
Displaying Level 2 database

asbr_east.00-00 0x53 0xe3f5 753 L1L2
abr_dfw.00-00 0x57 0x94ff 978 L1L2
abr_dfw.03-00 0x50 0x14f1 614 L1L2
Level (2) LSP Count : 3
=====
*A:ALA-A#

*A:Dut-B# show router isis database Dut-A.00-00 detail
=====
Rtr Base ISIS Instance 0 Database (detail)
=====
Displaying Level 1 database

Level (1) LSP Count : 0

Displaying Level 2 database

LSP ID : Dut-A.00-00 Level : L2
Sequence : 0x6 Checksum : 0xb7c4 Lifetime : 1153
Version : 1 Pkt Type : 20 Pkt Ver : 1
Attributes: L1L2 Max Area : 3
SysID Len : 6 Used Len : 311 Alloc Len : 311
```

```

TLVs :
 Area Addresses:
 Area Address : (2) 30.31
 Supp Protocols:
 Protocols : IPv4
 IS-Hostname : Dut-A
 Router ID :
 Router ID : 10.20.1.1
 I/F Addresses :
 I/F Address : 10.20.1.1
 I/F Address : 10.10.1.1
 I/F Address : 10.10.2.1
 TE IS Nbrs :
 Nbr : Dut-B.01
 Default Metric : 1000
 Sub TLV Len : 98
 IF Addr : 10.10.1.1
 MaxLink BW: 100000 kbps
 Resvble BW: 100000 kbps
 Unresvd BW:
 BW[0] : 10000 kbps
 BW[1] : 40000 kbps
 BW[2] : 40000 kbps
 BW[3] : 40000 kbps
 BW[4] : 50000 kbps
 BW[5] : 50000 kbps
 BW[6] : 50000 kbps
 BW[7] : 10000 kbps
 Admin Grp : 0x0
 TE Metric : 1000
 SUBTLV BW CONSTS : 8
 BW Model : 1
 BC[0] : 10000 kbps
 BC[1] : 0 kbps
 BC[2] : 40000 kbps
 BC[3] : 0 kbps
 BC[4] : 0 kbps
 BC[5] : 50000 kbps
 BC[6] : 0 kbps
 BC[7] : 0 kbps
 TE IP Reach :
 Default Metric : 0
 Control Info: , prefLen 32
 Prefix : 10.20.1.1
 Default Metric : 1000
 Control Info: , prefLen 24
 Prefix : 10.10.1.0
 Default Metric : 1000
 Control Info: , prefLen 24
 Prefix : 10.10.2.0

```

```

Level (2) LSP Count : 1

```

```

Flags : D = Prefix Leaked Down
 : N = Node Flag
 : R = Re-advertisement Flag
 : S = Sub-TLVs Present
 : X = External Prefix Flag

```

```
=====
*A:Dut-B#
```

## hostname

|                    |                                                                                  |
|--------------------|----------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hostname</b>                                                                  |
| <b>Context</b>     | show>router>isis                                                                 |
| <b>Description</b> | This command displays the hostname database. There are no options or parameters. |
| <b>Output</b>      | IS-IS Hostname Output                                                            |

The following table describes output fields for IS-IS hostname output.

**Table 38 IS-IS Hostname Output Fields**

| Label     | Description                                  |
|-----------|----------------------------------------------|
| System-id | System identifier mapped to hostname.        |
| Hostname  | Hostname for the specific <i>system-id</i> . |
| Type      | The type of entry (static or dynamic).       |

### Sample Output

```
A:ALA-A# show router isis hostname
=====
Rtr Base ISIS Instance 0 Hostnames
=====
System Id Hostname

1800.0000.0002 core_west
1800.0000.0005 core_east
1800.0000.0008 asbr_west
1800.0000.0009 asbr_east
1800.0000.0010 abr_sjc
1800.0000.0011 abr_lax
1800.0000.0012 abr_nyc
1800.0000.0013 abr_dfw
1800.0000.0015 dist_oak
1800.0000.0018 dist_nj
1800.0000.0020 acc_nj
1800.0000.0021 acc_ri
1800.0000.0027 dist_arl
1800.0000.0028 dist_msq
1800.0000.0029 acc_arl
1800.0000.0030 acc_msq
=====
A:ALA-A#
```

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ] [ <b>detail</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command shows IS-IS interface information. When no <i>ip-addr</i> or the <i>ip-int-name</i> is specified, all interfaces are listed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>ip-address</i> — Only displays the interface information associated with the specified IP address.</p> <p><b>Values</b></p> <p>ipv4-address</p> <ul style="list-style-type: none"> <li>• a.b.c.d (host bits must be 0)</li> </ul> <p>ipv6-address</p> <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> <li>• d: [0 to 255]D</li> </ul> <p><i>ip-int-name</i> — Only displays the interface information associated with the specified IP interface name.</p> <p><b>detail</b> — All output is displayed in the detailed format.</p> <p><i>mt-id-number</i> — Specifies the multi topology ID number for this route.</p> |
| <b>Output</b>      | IS-IS Interface Output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

The following table describes IS-IS interface output fields.

**Table 39 IS-IS Interface Output Fields**

| Label        | Description                                                                            |
|--------------|----------------------------------------------------------------------------------------|
| Interface    | The interface name.                                                                    |
| Level        | Specifies the interface level (1, 2, or 1 and 2).                                      |
| CirID        | Specifies the circuit identifier.                                                      |
| Oper State   | Up — The interface is operationally up.<br>Down — The interface is operationally down. |
| L1/L2 Metric | Interface metric for Level 1 and Level 2, if none are set to 0.                        |

### Sample Output

```
A:ALA-A# show router isis interface
=====
Rtr Base ISIS Instance 0 Interfaces
=====
```

| Interface | Level | CircID | Oper State | L1/L2 Metric |
|-----------|-------|--------|------------|--------------|
| system    | L1L2  | 1      | Up         | 10/10        |
| if2/1     | L2    | 8      | Up         | -/10         |
| if2/2     | L1    | 5      | Up         | 10/-         |
| if2/3     | L1    | 6      | Up         | 10/-         |
| if2/4     | L1    | 7      | Up         | 10/-         |
| if2/5     | L2    | 2      | Up         | -/10         |
| lag-1     | L2    | 3      | Up         | -/10         |
| if2/8     | L2    | 4      | Up         | -/10         |

Interfaces : 8

A:ALA-A#

\*A:Dut-C# show router isis interface "system" detail

Rtr Base ISIS Instance 0 Interfaces

```

Interface : system Level Capability: L1L2
Oper State : Up Admin State : Up
Auth Keychain : Disabled
Auth Type : None Auth State : Enabled
Circuit Id : 1 Retransmit Int. : 5
Type : Pt-to-Pt LSP Pacing Int. : 100
Oper Type : Pt-to-Pt CSNP Int. : 10
Mesh Group : Inactive BER : none
LFA NH Template : None Bfd Enabled : No
Topology : IPv4-Unicast, IPv6-Unicast, IPv4-Multicast, IPv6-Multicast
Te Metric : 0 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0
Route Tag : None LFA : Included
Default Instance: N/A LFA :
IPv4 Node SID : Index 1003 IPv6 Node SID : none

Level : 1 Adjacencies : 0
Auth Keychain : Disabled
Auth Type : None Metric : 0
Hello Timer : 3 IPv6-Ucast-Met : 0
Priority : 64 IPv6-Mcast-Met : 0
Passive : No IPv4-Mcast-Met : 0
SD-Offset : 0 SF-Offset : 0
Hello Mult. : 2

Level : 2 Adjacencies : 0
Auth Keychain : Disabled
Auth Type : None Metric : 0
Hello Timer : 3 IPv6-Ucast-Met : 0
Priority : 64 IPv6-Mcast-Met : 0
Passive : No IPv4-Mcast-Met : 0
SD-Offset : 0 SF-Offset : 0
Hello Mult. : 2

```

## lfa-coverage

|                    |                                                                 |
|--------------------|-----------------------------------------------------------------|
| <b>Syntax</b>      | <b>lfa-coverage</b>                                             |
| <b>Context</b>     | show>router>isis                                                |
| <b>Description</b> | This command displays IS-IS LFA coverage information.           |
| <b>Output</b>      | The following output is an example of LFA coverage information. |

### Sample Output

```
*A:SR# show router isis lfa-coverage
=====
Rtr Base ISIS Instance 0 LFA Coverage
=====
Topology Level Node IPv4

IPV4 Unicast L1 4/4 (100%) 826/826 (100%)
IPV4 Unicast L2 2/2 (100%) 826/826 (100%)
IPV6 Unicast L1 3/3 (100%) 0/0 (0%)
IPV6 Unicast L2 0/0 (0%) 0/0 (0%)
=====
*A:SR#

*A:SRR>config>router>isis# show router isis lfa-coverage
=====
LFA Coverage
=====
Topology Level Node IPv4 IPv6

IPV4 Unicast L1 3/4 (75%) 1484/1975 (75%) 0/0 (0%)
IPV4 Unicast L2 3/3 (100%) 1484/1975 (75%) 0/0 (0%)
=====
*A:SRR>config>router>isis#
```

## link-group-member-status

|                    |                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>link-group-member-status</b> <i>name</i> [ <i>level level</i> ]                                     |
| <b>Context</b>     | show>router>isis                                                                                       |
| <b>Description</b> | This command displays IS-IS link-group-member status.                                                  |
| <b>Parameters</b>  | <i>name</i> — Up to 32 characters.<br><i>level</i> — Specifies the interface level (1, 2, or 1 and 2). |
| <b>Output</b>      | The following output is an example of link group member status information.                            |

### Sample Output

```
A:cses-V94# show router isis link-group-member-status
- link-group-member-status <name> [level <level>]

<name> : [32 chars max]
<level> : 1|2

A:cses-V94# show router isis link-group-member-status "toDutB"

=====
Rtr Base ISIS Instance 0 Link-Group Member
=====
Link-group I/F name Level State

toDutB ip-10.10.12.3 L1 Up
toDutB ip-10.10.3.3 L1 Up
toDutB ip-10.10.12.3 L2 Up
toDutB ip-10.10.3.3 L2 Up

Legend: BER = bitErrorRate
=====
A:cses-V94#
```

## link-group-status

- Syntax** `link-group-status name [level level]`
- Context** `show>router>isis`
- Description** This command displays IS-IS link-group status.
- Parameters** *name* — Specifies the link-group name.  
*level* — Specifies the interface level (1, 2, or 1 and 2).
- Output** The following output is an example of link group status.

### Sample Output

```
A:cses-V94# show router isis link-group-status

=====
Rtr Base ISIS Instance 0 Link-Group Status
=====
Link-group Mbrs Oper Revert Active Level State
 Mbr Mbr Mbr Mbr

toDutB 2 2 2 2 L1 normal
toDutB 2 2 2 2 L2 normal
toDutE 2 2 2 2 L1 normal
toDutE 2 2 2 2 L2 normal
=====
A:cses-V94#
```

## mapping-server

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mapping-server</b> [ <b>prefix</b> <i>ip-address[/mask]</i> ] [ <b>index</b> <i>index</i> ] [ <b>level</b> <i>level</i> ] [ <b>flag</b> { <i>s</i> }]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command displays IS-IS mapping-server information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>ip-address[/mask]</i> — Specifies the IP address and mask of a prefix that has received a node-sid in a SID/Label binding TLV.</p> <p><b>Values</b>     <i>ip-address</i>: <b>a.b.c.d.</b> (host bits must be 0)<br/>                  <i>mask</i>: <b>0 to 32</b></p> <p><i>index</i> — Specifies the node-sid index value for the generated SID/Label binding TLV.</p> <p><b>Values</b>     0 to 4294967295</p> <p><b>Default</b>     none</p> <p><i>level</i> — Specifies a match on the mapping server's own flooding scope for the generated SID/Label binding TLV.</p> <p><b>Values</b>     1, 2, 1/2</p> <p><b>flag</b> — Specifies a match on the flooding scope of the generated SID/Label binding TLV.</p> <p><b>Values</b>     s — Specifies to match on the S flag value of 1. A SID/Label Binding TLV with the S flag set is flooded across the entire IS-IS routing domain, except across another IS-IS instance. If the S flag is not set (value of zero), the SID/Label Binding TLV is not leaked between levels.</p> |
| <b>Output</b>      | The following output is an example of mapping service information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Sample Output

```
*A:Dut-C# show router isis mapping-server
=====
Rtr Base ISIS Instance 0 Mapping Server
=====
Index Prefix Range Flags Level

1000 10.20.1.4/32 1 - L1L2
1001 10.20.1.5/32 1 - L1L2
1002 10.20.1.6/32 1 - L1L2

No. of Mapping Server Sid-Maps : 3
=====
```

## prefix-sids

|               |                                                                                                                                                                                                                                                                                                             |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>prefix-sids</b> [ <b>ipv4-unicast</b>   <b>ipv6-unicast</b>   <b>ipv4-multicast</b>   <b>ipv6-multicast</b>   <b>mt</b> <i>mt-id-number</i> ] [ <b>ip-prefix</b> / <i>prefix-length</i> ]] [ <b>sid</b> <i>sid</i> ] [ <b>adv-router</b> <i>system-id</i>   <i>hostname</i> ] [ <b>srms</b>   <b>no-</b> |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



**srms]**

|                    |                                                               |
|--------------------|---------------------------------------------------------------|
| <b>Context</b>     | show>router>isis                                              |
| <b>Description</b> | This command displays IS-IS prefix SIDs.                      |
| <b>Output</b>      | The following output is an example of prefix SID information. |

**Sample Output**

```
*A:Dut-C# show router isis prefix-sids
=====
Rtr Base ISIS Instance 0 Prefix/SID Table
=====
```

| Prefix       | SID  | Lvl/Typ | SRMS<br>MT | AdvRtr<br>Flags |
|--------------|------|---------|------------|-----------------|
| 10.0.0.1/32  | 1    | 2/Int.  | N          | Dut-B           |
|              |      |         | 0          | RNnP            |
| 10.0.0.1/32  | 1    | 2/Int.  | N          | Dut-C           |
|              |      |         | 0          | RNnP            |
| 10.0.0.1/32  | 1    | 1/Int.  | N          | Dut-D           |
|              |      |         | 0          | NnP             |
| 10.0.0.1/32  | 1    | 2/Int.  | N          | Dut-D           |
|              |      |         | 0          | NnP             |
| 10.0.0.1/32  | 1    | 2/Int.  | N          | Dut-E           |
|              |      |         | 0          | RNnP            |
| 10.20.1.2/32 | 1002 | 1/Int.  | N          | Dut-B           |
|              |      |         | 0          | NnP             |
| 10.20.1.2/32 | 1002 | 2/Int.  | N          | Dut-B           |
|              |      |         | 0          | NnP             |
| 10.20.1.2/32 | 1002 | 2/Int.  | N          | Dut-C           |
|              |      |         | 0          | RNnP            |
| 10.20.1.2/32 | 1002 | 2/Int.  | N          | Dut-D           |
|              |      |         | 0          | RNnP            |
| 10.20.1.2/32 | 1002 | 2/Int.  | N          | Dut-E           |
|              |      |         | 0          | RNnP            |
| 10.20.1.3/32 | 1003 | 2/Int.  | N          | Dut-B           |
|              |      |         | 0          | RNnP            |
| 10.20.1.3/32 | 1003 | 1/Int.  | N          | Dut-C           |
|              |      |         | 0          | NnP             |
| 10.20.1.3/32 | 1003 | 2/Int.  | N          | Dut-C           |
|              |      |         | 0          | NnP             |
| 10.20.1.3/32 | 1003 | 2/Int.  | N          | Dut-D           |
|              |      |         | 0          | RNnP            |
| 10.20.1.3/32 | 1003 | 2/Int.  | N          | Dut-E           |
|              |      |         | 0          | RNnP            |
| 10.20.1.4/32 | 1004 | 2/Int.  | N          | Dut-B           |
|              |      |         | 0          | RNnP            |
| 10.20.1.4/32 | 1004 | 2/Int.  | N          | Dut-C           |
|              |      |         | 0          | RNnP            |
| 10.20.1.4/32 | 1004 | 1/Int.  | N          | Dut-D           |
|              |      |         | 0          | NnP             |
| 10.20.1.4/32 | 1004 | 2/Int.  | N          | Dut-D           |
|              |      |         | 0          | NnP             |
| 10.20.1.4/32 | 1004 | 2/Int.  | N          | Dut-E           |
|              |      |         | 0          | RNnP            |
| 10.20.1.5/32 | 1005 | 2/Int.  | N          | Dut-B           |

---

|                        |       |        |   |         |      |
|------------------------|-------|--------|---|---------|------|
|                        |       |        |   | 0       | RNnP |
| 10.20.1.5/32           | 1005  | 2/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | RNnP |
| 10.20.1.5/32           | 1005  | 2/Int. | N | Dut - D |      |
|                        |       |        |   | 0       | RNnP |
| 10.20.1.5/32           | 1005  | 1/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | NnP  |
| 10.20.1.5/32           | 1005  | 2/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:1/128 | 10001 | 1/Int. | N | Dut - A |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:1/128 | 10001 | 2/Int. | N | Dut - A |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:1/128 | 10001 | 2/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:1/128 | 10001 | 2/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | RNnP |
| v:10:20:1:1/128        | 10001 | 2/Int. | N | Dut - F |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:2/128 | 10002 | 1/Int. | N | Dut - B |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:2/128 | 10002 | 2/Int. | N | Dut - B |      |
|                        |       |        |   | 0       | NnP  |
| v:10:20:1:2/128        | 10002 | 2/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:2/128 | 10002 | 2/Int. | N | Dut - D |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:2/128 | 10002 | 2/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:2/128 | 10002 | 2/Int. | N | Dut - F |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - A |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - B |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:3/128 | 10003 | 1/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - D |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:3/128 | 10003 | 2/Int. | N | Dut - F |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:4/128 | 10007 | 2/Int. | N | Dut - B |      |
|                        |       |        |   | 0       | RNnP |
| v:10:20:1:4/128        | 10007 | 2/Int. | N | Dut - C |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:4/128 | 10007 | 1/Int. | N | Dut - D |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:4/128 | 10007 | 2/Int. | N | Dut - D |      |
|                        |       |        |   | 0       | NnP  |
| 2001:db8:10:20:1:4/128 | 10007 | 2/Int. | N | Dut - E |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:4/128 | 10007 | 2/Int. | N | Dut - F |      |
|                        |       |        |   | 0       | RNnP |
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut - A |      |
|                        |       |        |   | 0       | RNnP |

---

|                        |       |        |   |       |
|------------------------|-------|--------|---|-------|
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:5/128 | 10005 | 1/Int. | N | Dut-E |
|                        |       |        | 0 | NnP   |
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | NnP   |
| 2001:db8:10:20:1:5/128 | 10005 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-A |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | RNnP  |
| 2001:db8:10:20:1:6/128 | 10006 | 1/Int. | N | Dut-F |
|                        |       |        | 0 | NnP   |
| 2001:db8:10:20:1:6/128 | 10006 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:1/128 | 10011 | 1/Int. | N | Dut-A |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:1/128 | 10011 | 2/Int. | N | Dut-A |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:1/128 | 10011 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:1/128 | 10011 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:1/128 | 10011 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:2/128 | 10012 | 1/Int. | N | Dut-B |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:2/128 | 10012 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:2/128 | 10012 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:2/128 | 10012 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:2/128 | 10012 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:2/128 | 10012 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | N | Dut-A |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:3/128 | 10013 | 1/Int. | N | Dut-C |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | N | Dut-E |

|                        |       |        |   |       |
|------------------------|-------|--------|---|-------|
| 2001:db8:20:20:1:3/128 | 10013 | 2/Int. | 0 | RNnP  |
|                        |       |        | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:4/128 | 10014 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:4/128 | 10014 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:4/128 | 10014 | 1/Int. | N | Dut-D |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:4/128 | 10014 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:4/128 | 10014 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:4/128 | 10014 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-A |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:5/128 | 10015 | 1/Int. | N | Dut-E |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:5/128 | 10015 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-A |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-B |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-C |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-D |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-E |
|                        |       |        | 0 | RNnP  |
| 2001:db8:20:20:1:6/128 | 10016 | 1/Int. | N | Dut-F |
|                        |       |        | 0 | NnP   |
| 2001:db8:20:20:1:6/128 | 10016 | 2/Int. | N | Dut-F |
|                        |       |        | 0 | NnP   |

-----  
No. of Prefix/SIDs: 25

Flags: R = Re-advertisement

N = Node-SID

nP = no penultimate hop POP

E = Explicit-Null

V = Prefix-SID carries a value

L = value/index has local significance

=====

\*A:Dut-C#

---

routes

|                    |                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>routes</b> [ipv4-unicast   ipv6-unicast   ipv4-multicast   ipv6-multicast   mt <i>mt-id-number</i> ] [ <i>ip-prefix[/prefix-length]</i> ] [alternative] [exclude-shortcut] [detail]                                                                                                                                                                                             |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command displays the routes in the IS-IS route table.                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <b>ipv4-unicast</b> — Displays IPv4 unicast parameters.<br><b>ipv6-unicast</b> — Displays IPv6 unicast parameters.<br><b>mt</b> <i>mt-id-number</i> — Displays multi-topology parameters.<br><b>Values</b> 0, 2<br><b>alternative</b> — Displays LFA details.<br><b>exclude-shortcut</b> — Displays the routes without shortcuts<br><b>detail</b> — Displays detailed information. |
| <b>Output</b>      | The following output is an example of IS-IS route information.                                                                                                                                                                                                                                                                                                                     |

**Sample Output**

```
*A:Dut-C# show router isis routes
=====
Rtr Base ISIS Instance 0 Route Table
=====
Prefix[Flags] Metric Lvl/Typ Ver. SysID/Hostname
NextHop MT AdminTag/SID[F]

1.1.2.0/24 20 1/Int. 0 Dut-A
 1.1.3.1 0
1.1.2.0/24 20 1/Int. 0 Dut-B
 1.2.3.2 0
1.1.3.0/24 10 1/Int. 0 Dut-C
 0.0.0.0 0
1.2.3.0/24 10 1/Int. 0 Dut-C
 0.0.0.0 0
1.2.4.0/24 20 1/Int. 0 Dut-B
 1.2.3.2 0
1.3.5.0/24 10 1/Int. 0 Dut-C
 0.0.0.0 0
1.4.5.0/24 20 1/Int. 0 Dut-E
 1.3.5.5 0
1.4.6.0/24 30 1/Int. 0 Dut-B
 1.2.3.2 0
1.4.6.0/24 30 1/Int. 0 Dut-E
 1.3.5.5 0
10.20.1.1/32 10 1/Int. 0 Dut-A
 1.1.3.1 0
10.20.1.2/32 10 1/Int. 0 Dut-B
 1.2.3.2 0
10.20.1.3/32 0 1/Int. 0 Dut-C
 0.0.0.0 0
```

|              |    |        |   |       |
|--------------|----|--------|---|-------|
| 10.20.1.4/32 | 20 | 1/Int. | 0 | Dut-B |
| 1.2.3.2      |    |        | 0 | 0     |
| 10.20.1.4/32 | 20 | 1/Int. | 0 | Dut-E |
| 1.3.5.5      |    |        | 0 | 0     |
| 10.20.1.5/32 | 10 | 1/Int. | 0 | Dut-E |
| 1.3.5.5      |    |        | 0 | 0     |
| 10.20.1.6/32 | 30 | 1/Int. | 0 | Dut-B |
| 1.2.3.2      |    |        | 0 | 0     |
| 10.20.1.6/32 | 30 | 1/Int. | 0 | Dut-E |
| 1.3.5.5      |    |        | 0 | 0     |

-----  
No. of Routes: 17 (17 paths)  
-----

Flags : L = LFA nexthop available  
SID[F] : R = Re-advertisement  
N = Node-SID  
nP = no penultimate hop POP  
E = Explicit-Null  
V = Prefix-SID carries a value  
L = value/index has local significance

=====

\*A:Dut-C#

\*A:Dut-C# show router isis routes ipv4-unicast

=====

Rtr Base ISIS Instance 0 Route Table

| Prefix[Flags]<br>NextHop | Metric | Lvl/Typ | Ver.<br>MT | SysID/Hostname<br>AdminTag/SID[F] |
|--------------------------|--------|---------|------------|-----------------------------------|
| 1.2.3.0/24               | 10     | 1/Int.  | 8          | Dut-C                             |
| 0.0.0.0                  |        |         | 0          | 0                                 |
| 1.2.4.0/24               | 20     | 1/Int.  | 8          | Dut-D                             |
| 1.3.4.4                  |        |         | 0          | 0                                 |
| 1.2.5.0/24               | 20     | 1/Int.  | 8          | Dut-B                             |
| 1.2.3.2                  |        |         | 0          | 0                                 |
| 1.3.4.0/24               | 10     | 1/Int.  | 8          | Dut-C                             |
| 0.0.0.0                  |        |         | 0          | 0                                 |
| 1.3.5.0/24 [L]           | 30     | 1/Int.  | 11         | Dut-B                             |
| 1.2.3.2                  |        |         | 0          | 0                                 |
| 1.4.5.0/24               | 20     | 1/Int.  | 8          | Dut-D                             |
| 1.3.4.4                  |        |         | 0          | 0                                 |
| 2.3.4.0/24               | 40     | 1/Int.  | 11         | Dut-D                             |
| 1.3.4.4                  |        |         | 0          | 0                                 |
| 4.0.0.1/32               | 10     | 1/Int.  | 8          | Dut-D                             |
| 1.3.4.4                  |        |         | 0          | 0/1 [NnP]                         |
| 10.20.1.2/32             | 10     | 1/Int.  | 8          | Dut-B                             |
| 1.2.3.2                  |        |         | 0          | 0/1002 [NnP]                      |
| 10.20.1.3/32             | 0      | 1/Int.  | 5          | Dut-C                             |
| 0.0.0.0                  |        |         | 0          | 0/1003 [NnP]                      |
| 10.20.1.4/32             | 10     | 1/Int.  | 8          | Dut-D                             |
| 1.3.4.4                  |        |         | 0          | 0/1004 [NnP]                      |
| 10.20.1.5/32 [L]         | 20     | 1/Int.  | 11         | Dut-B                             |
| 1.2.3.2                  |        |         | 0          | 0/1005 [NnP]                      |
| 10.21.1.2/32             | 10     | 1/Int.  | 8          | Dut-B                             |
| 1.2.3.2                  |        |         | 0          | 0                                 |
| 10.21.1.3/32             | 0      | 1/Int.  | 5          | Dut-C                             |
| 0.0.0.0                  |        |         | 0          | 0                                 |
| 10.21.1.4/32             | 10     | 1/Int.  | 8          | Dut-D                             |

```

1.3.4.4 0 0
10.21.1.5/32 [L] 20 1/Int. 11 Dut-B
1.2.3.2 0 0

No. of Routes: 16 (16 paths)

Flags : L = LFA nexthop available
SID[F] : R = Re-advertisement
 N = Node-SID
 nP = no penultimate hop POP
 E = Explicit-Null
 V = Prefix-SID carries a value
 L = value/index has local significance
=====

*A:Dut-C# show router isis routes ipv4-unicast alternative
=====
Rtr Base ISIS Instance 0 Route Table (alternative)
=====
Prefix[Flags] Metric Lvl/Typ Ver. SysID/Hostname
NextHop MT AdminTag/SID[F]
Alt-Nexthop Alt- Alt-Type
Metric

1.2.3.0/24 10 1/Int. 8 Dut-C
0.0.0.0 0 0
1.2.4.0/24 20 1/Int. 8 Dut-D
1.3.4.4 0 0
1.2.5.0/24 20 1/Int. 8 Dut-B
1.2.3.2 0 0
1.3.4.0/24 10 1/Int. 8 Dut-C
0.0.0.0 0 0
1.3.5.0/24 30 1/Int. 11 Dut-B
1.2.3.2 0 0
1.3.4.4(L) 30 NP
1.4.5.0/24 20 1/Int. 8 Dut-D
1.3.4.4 0 0
2.3.4.0/24 40 1/Int. 11 Dut-D
1.3.4.4 0 0
4.0.0.1/32 10 1/Int. 8 Dut-D
1.3.4.4 0 0/1 [NnP]
10.20.1.2/32 10 1/Int. 8 Dut-B
1.2.3.2 0 0/1002 [NnP]
10.20.1.3/32 0 1/Int. 5 Dut-C
0.0.0.0 0 0/1003 [NnP]
10.20.1.4/32 10 1/Int. 8 Dut-D
1.3.4.4 0 0/1004 [NnP]
10.20.1.5/32 20 1/Int. 11 Dut-B
1.2.3.2 0 0/1005 [NnP]
1.3.4.4(L) 20 NP
10.21.1.2/32 10 1/Int. 8 Dut-B
1.2.3.2 0 0
10.21.1.3/32 0 1/Int. 5 Dut-C
0.0.0.0 0 0
10.21.1.4/32 10 1/Int. 8 Dut-D
1.3.4.4 0 0
10.21.1.5/32 20 1/Int. 11 Dut-B
1.2.3.2 0 0
1.3.4.4(L) 20 NP

```

```

No. of Routes: 16 (18 paths)
Flags : L = Loop-Free Alternate nexthop
Alt-Type : LP = linkProtection, NP = nodeProtection
SID[F] : R = Re-advertisement
 N = Node-SID
 nP = no penultimate hop POP
 E = Explicit-Null
 V = Prefix-SID carries a value
 L = value/index has local significance
=====
*A:Dut-C#
```

The following table describes IS-IS route output fields.

**Table 40** IS-IS Route Output Fields

| Label     | Description                                                                        |
|-----------|------------------------------------------------------------------------------------|
| Prefix    | The route prefix and mask.                                                         |
| Metric MT | The route’s metric.                                                                |
| Lvl/Type  | Specifies the level (1 or 2) and the route type, Internal (Int) or External (Ext). |
| Version   | SPF version that generated route.                                                  |
| Nexthop   | System ID of nexthop, give hostname if possible.                                   |
| Hostname  | Hostname for the specific <i>system-id</i> .                                       |

spf-log

- Syntax** `spf-log [detail]`
- Context** `show>router>isis`
- Description** This command displays IS-IS SPF log information.
- Parameters** `detail` — Displays detailed spf-log information.
- Output** The following output is an example of SPF log information.

**Sample Output**

```
*A:Dut-C# show router isis spf-log

=====
Rtr Base ISIS Instance 0 SPF Log
=====
When Duration L1 Nodes L2 Nodes Event Count Type
```



```

01/26/2015 11:22:19 <0.01s - - - pLfa
01/26/2015 11:22:19 <0.01s - - - rLfa
01/26/2015 11:22:20 <0.01s 4 4 25 Reg
01/26/2015 11:22:20 <0.01s - - - Lfa
01/26/2015 11:22:20 <0.01s - - - rLfa
01/26/2015 11:22:24 <0.01s 4 4 11 Reg
01/26/2015 11:22:24 <0.01s - - - Lfa
01/26/2015 11:22:24 <0.01s - - - rLfa
01/26/2015 11:22:32 <0.01s 4 4 21 Reg
01/26/2015 11:22:32 <0.01s - - - Lfa
01/26/2015 11:22:32 <0.01s - - - rLfa
01/26/2015 11:22:33 <0.01s - - - pSpf
01/26/2015 11:22:33 <0.01s - - - pLfa
01/26/2015 11:22:33 <0.01s - - - rLfa
01/26/2015 11:22:41 <0.01s - - - pSpf
01/26/2015 11:22:41 <0.01s - - - pLfa
01/26/2015 11:22:41 <0.01s - - - rLfa
01/26/2015 11:22:51 <0.01s 4 4 4 Reg
01/26/2015 11:22:51 <0.01s - - - Lfa
01/26/2015 11:22:51 <0.01s - - - rLfa

Log Entries : 20
=====
*A:Dut-C#
```

A:SetupCLI# show router isis spf-log detail

Rtr Base ISIS Instance 0 SPF Log

```
=====
When : 10/01/2011 03:40:25 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 78
SPF Type : Reg
Reason : LSPCONTENT

When : 10/01/2011 03:40:26 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 1
SPF Type : Reg
Reason : LSPCONTENT

When : 10/01/2011 03:40:25 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 25
SPF Type : Reg
Reason : NEWAREA NEWREACH LSPCONTENT MANUALREQ

When : 10/01/2011 03:40:27 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 1
SPF Type : Reg
Reason : LSPCONTENT

When : 10/01/2011 03:40:27 Duration : <0.01s
L1 Nodes : 0 L2 Nodes : 0
Trigger LSP: SetupCLI.00-00 Event Count : 1
```

```

SPF Type : Lfa
Reason : LSPCONTENT

When : 10/01/2011 03:40:25 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 75
SPF Type : Reg
Reason : LSPCONTENT

When : 10/01/2011 03:40:27 Duration : <0.01s
L1 Nodes : 1 L2 Nodes : 1
Trigger LSP: SetupCLI.00-00 Event Count : 1
SPF Type : Reg
Reason : LSPCONTENT
=====
A:SetupCLI#

```

The following table describes the IS-IS SPF log output fields.

**Table 41** IS-IS SPF Log Output Fields

| Label       | Description                                                                     |
|-------------|---------------------------------------------------------------------------------|
| When        | Displays the timestamp when the SPF run started on the system.                  |
| Duration    | Displays the time (in hundredths of a second) required to complete the SPF run. |
| L1 Nodes    | Displays the number of Level 1 nodes involved in the SPF run.                   |
| L2 Nodes    | Displays the number of Level 2 nodes involved in the SPF run.                   |
| Event Count | Displays the number of SPF events that triggered the SPF calculation.           |
| Type        | Displays the SPF type, Reg (regular) or Lfa (loop free alternative).            |
| Trigger LSP | Displays the LSP that triggered the SPF run.                                    |

**Table 41 IS-IS SPF Log Output Fields (Continued)**

| Label  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reason | <p>Displays the reason(s) for the SPF run.</p> <p>NEWADJ: An adjacency changed.</p> <p>NEWLSP: A new LSP was received.</p> <p>NEWAREA: An area changed.</p> <p>NEWREACH: A prefix changed.</p> <p>ECMPCHANGED: An ECMP path changed.</p> <p>NEWMETRIC: A prefix metric changed.</p> <p>RESTART: The graceful restart exited.</p> <p>LSPEXPIRED: An LSP expired.</p> <p>DBCHANGED: The LSP database was cleared by an administrator.</p> <p>LSPCONTENT: The content of an LSP changed.</p> <p>NEWPREF: The external route preference changed.</p> <p>NEWNLPID: The routed protocols (IPv4 or IPv6) changed.</p> <p>MANUALREQ: An SPF calculation was requested by an administrator.</p> <p>ADMINTAGCHANGED: An administrative tag changed.</p> <p>TUNNELCHANGED: An MPLS tunnel changed.</p> |

## sr-adj-sets

- Syntax** **sr-adj-sets** [*adjacency-set-id*] [**detail**]
- Context** show>router>isis
- Description** This command displays IS-IS SR adjacency set log information.
- Parameters** *adjacency-set-id* — Displays SR adjacency set ID information.
- Values** 1 to 4294967295
- detail** — Displays detailed SR adjacency set information.
- Output** The following output is an example of **sr-adj-sets** log information.

### Sample Output

```
*A:Dut-C# show router isis sr-adj-sets

=====
Rtr Base ISIS Instance 0 Adjacency Sets
=====
ID State Label Label Members Members Neighbor
```

```

 Value Type Total Active System Id

1003 active 524285 DL 1 1 0100.2000.1001
1006 active 524284 DL 1 1 0100.2000.1005
1015 active 18448 SL 1 1 0100.2000.1001
1018 active 18451 SL 1 1 0100.2000.1005
1027 active 524283 DL 1 1 0100.2000.1001
1030 active 524282 DL 1 1 0100.2000.1005
1039 active 18460 SL 1 1 0100.2000.1001
1042 active 18463 SL 1 1 0100.2000.1005
1050 active 524281 DL 2 2 none
1056 active 18471 SL 2 2 none
1063 active 524280 DL 1 1 0100.2000.1001
1066 active 524279 DL 1 1 0100.2000.1005
1075 active 18478 SL 1 1 0100.2000.1001
1078 active 18481 SL 1 1 0100.2000.1005
1087 active 524278 DL 1 1 0100.2000.1001
1090 active 524277 DL 1 1 0100.2000.1005
1099 active 18490 SL 1 1 0100.2000.1001
1102 active 18493 SL 1 1 0100.2000.1005
1110 active 524276 DL 2 2 none
1116 active 18501 SL 2 2 none

DL: dynamic local
SL: static local

No. of SR Adjacency Sets: 20
=====

```

The following table describes the IS-IS **sr-adj-sets** output fields.

**Table 42** IS-IS SR Adjacency Set Output Fields

| Label                    | Description                                                      |
|--------------------------|------------------------------------------------------------------|
| ID                       | Displays the ID                                                  |
| State                    | Displays the state (active or inactive)                          |
| Label Value              | Displays the label value                                         |
| Label Type               | Displays the label type, Dynamic Local (DL) or Static Local (SL) |
| Members Total            | Displays the number of total members                             |
| Members Active           | Displays the number of active members                            |
| Neighbor System Id       | Displays the neighbor system ID                                  |
| No. of SR Adjacency Sets | Displays the number of SR adjacency sets                         |

## statistics

|                    |                                                                       |
|--------------------|-----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                                                     |
| <b>Context</b>     | show>router>isis                                                      |
| <b>Description</b> | This command displays information regarding IS-IS traffic statistics. |
| <b>Output</b>      | The following output is an example of IS-IS statistics output.        |

### Sample Output

```
*A:Dut-C# show router isis statistics

=====
Rtr Base ISIS Instance 0 Statistics
=====
ISIS Instance : 0
Purge Initiated : 0 LSP Regens. : 17
Sid SRGB err : 0 Sid dupl err : 0

CSPF Statistics
Requests : 0 Request Drops : 0
Paths Found : 0 Paths Not Found : 0

SPF Statistics
SPF Runs : 7
 Last runTimeStamp: 01/26/2015 11:22:50
Partial SPF Runs : 3
 Last runTimeStamp: 01/26/2015 11:22:51

LFA Statistics
LFA Runs : 7
 Last runTimeStamp: 01/26/2015 11:22:51
Partial LFA Runs : 3
 Last runTimeStamp: 01/26/2015 11:22:41

RLFA Statistics
RLFA Runs : 10
 Last runTimeStamp: 01/26/2015 11:22:51

PDU Type Received Processed Dropped Sent Retransmitted

LSP 164 164 0 151 0
IIH 146 146 0 147 0
CSNP 288 288 0 291 0
PSNP 71 71 0 74 0
Unknown 0 0 0 0 0
=====
*A:Dut-C#
```

This table describes IS-IS statistics output fields.

**Table 43 IS-IS Statistics Output Fields**

| Label           | Description                                                                                     |
|-----------------|-------------------------------------------------------------------------------------------------|
| Purge Initiated | The number of times purges have been initiated.                                                 |
| SPF Runs        | The number of times shortest path first calculations have been made.                            |
| LSP Regens      | The count of LSP regenerations.                                                                 |
| Requests        | The number of CSPF requests made to the protocol.                                               |
| Paths Found     | The number of responses to CSPF requests for which paths satisfying the constraints were found. |
| PDU Type        | The PDU type.                                                                                   |
| Received        | The count of link state PDUs received by this instance of the protocol.                         |
| Processed       | The count of link state PDUs processed by this instance of the protocol.                        |
| Dropped         | The count of link state PDUs dropped by this instance of the protocol.                          |
| Sent            | The count of link state PDUs sent out by this instance of the protocol.                         |
| Retransmitted   | The count of link state PDUs that had to be retransmitted by this instance of the protocol.     |

## status

**Syntax**     **status**

**Context**     show>router>isis

**Description**     This command displays information regarding IS-IS status.

**Output**     IS-IS Status Output

The following table describes IS-IS status output fields.

**Table 44 IS-IS Status Output Fields**

| Label     | Description                                     |
|-----------|-------------------------------------------------|
| System Id | The configured and operational IS-IS system ID. |

**Table 44 IS-IS Status Output Fields (Continued)**

| Label                     | Description                                                                                                                                                                     |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Router Id                 | The configured and operational IS-IS router ID.                                                                                                                                 |
| Admin State<br>Oper State | Up — IS-IS is administratively up.                                                                                                                                              |
|                           | Down — IS-IS is administratively down.                                                                                                                                          |
| Ipv4 Routing              | Enabled — IPv4 routing is enabled.<br>Disabled — IPv4 routing is disabled.                                                                                                      |
| Ipv6 Routing              | Disabled — IPv6 routing is disabled.<br>Enabled, Native — IPv6 routing is enabled.<br>Enabled, Multi-topology — Multi-topology TLVs for IPv6 routing is enabled.                |
| Multi-topology            | Disabled — Multi-topology TLVs for IPv6 routing is disabled.<br>Enabled — Multi-topology TLVs for IPv6 routing is enabled.                                                      |
| Last Enabled              | The date and time when IS-IS was last enabled in the router.                                                                                                                    |
| Level Capability          | The routing level for the IS-IS routing process.                                                                                                                                |
| Authentication Check      | True — All IS-IS mismatched protocol packets are rejected.<br>False — Authentication is performed on received IS-IS protocol packets but mismatched packets are not rejected.   |
| Authentication Type       | The method of authentication used to verify the authenticity of packets sent by neighboring routers on an IS-IS interface.                                                      |
| Traffic Engineering       | Enabled — TE is enabled for the router.<br>Disabled — TE is disabled so that TE metrics are not generated and are ignored when received by this node.                           |
| Graceful Restart          | Enabled — Graceful restart is enabled for this instance of IS-IS on the router.<br>Disabled — Graceful restart capability is disabled for this instance of IS-IS on the router. |
| Ldp Sync Admin State      | Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol.                                      |
| LFA NH Template           | Indicates the LFA template that is applied for the configured LFA policies.                                                                                                     |
| LFA Policies              | Indicates the configured LFA policies.                                                                                                                                          |
| Loopfree-Alternate        | Indicates if Loop-Free Alternate (LFA) is enabled or disabled in this IS-IS instance.                                                                                           |

**Sample Output**

```

=====
*A:Dut-A# show router isis status
=====
Rtr Base ISIS Instance 0 Status
=====
ISIS Cfg System Id : 0000.0000.0000
ISIS Oper System Id : 0100.2000.1001
ISIS Cfg Router Id : 0.0.0.0
ISIS Oper Router Id : 10.20.1.1
ASN : 0
Admin State : Up
Oper State : Up
Ipv4 Routing : Enabled
Ipv6 Routing : Disabled
Mcast Ipv4 Routing : Enabled, Native
Mcast Ipv6 Routing : Disabled
Last Enabled : 07/06/2017 21:29:26
Level Capability : L1L2
Authentication Check : True
Auth Keychain : Disabled
Authentication Type : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Enabled
Graceful Restart : Disabled
GR Helper Mode : Disabled
Overload-On-Boot Tim*: 0
Overload Max-Metric : False
Overload-On-Boot Max*: False
Ovl Export Interlevel: Disabled
Ovl Export External : Disabled
LSP Lifetime : 1200
LSP Refresh Half Int*: 600 (Config) 600 (Oper)
LSP Wait : 5000 ms (Max) 10 ms (Initial) 1000 ms (Second)
LSP MTU Size : 1492 (Config)
L1 LSP MTU Size : 1492 (Config) 1492 (Oper)
L2 LSP MTU Size : 1492 (Config) 1492 (Oper)
L1 MAX LSP MTU Size : 9187
L2 MAX LSP MTU Size : 1497
Adjacency Check : loose
L1 Auth Keychain : Disabled
L1 Auth Type : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticati*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference : 15
L1 Ext. Preference : 160
L1 Wide Metrics : Disabled
L1 LSDB Overload : Disabled
L1 LSPs : 1
L1 Default Metric : 10
L1 IPv6 Def Metric : 10
L1 Mcast IPv4 Def Me*: 10
L1 Mcast IPv6 Def Me*: 10
L1 Adv Router Cap : Enabled
Last SPF : 07/06/2017 21:30:10

```



---

```

SPF Wait : 10000 ms (Max) 1000 ms (Initial) 1000 ms (Second)
Multi-topology : Disabled
IPv6-Unicast MT2 : Disabled
IPv4-Multicast MT3 : Disabled
IPv6-Multicast MT4 : Disabled
Area Addresses : 30.31
Total Exp Routes(L1) : 0
Standard Multi-Insta*: Disabled
IID TLV : Disabled
Prefix Attributes TLV: Disabled
All-L1-MacAddr (Cfg) : 01:80:c2:00:00:14
L2 Auth Keychain : Disabled
L2 Auth Type : none
L2 CSNP-Authenticati*: Enabled
L2 HELLO-Authenticat*: Enabled
L2 PSNP-Authenticati*: Enabled
L2 Preference : 18
L2 Ext. Preference : 165
L2 Wide Metrics : Enabled
L2 LSDB Overload : Disabled
L2 LSPs : 6
L2 Default Metric : 10
L2 IPv6 Def Metric : 10
L2 Mcast IPv4 Def Me*: 10
L2 Mcast IPv6 Def Me*: 10
L2 Adv Router Cap : Enabled
Export Policies : None
Import Policies : None
LFA Policies : None
Multicast Import : None
Advertise-Passive-On*: Disabled
Ignore Attached Bit : Disabled
Suppress Attached Bit: Disabled
Default Route Tag : None
Rib Prio List High : None
Rib Prio Tag High : None
Ldp Sync Admin State : Up
LDP-over-RSVP : Disabled
IGP-Shortcut : Disabled
IPv4 IGP SC Tunn-Nhop: Res-Disabled
IPv6 IGP SC Tunn-Nhop: Res-Disabled
SRv4 IGP SC Tunn-Nhop: Res-Disabled
SRv6 IGP SC Tunn-Nhop: Res-Disabled
Advertise-Tunnel-Link: Disabled
Export Limit : 0
Exp Lmt Log Percent : 0
Total Exp Routes(L2) : 0
All-L2-MacAddr (Cfg) : 01:80:c2:00:00:15
Loopfree-Alternate : Disabled
Remote-LFA : Disabled
Max PQ Cost : 4261412864
TI-LFA : Disabled
Max SR FRR Labels : 2
L1 LFA : Included
L2 LFA : Included
Advertise Router Cap : Area
Hello Padding : Disabled
L1 Hello Padding : Disabled
L2 Hello Padding : Disabled

```

```

Ignore Lsp Errors : Disabled
Ignore Narrow Metric : Disabled
Reference Bandwidth : 0
Ucast Import Disable : None
Segment Routing : Up
Mapping Server : Down
Purge Originator Id : Disabled
Entropy Label : Enabled
Override ELC : Enabled

```

=====

\* indicates that the corresponding row element may have been truncated.

\*A:Dut-A#

The following sample output and detail sample output show LFA policies configured in the **configure router isis** context.

\*A:SRR# show router isis status

=====

Rtr Base ISIS Instance 0 Status

=====

```

ISIS Cfg System Id : 0000.COAS.0001
ISIS Oper System Id : 0000.COAS.0001
ISIS Cfg Router Id : 0.0.0.0
ISIS Oper Router Id : 10.20.1.3
Admin State : Up
Oper State : Up
Ipv4 Routing : Enabled
Ipv6 Routing : Disabled
Mcast Ipv4 Routing : Enabled, Native
Mcast Ipv6 Routing : Disabled
Last Enabled : 04/29/2014 15:14:33
Level Capability : L1
Authentication Check : True
Auth Keychain : Disabled
Authentication Type : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Disabled
Graceful Restart : Disabled
GR Helper Mode : Disabled
Overload-On-Boot Tim* : 0
Overload Max-Metric : False
Overload-On-Boot Max* : False
LSP Lifetime : 1200
LSP Refresh Interval : 600
LSP Wait : 5 sec (Max) 0 sec (Initial) 1 sec (Second)
LSP MTU Size : 1492 (Config) 1492 (Oper)
Adjacency Check : loose
L1 Auth Keychain : Disabled
L1 Auth Type : none
L1 CSNP-Authenticati* : Enabled
L1 HELLO-Authenticati* : Enabled
L1 PSNP-Authenticati* : Enabled
L1 Preference : 15
L1 Ext. Preference : 160
L1 Wide Metrics : Enabled
L1 LSDB Overload : Disabled

```

```

L1 LSPs : 5
L1 Default Metric : 10
L1 IPv6 Def Metric : 10
L1 Mcast IPv4 Def Me*: 10
L1 Mcast IPv6 Def Me*: 10
Last SPF : 04/29/2014 15:22:13
SPF Wait : 10 sec (Max) 1000 ms (Initial) 1000 ms (Second)
Multi-topology : Disabled
IPv6-Unicast MT2 : Disabled
IPv4-Multicast MT3 : Disabled
IPv6-Multicast MT4 : Disabled
Area Addresses : 49.0001
Total Exp Routes(L1) : 1
IID TLV : Disabled
Prefix Attributes TLV: Disabled
All-L1-MacAddr : 01:80:c2:00:00:14
L2 Auth Keychain : Disabled
L2 Auth Type : none
L2 CSNP-Authenticati*: Enabled
L2 HELLO-Authenticati*: Enabled
L2 PSNP-Authenticati*: Enabled
L2 Preference : 18
L2 Ext. Preference : 165
L2 Wide Metrics : Disabled
L2 LSDB Overload : Disabled
L2 LSPs : 0
L2 Default Metric : 10
L2 IPv6 Def Metric : 10
L2 Mcast IPv4 Def Me*: 10
L2 Mcast IPv6 Def Me*: 10
Export Policies : static
LFA Policies : pol1
 : pol2
 : pol3
 : pol4
 : pol5
Multicast Import : None
Advertise-Passive-On*: Disabled
Suppress Default : Disabled
Default Route Tag : None
Ldp Sync Admin State : Up
LDP-over-RSVP : Disabled
IGP-Shortcut : Disabled
Advertise-Tunnel-Link: Disabled
Export Limit : 0
Exp Lmt Log Percent : 0
Total Exp Routes(L2) : 0
All-L2-MacAddr : 01:80:c2:00:00:15
Loopfree-Alternate : Enabled
L1 LFA : Included
L2 LFA : Included
Advertise Router Cap : disable
Hello Padding : disable
=====
* indicates that the corresponding row element may have been truncated.
*A:SRR#

*A:SRR# show router isis interface "DUTC_TO_DUTE.1.0" detail

```

```

=====
Rtr Base ISIS Instance 0 Interfaces (detail)
=====

Interface : DUTC_TO_DUTE.1.0 Level Capability: L1L2
Oper State : Up Admin State : Up
Auth Keychain : Disabled
Auth Type : None Auth State : Enabled
Circuit Id : 3 Retransmit Int. : 5
Type : Broadcast LSP Pacing Int. : 100
Oper Type : Broadcast CSNP Int. : 10
Mesh Group : Inactive BER : none
LFA NH Template : "template1" Bfd Enabled : No
Topology : IPv4-Unicast, IPv6-Unicast, IPv4-Multicast, IPv6-Multicast
Te Metric : 0 Te State : Down
Admin Groups : None
Ldp Sync : outOfService Ldp Sync Wait : Disabled
Ldp Timer State : Disabled Ldp Tm Left : 0
Route Tag : None LFA : Included

Level : 1 Adjacencies : 0
Desg. IS : Dut-C
Auth Keychain : Disabled
Auth Type : None Metric : 10
Hello Timer : 9 IPv6-Ucast-Met : 10
Priority : 64 IPv6-Mcast-Met : 10
Passive : No IPv4-Mcast-Met : 10
SD-Offset : 0 SF-Offset : 0
Hello Mult. : 3

Level : 2 Adjacencies : 0
Desg. IS : Dut-C
Auth Keychain : Disabled
Auth Type : None Metric : 10
Hello Timer : 9 IPv6-Ucast-Met : 10
Priority : 64 IPv6-Mcast-Met : 10
Passive : No IPv4-Mcast-Met : 10
SD-Offset : 0 SF-Offset : 0
Hello Mult. : 3

=====
*A:SRR#

```

## summary-address

- Syntax** `summary-address [ip-address [/prefix-length]]`
- Context** `show>router>isis`
- Description** This command displays IS-IS summary addresses information.
- Output** Router IS-IS Summary Address Output
- The following table describes the IS-IS summary address output fields.

**Table 45 IS-IS Summary Address Output Fields**

| Label   | Description                                                           |
|---------|-----------------------------------------------------------------------|
| Address | The IP address.                                                       |
| Level   | Specifies the IS-IS level from which the prefix should be summarized. |

**Sample Output**

```
A:ALA-48# show router isis summary-address
=====
Rtr Base ISIS Instance 0 Summary Address
=====
Address Level

1.0.0.0/8 L1
2.1.0.0/24 L1L2
3.1.2.3/32 L2

Summary Addresses : 3
=====
A:ALA-48#
```

topology

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>topology</b> [ipv4-unicast   ipv6-unicast   ipv4-multicast   ipv6-multicast   mt <i>mt-id-number</i> ] [ <i>lfa</i> ] [detail]                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | show>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command shows IS-IS topology information.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <b>ipv4-unicast</b> — Displays IPv4 unicast parameters.<br><b>ipv6-unicast</b> — Displays IPv6 unicast parameters.<br><b>ipv4-multicast</b> — Displays IPv4 multicast parameters.<br><b>ipv6-multicast</b> — Displays IPv6 multicast parameters.<br><b>mt</b> <i>mt-id-number</i> — Displays multi-topology parameters.<br><b>Values</b> 0, 2, 3, 4<br><b>lfa</b> — Displays LFA (loop free alternative) information.<br><b>detail</b> — Displays detailed topology information. |
| <b>Output</b>      | Router IS-IS Topology Output<br><br>The following table describes the IS-IS topology output fields.                                                                                                                                                                                                                                                                                                                                                                              |

**Table 46 IS-IS Topology Output Fields**

| Label         | Description                          |
|---------------|--------------------------------------|
| Node          | Displays the IP address.             |
| Interface     | Displays the interface name.         |
| Nexthop       | Displays the nexthop IP address.     |
| LFA Interface | Displays the LFA interface name.     |
| LFA Nexthop   | Displays the LFA nexthop IP address. |

**Sample Output**

```

*A:Dut-A# show router isis topology
=====
Rtr Base ISIS Instance 0 Topology Table
=====
Node Interface Nexthop

IS-IS IP paths (MT-ID 0), Level 1

Dut-B.00 ip-3FFE::A0A:101 Dut-B
Dut-B.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.00 ip-3FFE::A0A:101 Dut-B
Dut-CA.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.02 ip-3FFE::A0A:101 Dut-B
Dut-CA.05 ip-3FFE::A0A:101 Dut-B
Dut-DA.00 ip-3FFE::A0A:101 Dut-B
Dut-DA.01 ip-3FFE::A0A:101 Dut-B
Dut-E.00 ip-3FFE::A0A:101 Dut-B
Dut-F.00 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.01 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.02 ies-1-3FFE::A0A:1501 Dut-F

IS-IS IPv6 paths (MT-ID 2), Level 1

Dut-B.00 ip-3FFE::A0A:101 Dut-B
Dut-B.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.00 ip-3FFE::A0A:101 Dut-B
Dut-CA.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.02 ip-3FFE::A0A:101 Dut-B
Dut-CA.05 ip-3FFE::A0A:101 Dut-B
Dut-DA.00 ip-3FFE::A0A:101 Dut-B
Dut-DA.01 ip-3FFE::A0A:101 Dut-B
Dut-E.00 ip-3FFE::A0A:101 Dut-B
Dut-F.00 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.01 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.02 ies-1-3FFE::A0A:1501 Dut-F

IS-IS IP paths (MT-ID 0), Level 2

Dut-B.00 ip-3FFE::A0A:101 Dut-B
Dut-B.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.00 ip-3FFE::A0A:101 Dut-B

```

```
Dut-CA.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.02 ip-3FFE::A0A:101 Dut-B
Dut-CA.05 ip-3FFE::A0A:101 Dut-B
Dut-DA.00 ip-3FFE::A0A:101 Dut-B
Dut-DA.01 ip-3FFE::A0A:101 Dut-B
Dut-E.00 ip-3FFE::A0A:101 Dut-B
Dut-F.00 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.01 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.02 ies-1-3FFE::A0A:1501 Dut-F

IS-IS IPv6 paths (MT-ID 2), Level 2

Dut-B.00 ip-3FFE::A0A:101 Dut-B
Dut-B.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.00 ip-3FFE::A0A:101 Dut-B
Dut-CA.01 ip-3FFE::A0A:101 Dut-B
Dut-CA.02 ip-3FFE::A0A:101 Dut-B
Dut-CA.05 ip-3FFE::A0A:101 Dut-B
Dut-DA.00 ip-3FFE::A0A:101 Dut-B
Dut-DA.01 ip-3FFE::A0A:101 Dut-B
Dut-E.00 ip-3FFE::A0A:101 Dut-B
Dut-F.00 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.01 ies-1-3FFE::A0A:1501 Dut-F
Dut-F.02 ies-1-3FFE::A0A:1501 Dut-F
=====
*A:Dut-A#
```

### 4.8.2.2 Clear Commands

#### isis

|                    |                                                                              |
|--------------------|------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>isis</b> [ <i>isis-instance</i> ]                                         |
| <b>Context</b>     | clear>router>isis                                                            |
| <b>Description</b> | This command enables the context to clear and reset IS-IS protocol entities. |
| <b>Parameters</b>  | <i>isis-instance</i> — Specifies the instance ID for the IS-IS instance.     |
| <b>Values</b>      | 1 to 31                                                                      |

#### adjacency

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>adjacency</b> [ <i>system-id</i> ]                                         |
| <b>Context</b>     | clear>router>isis                                                             |
| <b>Description</b> | This command clears and resets the entries from the IS-IS adjacency database. |

---

|                   |                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>system-id</i> — When the system ID is entered, only the specified entries are removed from the IS-IS adjacency database. |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------|

## database

|                    |                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>database</b> [ <i>system-id</i> ]                                                                                         |
| <b>Context</b>     | clear>router>isis                                                                                                            |
| <b>Description</b> | This command removes the entries from the IS-IS link-state database which contains information about PDUs.                   |
| <b>Parameters</b>  | <i>system-id</i> — When the system ID is entered, only the specified entries are removed from the IS-IS link-state database. |

## export

|                    |                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export</b>                                                                                                               |
| <b>Context</b>     | clear>router>isis                                                                                                           |
| <b>Description</b> | This command re-evaluates route policies participating in the export mechanism, either as importers or exporters of routes. |

## overload

|                    |                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>overload</b> {rtm   fib   prefix-limit}                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | clear>router>isis                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command clears the IS-IS overload.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <b>rtm</b> — Clears the overload because IS-IS reached the configured maximum route limit set with <b>maximum-routes</b> or <b>maximum-ipv6-routes</b> in a VPRN.<br><b>fib</b> — Clears the overload because adding routes to the hardware FIB failed.<br><b>prefix-limit</b> — Clears the overload when IS-IS has reached the configured maximum prefix limit set with the <b>prefix-limit</b> command. |

## spf-log

|                    |                                  |
|--------------------|----------------------------------|
| <b>Syntax</b>      | <b>spf-log</b>                   |
| <b>Context</b>     | clear>router>isis                |
| <b>Description</b> | This command clears the SPF log. |



statistics

|                    |                                                  |
|--------------------|--------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                                |
| <b>Context</b>     | clear>router>isis                                |
| <b>Description</b> | This command clears and resets IS-IS statistics. |

4.8.2.3 Debug Commands

isis

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>isis</b> [ <i>isis-instance</i> ]                                          |
| <b>Context</b>     | debug>router                                                                  |
| <b>Description</b> | This command enables the context to debug IS-IS protocol entities.            |
| <b>Parameters</b>  | <i>isis-instance</i> — Specifies the IS-IS instance.<br><b>Values</b> 1 to 31 |
| <b>Output</b>      | Sample output for the command is shown below.                                 |

Sample Output

```
*A:Dut-C# /tools dump router isis sr-database prefix 10.20.1.5 detail
=====
Rtr Base ISIS Instance 0 SR Database
=====
103 474390 10.20.1.5 LfaNhops 1 0 15 1000 1 1
 1492 1500 1500 0 0 1 1 0100.2000.1005 SR_ERR_OK
 IP:10.10.5.5 gifId:3 ifId:4 protectId:7 numLabels:1 outLbl:474390 isAdv:1 is
LfaX:0
 IP:10.10.12.2 gifId:5 ifId:6 protectId:0 numLabels:2 outLbl1:474389 outLbl2:
474390 numLfaNhops:1 isAdv:0

D = duplicate pending
xL = exclude from LFA
rL = remote LFA
Act = tunnel active
LDP = LDP FEC is the SID NH for SR-LDP stitching
=====

*A:Dut-C# /tools dump router isis sr-database nh-type ldp detail
=====
Rtr Base ISIS Instance 0 SR Database
=====
SID Label Prefix Last-act Lev MT TnlPref Metric IpNh SrNh
 Mtu MtuPrim MtuBk D xL rL Act AdvSystemId SrErr

```

```
1000 475287 10.20.1.4 AddTnl 1 0 15 0 1 1
 0 0 0 0 0 0 1 0100.2000.1004 SR_ERR_OK
 LDP: IP:10.20.1.4 tnlId:65546 tnlTyp:2
1001 475288 10.20.1.5 AddTnl 1 0 15 0 1 1
 0 0 0 0 0 0 1 0100.2000.1005 SR_ERR_OK
 LDP: IP:10.20.1.5 tnlId:65548 tnlTyp:2
1002 475289 10.20.1.6 AddTnl 1 0 15 0 1 1
 0 0 0 0 0 0 1 0100.2000.1006 SR_ERR_OK
 LDP: IP:10.20.1.6 tnlId:65549 tnlTyp:2

D = duplicate pending
xL = exclude from LFA
rL = remote LFA
Act = tunnel active
LDP = LDP FEC is the SID NH for SR-LDP stitching
=====
```

adjacency

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax      | [no] adjacency [ <i>ip-int-name</i>   <i>ip-address</i>   <i>nbr-system-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Context     | debug>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description | This command enables debugging for IS-IS adjacency.<br><br>The <b>no</b> form of the command disables debugging.                                                                                                                                                                                                                                                                                                                                                                                                         |
| Parameters  | <i>ip-address</i> — When specified, only adjacencies with the specified interface address are debugged.<br><br><b>Values</b> ipv4-address:<br>• a.b.c.d (host bits must be 0)<br>ipv6-address:<br>• x:x:x:x:x:x:x (eight 16-bit pieces)<br>• x:x:x:x:x:d.d.d.d<br>• x: [0 to FFFF]H<br>• d: [0 to 255]D<br><br><i>ip-int-name</i> — When specified, only adjacencies with the specified interface name are debugged.<br><br><i>nbr-system-id</i> — When specified, only the adjacency with the specified ID is debugged. |

cspf

|             |                                                |
|-------------|------------------------------------------------|
| Syntax      | [no] cspf                                      |
| Context     | debug>router>isis                              |
| Description | This command enables debugging for IS-IS cspf. |

The **no** form of the command disables debugging.

## graceful-restart

|                    |                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] graceful-restart</b>                                                                                            |
| <b>Context</b>     | debug>router>isis                                                                                                       |
| <b>Description</b> | This command enables debugging for IS-IS graceful-restart.<br><br>The <b>no</b> form of the command disables debugging. |

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i>   <i>ip-address</i> ]<br><b>no interface</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | debug>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command enables debugging for IS-IS interface.<br><br>The <b>no</b> form of the command disables debugging.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>ip-address</i> — When specified, only the interface with the specified interface address is debugged.<br><br><div style="margin-left: 40px;"> <b>Values</b> <div style="margin-left: 20px;">                     ipv4-address:                     <ul style="list-style-type: none"> <li>• a.b.c.d (host bits must be 0)</li> </ul>                     ipv6-address:                     <ul style="list-style-type: none"> <li>• x:x:x:x:x:x:x (eight 16-bit pieces)</li> <li>• x:x:x:x:x:d.d.d.d</li> <li>• x: [0 to FFFF]H</li> <li>• d: [0 to 255]D</li> </ul> </div> </div> <i>ip-int-name</i> — When specified, only the interface with the specified interface name is debugged. |

## leak

|                    |                                                     |
|--------------------|-----------------------------------------------------|
| <b>Syntax</b>      | <b>leak</b> [ <i>ip-address</i> ]<br><b>no leak</b> |
| <b>Context</b>     | debug>router>isis                                   |
| <b>Description</b> | This command enables debugging for IS-IS leaks.     |

The **no** form of the command disables debugging.

**Parameters** *ip-address* — When specified, only the specified address is debugged for IS-IS leaks.

**Values**

ipv4-address:

- a.b.c.d (host bits must be 0)

ipv6-address:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

## lsdb

**Syntax** [**no**] **lsdb** [*level-number*] [*system-id* | *lsp-id*]

**Context** debug>router>isis

**Description** This command enables debugging for Link State DataBase (LSDB).

The **no** form of the command disables debugging.

**Parameters** *system-id* — When specified, only the specified system-id is debugged. Host name up to 38 characters.

*lsp-id* — When specified, only the specified lsp-id is debugged. Hostname up to 38 characters.

*level-number* — Specifies the interface level (1, 2, or 1 and 2).

## misc

**Syntax** [**no**] **misc**

**Context** debug>router>isis

**Description** This command enables debugging for IS-IS misc.

The **no** form of the command disables debugging.

## packet

**Syntax** **packet** [*packet-type*] [*ip-int-name* | *ip-address*] [**detail**]

**Context** debug>router>isis

**Description** This command enables debugging for IS-IS packets.

The **no** form of the command disables debugging.

**Parameters** *ip-address* — When specified, only packets with the specified interface address are debugged.

**Values**     ipv4-address:  
                  • a.b.c.d (host bits must be 0)  
                  ipv6-address:  
                  • x:x:x:x:x:x:x (eight 16-bit pieces)  
                  • x:x:x:x:x:d.d.d.d  
                  • x: [0 to FFFF]H  
                  • d: [0 to 255]D

*ip-int-name* — When specified, only packets with the specified interface name are debugged.

*packet-type* — When specified, only packets of the specified type are debugged.

**Values**     ptop-hello | l1-hello | l2-hello | l1-psnp | l2-psnp | l1-csnp | l2-csnp |  
                  l1-lsp | l2-lsp

**detail** — All output is displayed in the detailed format.

## rtm

**Syntax**     **rtm** [*ip-address*]  
                  **no rtm**

**Context**     debug>router>isis

**Description**     This command enables debugging for IS-IS route table manager (RTM).

The **no** form of the command disables debugging.

**Parameters** *ip-address* — The specified IP address.

**Values**     ipv4-address:  
                  • a.b.c.d (host bits must be 0)  
                  ipv6-address:  
                  • x:x:x:x:x:x:x (eight 16-bit pieces)  
                  • x:x:x:x:x:d.d.d.d  
                  • x: [0 to FFFF]H  
                  • d: [0 to 255]D

## spf

**Syntax**     [**no**] **spf** [*level-number*] [*system-id*]

---

|                    |                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | debug>router>isis                                                                                                                                                                                         |
| <b>Description</b> | This command enables debugging for IS-IS SFP.<br><br>The <b>no</b> form of the command disables debugging.                                                                                                |
| <b>Parameters</b>  | <i>system-id</i> — When specified, only the specified system-id is debugged. A 6-octet system identifier (xxxx.xxxx.xxxx).<br><br><i>level-number</i> — Specifies the interface level (1, 2, or 1 and 2). |

## summary

|                    |                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>summary</b> [ <i>ip-address</i> ]<br><b>no summary</b>                                                                   |
| <b>Context</b>     | debug>router>isis                                                                                                           |
| <b>Description</b> | This command enables debugging for ISIS summary addresses.<br><br>The <b>no</b> form of the command disables the debugging. |
| <b>Parameters</b>  | <i>ip-address</i> — When specified, only packets with the specified address are debugged.                                   |

## tunnel-endpoint

|                    |                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-endpoint</b> [ <b>tunnel-spf</b> ] [ <b>tunnel-leak</b> <i>ip-address</i> ]<br><b>no tunnel-endpoint</b>                                  |
| <b>Context</b>     | debug>router>isis                                                                                                                                   |
| <b>Description</b> | This command enables debugging for an ISIS tunnel endpoint<br><br>The <b>no</b> form of the command disables the debugging.                         |
| <b>Parameters</b>  | <b>tunnel-spf</b> — Debugs tunnel SPF information.<br><br><i>ip-address</i> — When specified, only packets with the specified address are debugged. |

### 4.8.2.4 Tools Commands

## sr-adjacencies

|                |                                                                                    |
|----------------|------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>sr-adjacencies</b> [ <i>ip-int-name</i>   <i>ip-address</i> ] [ <b>detail</b> ] |
| <b>Context</b> | tools>dump>router>isis                                                             |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command displays ISIS SR adjacency information.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>ip-int-name</i> — Specifies the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>ip-address</i> — Specifies the interface's IP address.</p> <p><b>Values</b>      a.b.c.d</p> |

## sr-database

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sr-database</b> [ <i>sid sid</i> ] [ <i>prefix ip-address</i> ] [ <i>nh-type {sr   bgp   rsvp}</i> ] [ <i>ipv4-unicast   ipv6-unicast</i> ] [ <i>detail</i> ]                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | tools>dump>router>isis                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | This command displays ISIS SR database information.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>sid</i> — Displays information for the specific segment identifier.</p> <p><b>Values</b>      0 to 524287</p> <p><i>ip-address</i> — Specifies the IP address in dotted decimal notation.</p> <p><b>Values</b>      ipv4-address:<br/>                  • a.b.c.d<br/>                  ipv6-address:<br/>                  • x:x:x:x:x:x:x [eight 16-bit pieces]<br/>                  • x:x:x:x:x:d.d.d.d<br/>                  • x: [0 to FFFF]H<br/>                  • d: [0 to 255]D</p> |

## Output

### Sample Output

```
*A:Dut-E# tools dump router isis sr-database
=====
Rtr Base ISIS Instance 0 SR Database
=====
SID Label Prefix Last-act Lev MT RtmPref TtmPref Metric IpNh SrNh
 Mtu MtuPrim MtuBk D xL LT Act AdvSystemId SrErr

0 60000 1.0.33.3 LfaNhops 1 0 15 11 100 1 1
 1484 1500 1500 0 0 T +R 0100.2000.1003 SR_ERR_OK
1 60001 1.0.44.4 LfaNhops 1 0 15 11 200 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1004 SR_ERR_OK
2 60002 1.0.55.5 Local - - - - - - -
 - - - 0 - - +R 0100.2000.1005 SR_ERR_OK
```

---

```

3 60003 1.0.66.6 LfaNhops 1 0 15 11 100 1 1
 1484 1500 1500 0 0 T +R 0100.2000.1006 SR_ERR_OK
4 60004 1.0.11.1 LfaNhops 1 0 15 11 200 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1001 SR_ERR_OK
5 60005 1.0.22.2 LfaNhops 1 0 15 11 300 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1002 SR_ERR_OK
6 60006 10.20.1.3 LfaNhops 1 0 15 11 100 1 1
 1484 1500 1500 0 0 T +R 0100.2000.1003 SR_ERR_OK
7 60007 10.20.1.4 LfaNhops 1 0 15 11 200 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1004 SR_ERR_OK
8 60008 10.20.1.5 Local - - - - - - -
 - - - 0 - - +R 0100.2000.1005 SR_ERR_OK
9 60009 10.20.1.6 LfaNhops 1 0 15 11 100 1 1
 1484 1500 1500 0 0 T +R 0100.2000.1006 SR_ERR_OK
10 60010 10.20.1.1 LfaNhops 1 0 15 11 200 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1001 SR_ERR_OK
11 60011 10.20.1.2 LfaNhops 1 0 15 11 300 1 1
 1484 1500 1500 0 0 Tn +R 0100.2000.1002 SR_ERR_OK

```

No. of Entries: 12

-----

Lev = route level

IpNh = number of IP next-hops

SrNh = number of SR-tunnel next-hops

D = duplicate pending

xL = exclude from LFA

LT = LFA type (L:LFA, R:RLFA, T:TILFA, n:nodeProtection)

Act = tunnel active state (R:reported, F:failed, +:SR-ack)

=====



---

## 5 BGP

### 5.1 BGP Overview

Border Gateway Protocol (BGP) is an inter-Autonomous System routing protocol. An Autonomous System (AS) is a set of routers managed and controlled by a common technical administration. BGP-speaking routers establish BGP sessions with other BGP-speaking routers and use these sessions to exchange BGP routes. A BGP route provides information about a network path that can reach an IP prefix or other type of destination. The path information in a BGP route includes the list of ASes that must be traversed to reach the route source; this allows inter-AS routing loops to be detected and avoided. Other path attributes that may be associated with a BGP route include the Local Preference, Origin, Next-Hop, Multi-Exit Discriminator (MED) and Communities. These path attributes can be used to implement complex routing policies.

The primary use of BGP was originally Internet IPv4 routing but multi-protocol extensions to BGP have greatly expanded its applicability. Now BGP is used for many purposes, including:

- Internet IPv6 routing
- Inter-domain multicast support
- L3 VPN signaling (unicast and multicast)
- L2 VPN signaling (BGP auto-discovery for LDP-VPLS, BGP-VPLS, BGP-VPWS, multi-segment pseudowire routing, EVPN)
- Setup of inter-AS MPLS LSPs
- Distribution of flow specification rules (filters/ACLs)

The next sections provide information about BGP sessions, BGP network design, BGP messages and BGP path attributes.


## 5.2 BGP Sessions

A BGP session is a TCP connection formed between two BGP routers over which BGP messages are exchanged. There are three types of BGP sessions: internal BGP (IBGP), external BGP (EBGP), and confederation external BGP (confed-EBGP).

An IBGP session is formed when the two BGP routers belong to the same Autonomous System. Routes received from an IBGP peer are not advertised to other IBGP peers unless the router is a route reflector. The two routers that form an IBGP session are usually not directly connected. [Figure 30](#) shows an example of two Autonomous Systems that use BGP to exchange routes. In this example the router ALA-A forms IBGP sessions with ALA-B and ALA-C.

An EBGP session is formed when the two BGP routers belong to different Autonomous Systems. Routes received from an EBGP peer can be advertised to any other peer. The two routers that form an EBGP session are often directly connected but multi-hop EBGP sessions are also possible. When a route is advertised to an EBGP peer the Autonomous System number(s) of the advertising router are added to the AS Path attribute. In the example of [Figure 30](#) the router ALA-A forms an EBGP session with ALA-D.

**Figure 30 BGP Sessions**



OSRG053

A confederation EBGP session is formed when the two BGP routers belong to different member AS of the same confederation. More details about BGP confederations are provided in the section titled [BGP Confederations](#).

SR OS supports both statically configured and dynamic (unconfigured) BGP sessions. Dynamic sessions are supported by configuring one or more **prefix** commands in the **dynamic-neighbor** CLI context of a BGP group. Statically configured BGP sessions are configured using the **neighbor** command. This command accepts either an IPv4 or IPv6 address, which allows the session transport to be IPv4 or IPv6. By default, the router is the **active** side of TCP connections to statically configured remote peers, meaning that as soon as a session leaves the *idle* state, the router attempts to set up an outgoing TCP connection to the remote neighbor in addition to listening on TCP port 179 for an incoming connection from the peer. If required, a statically configured BGP session can be configured for **passive** mode so that the router only listens for an incoming connection and does not attempt to set up the outgoing connection. The router always operates in passive mode with respect to its dynamic (unconfigured) sessions.

The source IP address used to set up the TCP connection to the statically configured or dynamic peer can be configured explicitly using the **local-address** command. If a **local-address** is not configured then the source IP address is determined as follows:

- If the neighbor's IP address belongs to a local subnet the source IP address is this router's IP address on that subnet
- If the neighbor's IP address does not belong to a local subnet the source IP address is this router's system IP address

In addition, it is possible to configure the local address with the name of the router interface. To configure the BGP local address to use the router interface's IP address information, the **local-address** command is used in conjunction with the router interface name (*ip-int-name*) parameter.

Configuring the router interface as the local address is available in both the **config>router>bgp>group** context and the **config>router>bgp>group-neighbor** context.

When the router interface is configured as the local address, BGP inherits the address from the interface as follows:

- BGPv4 sessions — the primary IPv4 address configured on the interface is used as the local address
- BGPv6 sessions — the primary IPv6 address configured on the interface is used as the local address

If the corresponding IPv4 or IPv6 address is not configured on the router interface, the BGP sessions that have this interface set as the local address are kept down until an interface address is configured on the router interface.

If the primary IPv4 or IPv6 address is changed on the router interface and that interface is being used as the local address for BGP, then BGP will bounce the link. This will remove all routes advertised using the previous address and start advertising those routes again using the newly configured IP address.

## 5.2.1 BGP Session States

A BGP session is in one of the following states at any given moment in time:

- *Idle*. This is the state of a BGP session when it is administratively disabled with the **shutdown** command, indicated with (*Shutdown*). In this state no incoming TCP connection is accepted from the peer. When the session is administratively enabled it transitions out of the *Idle* state immediately. When the session is restarted automatically it may not leave the *Idle* state immediately if **damp-peer-oscillations** is configured, **damp-peer-oscillations** holds a session in the *Idle* state for exponentially increasing amounts of time if the session is unstable and resets frequently.
- *Connect*. This is the state of a BGP session when the router, acting in *Active* mode, is attempting to establish an outbound TCP connection with the remote peer.
- *Active*. This is the state of a BGP session when the router is listening for an inbound TCP connection attempt from the remote peer.
- *OpenSent*. This is the state of a BGP session when the router has sent an OPEN message to its peer in reaction to successful setup of the TCP connection and is waiting for an OPEN message from the peer.
- *OpenConfirm*. This is the state of a BGP session after the router has received an acceptable OPEN message from the peer and sent a KEEPALIVE message in response and is waiting for a KEEPALIVE message from the peer. TCP connection collision procedures may be performed at this stage. Refer to RFC 4271 for more details.
- *Established*. This is the state of a BGP session after the router has received a KEEPALIVE message from the peer. In this state BGP can advertise and withdraw routes by sending UPDATE messages to its peer.

## 5.2.2 Detecting BGP Session Failures

If a router suspects that its peer at the other end of an established session has experienced a complete failure of both its control and data planes the router should divert traffic away from the failed peer as quickly as possible in order to minimize traffic loss. There are various mechanisms that the router can use to detect such failures, including:

- BGP session hold timer expiry. See the section titled [Keepalive Message](#) for more details about this mechanism.
- Peer tracking
- BFD
- Fast external failover

When any one or these mechanisms is triggered the session immediately returns to the *Idle* state and a new session is attempted. Peer tracking, BFD and fast external failover are described in more detail in the following sections.

### 5.2.2.1 Peer Tracking

When peer tracking is enabled on a session the neighbor IP address is tracked in the routing table; if a failure occurs and there is no longer any IP route matching the neighbor address or else if the longest prefix match (LPM) route is rejected by the configurable **peer-tracking-policy** then after a 1 second delay the session is taken down. By default peer-tracking is disabled on all sessions. The default peer-tracking policy allows any type of route to match the neighbor IP address except aggregate routes and LDP shortcut routes.

Peer tracking was introduced when BFD was not yet supported for peer failure detection. Now that BFD is available peer-tracking has less value and is used less often.



**Note:** Peer tracking should be used with caution. Peer tracking can tear a session down even if the loss of connectivity turns out to be short-lived — for example while the IGP protocol is re-converging. Next-hop tracking, which is always enabled, handles such temporary connectivity issues much more effectively.

---

### 5.2.2.2 Bidirectional Forwarding Detection (BFD)

SR OS also supports the option to setup an async-mode BFD session to a BGP neighbor so that failure of the BFD session can trigger immediate teardown of the BGP session. When BFD is enabled on a BGP session a 1-hop or multi-hop BFD session is setup to the neighbor IP address and the BFD parameters come from the BFD configuration of the interface associated with the **local-address**; for multi-hop sessions this is typically the system interface. With a 10 ms transmit interval and a multiplier of 3 BFD can detect a peer failure in a period of time as short of 30 ms.

### 5.2.2.3 Fast External Failover

Fast external failover applies only to single-hop EBGp sessions. When fast external failover is enabled on a single-hop EBGp session and the interface associated with the session goes down the BGP session is immediately taken down as well, even if other mechanisms such as the hold-timer have not yet indicated a failure.

## 5.2.3 High Availability BGP Sessions

A BGP session reset can be very disruptive – each router participating in the failed session must delete the routes it received from its peer, recalculate new best paths, update forwarding tables (depending on the types of routes), and send route withdrawals and advertisements to other peers. It makes sense then that session resets should be avoided as much as possible and when a session reset cannot be avoided the disruption to the network should be minimized.

To support these objectives, the BGP implementation in SR OS supports two key features:

- BGP high availability (HA)
- BGP graceful restart (GR)

BGP HA refers to the capability of a router with redundant CPMs to keep established BGP sessions up whenever a planned or unplanned CPM switchover occurs. A planned CPM switchover can occur during In-Service Software Upgrade (ISSU). An unplanned CPM switchover can occur if there is an unexpected failure of the primary CPM.

BGP HA is always enabled on routers with redundant CPMs; it cannot be disabled. BGP HA keeps the standby CPM in-sync with the primary CPM, with respect to BGP and associated TCP state, so that the standby CPM is ready to take over for the primary CPM at any time. The primary CPM is responsible for building and sending the BGP messages to peers but the standby CPM reliably receives a copy of all outgoing UPDATE messages so that it has a synchronized view of the RIB-OUT.

### 5.2.3.1 BGP Graceful Restart

Some BGP routers do not have redundant control plane processor modules or do not support BGP HA with the same quality or coverage as 7450 ESS, 7750 SR, or 7950 XRS routers. When dealing with such routers or certain error conditions, BGP graceful restart (GR) is a good option for minimizing the network disruption caused by a control plane reset.

BGP GR assumes that the router restarting its BGP sessions has the ability and architecture to continue packet forwarding throughout the control plane reset. If this is the case, then the peers of the restarting router act as helpers and “hide” the control plane reset from the rest of the network so that forwarding can continue uninterrupted. Forwarding based on stale routes and hiding the “staleness” from other routers is considered acceptable because the duration of the control plane outage is expected to be relatively short (a few minutes). For BGP GR to be used on a session, both routers must advertise the BGP GR capability during the OPEN message exchange; see the [BGP Advertisement](#) section for more details.

BGP GR is enabled on one or more BGP sessions by configuring the **graceful-restart** command in the global, group, or neighbor context. The command causes GR mode to be supported for the following active families:

- IPv4 unicast
- IPv6 unicast
- VPN-IPv4
- VPN-IPv6
- Label-IPv4
- Label-IPv6
- L2-VPN
- route-target (RTC)
- Flow-IPv4 (IPv4 flow-spec)
- Flow-IPv6 (IPv6 flow-spec)

Helper mode is activated when one of the following events affects an *Established* session:

- TCP socket error
- New inbound TCP connection from the peer
- Hold timer expiry
- Peer unreachable
- BFD down
- Sent NOTIFICATION message (only if **enable-notification** is configured under **graceful-restart**, and the peer set the “N” bit in its GR capability, and the NOTIFICATION is not a *Cease* with subcode *Hard Reset*)
- Received NOTIFICATION message (only if **enable-notification** is configured under **graceful-restart**, and the peer set the “N” bit in its GR capability, and the NOTIFICATION is not a *Cease* with subcode *Hard Reset*)

As soon as the failure is detected, the helping 7450 ESS, 7750 SR, or 7950 XRS router marks all the routes received from the peer as stale and starts a restart timer. The stale state is not factored into the BGP decision process, and it is not made visible to other routers in the network. The restart timer derives its initial value from the Restart Time carried in the last GR capability of the peer. The default advertised Restart Time is 300 seconds, but it can be changed using the **restart-time** command.

When the restart timer expires, helping stops if the session is not yet re-established. If the session is re-established before the restart timer expires and the new GR capability from the restarting router indicates that the forwarding state has been preserved, then helping continues and the peers exchange routes per the normal procedure.

When each router has advertised all its routes for a specific address family, it sends an **End-of-RIB** marker (EOR) for the address family. The EOR is a minimal UPDATE message with no reachable or unreachable NLRI for the AFI or SAFI. When the helping router receives an EOR, it deletes all remaining stale routes of the AFI or SAFI that were not refreshed in the most recent set of UPDATE messages. The maximum amount of time that routes can remain stale (before being deleted if they are not refreshed) is configurable using the **stale-routes-time**.



**Note:** If a second reset occurs before GR has successfully completed, the router will always abort the GR helper process, regardless of the failure trigger.



### 5.2.3.2 BGP Long-Lived Graceful Restart

SR OS supports Long-Lived Graceful Restart (LLGR). LLGR is supported for the same address families as normal GR, as described in [BGP Graceful Restart](#).

The LLGR procedures adhere to *draft-uttaro-idr-bgp-persistence-03*. LLGR is intended to handle more serious and longer-term outages than ordinary GR.

SR OS routers support LLGR in the context of both the restarting router (which experienced a restart or failure) and the helper or receiving router (which is a peer of the failed router). Both functionalities are enabled and disabled at the same time by adding the **long-lived** command under a **graceful-restart** configuration context.

When **long-lived** is applied to a session (and capability negotiation is not disabled), the OPEN message sent to the peer includes both the GR capability and the LLGR capability. Both capabilities list the same set of AFI/SAFI.

#### 5.2.3.2.1 LLGR Operations

If a BGP session protected by LLGR goes down due to a restart or failure of the peer, then the SR OS router activates GR+LLGR helper mode for all the protected AFI/SAFI. In GR+LLGR helper mode, the received routes of a particular AFI/SAFI are retained as stale routes for a maximum duration of:

*restart-time + LLGR-stale-time*

where:

- *restart-time* is signaled in the GR capability of the peer (but overridden, if necessary, by the locally-configured **helper-override-restart-time** command)
- *LLGR-stale-time* is signaled in the LLGR capability of the peer (but overridden, if necessary, by a locally-configured **helper-override-stale-time** command)

While the restart-timer is running, the SR OS router acts in the normal GR helper role. When the restart-timer elapses, the LLGR phase begins. When LLGR starts, the following occur.

1. The *LLGR-stale-time* starts to count down.
2. Stale routes marked with the NO\_LLGR community are immediately deleted.
3. Remaining stale routes are not preferred. The BGP best path selection algorithm is rerun with a new first step that prefers valid, non-stale LLGR routes over any stale LLGR routes.

4. If a de-preferenced stale route remains, the best and valid NH-reachable path for the NLRI is re-advertised, with an added LLGR\_STALE community, to peers that signaled support for the LLGR capability. The route may be withdrawn or re-advertised towards peers that do not support LLGR, subject to the configuration of the **advertise-stale-to-all-neighbors** command and **without-no-export** parameter.

LLGR ends for a particular AFI/SAFI when the *LLGR-stale-time* reaches zero. At that time, all remaining stale routes of the AFI/SAFI are deleted. The *LLGR-stale-time* is not stopped by re-establishment of the session with the failed peer; it continues until the EoR marker is received for the AFI/SAFI.

Stale routes may be deleted prior to expiration of the *LLGR-stale-time*. If the session with the failed peer comes back up and one of the following is true, then the stale routes should be deleted immediately:

- the GR or LLGR capability is missing
- the AFI/SAFI is missing from the LLGR capability
- the F bit = 0 for the AFI/SAFI

#### 5.2.3.2.2 Receiving Routes with LLGR\_STALE Community

When a router running SR OS Release 15.0.R4 or later receives a BGP route of any AFI/SAFI, with the LLGR\_STALE community, the decision process considers the route less preferred than any valid, non-stale LLGR route for that NLRI. This logic applies even if the router is not configured as long-lived. If a route with an LLGR\_STALE community is selected as the best path, then it is advertised to peers according to the configuration of the **advertise-stale-to-all-neighbors** command; if this command is absent (or the **long-lived** context is absent), then the route is advertised only to peers that advertised the LLGR capability.

## 5.2.4 BGP Session Security

### 5.2.4.1 TCP MD5 Authentication

The operation of a network can be compromised if an unauthorized system is able to form or hijack a BGP session and inject control packets by falsely representing itself as a valid neighbor. This risk can be mitigated by enabling TCP MD5 authentication on one or more of the sessions. When TCP MD5 authentication is enabled on a session every TCP segment exchanged with the peer includes a TCP option (19)

containing a 16-byte MD5 digest of the segment (more specifically the TCP/IP pseudo-header, TCP header and TCP data). The MD5 digest is generated and validated using an authentication key that must be known to both sides. If the received digest value is different from the locally computed one then the TCP segment is dropped, thereby protecting the router from spoofed TCP segments.

#### 5.2.4.2 TTL Security Mechanism

The TTL security mechanism (GTSM) relies on a simple concept to protect BGP infrastructure from spoofed IP packets. It recognizes the fact that the vast majority of EBGP sessions are established between directly-connected routers and therefore the IP TTL values in packets belonging to these sessions should have predictable values. If an incoming packet does not have the expected IP TTL value it is possible that it is coming from an unauthorized and potentially harmful source.

TTL security is enabled using the **ttl-security** command. This command requires a minimum TTL value to be specified. When TTL security is enabled on a BGP session the IP TTL values in packets that are supposedly coming from the peer are compared (in hardware) to the configured minimum value and if there is a discrepancy the packet is discarded and a log is generated. TTL security is used most often on single-hop EBGP sessions but it can be used on multihop EBGP and IBGP sessions as well.

To enable TTL security on a single-hop EBGP session, configure **ttl-security** and **multihop** to a value of 255. To enable TTL security on a multihop EBGP session, configure **ttl-security** and **multihop** to match the expected TTL of (255 - hop count). The TTL value for both EBGP peers must be manually configured to the same value, as there is no TTL negotiation.



**Note:** IP packets sent to an IBGP peer are originated with an IP TTL value of 64. IP packets to an EBGP peer are originated with an IP TTL value of 1, except if **multihop** is configured; in that case, the TTL value is taken from the **multihop** command.

#### 5.2.5 BGP Address Family Support for Different Session Types

When the base router has a neighbor identified by an IPv4 address, and therefore the transport of the BGP session uses IPv4 TCP, all MP-BGP address families available in SR OS are supported by that session.

---

When the base router has a neighbor identified by an IPv6 address, and therefore the transport of the BGP session uses IPv6 TCP, the following MP-BGP address families are supported:

- ipv4
- ipv6
- mcast-ipv4
- mcast-ipv6
- vpn-ipv4
- vpn-ipv6
- evpn
- flow-ipv6
- label-ipv4
- label-ipv6
- bgp-ls

When a VPRN has a neighbor identified by an IPv4 address, and therefore the transport is IPv4 TCP, the following MP-BGP address families are supported:

- ipv4
- ipv6
- mcast-ipv4
- mcast-ipv6
- flow-ipv4
- flow-ipv6
- label-ipv4

When a VPRN has a neighbor identified by an IPv6 address, and therefore the transport is IPv6 TCP, the following MP-BGP address families are supported:

- ipv4
- ipv6
- mcast-ipv4
- mcast-ipv6
- flow-ipv6

---

## 5.2.6 BGP Groups

In SR OS, every neighbor (and hence BGP session) is configured under a **group**. A group is a CLI construct that saves configuration effort when multiple peers have a similar configuration; in this situation the common configuration commands can be configured once at the group level and need not be repeated for every neighbor. A single BGP instance can support many groups and each group can support many peers. Most SR OS commands that are available at the **neighbor** level are also available at the **group** level.

## 5.3 BGP Design Concepts


BGP assumes that all routers within an Autonomous System can reach destinations external to the Autonomous System using efficient, loop-free intra-AS forwarding paths. This generally requires that all the routers within the AS have a consistent view of the best path to every external destination. This is especially true when each BGP router in the AS makes its own forwarding decisions based on its own BGP routing table. The basic BGP specification does not store any intra-AS path information in the AS Path attribute so basic BGP has no way to detect routing loops within an AS that arise from inconsistent best path selections.

There are 3 solutions for dealing the issues outlined above.

- Create a full-mesh of IBGP sessions within the AS as shown in [Figure 31](#). This ensures routing consistency but does not scale well because the number of sessions increases exponentially with the number of BGP routers in the AS.
- Use BGP route reflectors in the AS. Route reflection is described in the section titled [Route Reflection](#). BGP route reflectors allow for routing consistency with only a partial mesh of IBGP sessions within the AS.

Create a confederation of autonomous systems. BGP confederations are described in the section titled [BGP Confederations](#).

**Figure 31 Fully Meshed BGP Configuration**



al\_0138


### 5.3.1 Route Reflection

In a standard BGP configuration a BGP route learned from one IBGP peer is not re-advertised to another IBGP peer. This rule exists because of the assumption of a full IBGP mesh within the AS. As discussed in the previous section a full IBGP mesh imposes certain scaling challenges. BGP route reflection eliminates the need for a full IBGP mesh by allowing routers configured as *route reflectors* to re-advertise routes from one IBGP peer to another IBGP peer.

A route reflector provides route reflection service to IBGP peers called *clients*. Other IBGP peers of the RR are called *non-clients*. An RR and its *client* peers form a *cluster*. A large AS can be sub-divided into multiple clusters, each identified by a unique 32-bit *cluster ID*. Each cluster contains at least one route reflector which is responsible for redistributing routes to its clients. The *clients* within a cluster do not need to maintain a full IBGP mesh between each other; they only require IBGP sessions to the route reflector(s) in their cluster. (If the clients within a cluster are fully meshed consider using the **disable-client-reflect** functionality.) The *non-clients* in an AS must be fully meshed with each other.

Figure 32 depicts the same network as Figure 31 but with route reflectors deployed to eliminate the IBGP mesh between SR-B, SR-C, and SR-D. SR-A, configured as the route reflector, is responsible for reflection routes to its clients SR-B, SR-C, and SR-D. SR-E and SR-F are non-clients of the route reflector. As a result, a full mesh of IBGP sessions must be maintained between SR-A, SR-E and SR-F.

**Figure 32 BGP Configuration with Route Reflectors**



A router becomes a route reflector whenever it has one or more client IBGP sessions. A client IBGP session is created with the **cluster** command, which also indicates the cluster ID of the client. Typical practice is to use the router ID as the cluster ID, but this is not necessary.

Basic route reflection operation (without Add-Path configured) can be summarized as follows:

- If the best and valid path for an NLRI is learned from a *client* and **disable-client-reflect** is not configured then advertise that route to all *clients*, *non-clients* and EBGp peers (as allowed by policy). If the client that advertised the best and valid path is a neighbor to which the **split-horizon** command (at the **bgp**, **group** or **neighbor** level) applies then the route is not advertised back to the sending client. In the route that is reflected to *clients* and *non-clients*:
  - The route reflector adds an ORIGINATOR\_ID attribute if it did not already exist; the ORIGINATOR\_ID indicates the BGP identifier (router ID) of the *client* that originated the route.
  - The route reflector prepends the cluster ID of the *client* that advertised the route and then the cluster ID of the *client* receiving the route (if applicable) to the CLUSTER\_LIST attribute, creating the attribute if it did not previously exist.
- If the best and valid path for an NLRI is learned from a *client* and **disable-client-reflect** is configured then advertise that route to all *clients* in other clusters, *non-clients* and EBGp peers (as allowed by policy). In the route that is reflected to *clients* in other clusters and *non-clients*:
  - The route reflector adds an ORIGINATOR\_ID attribute if it did not already exist; the ORIGINATOR\_ID indicates the BGP identifier (router ID) of the *client* that originated the route.
  - The route reflector prepends the cluster ID of the *client* that advertised the route and then the cluster ID of the *client* receiving the route (if applicable) to the CLUSTER\_LIST attribute, creating the attribute if it did not previously exist.
- If the best and valid path for an NLRI is learned from a *non-client* then advertise that route to all *clients* and EBGp peers (as allowed by policy). In the route that is reflected to *clients*:
  - The route reflector adds an ORIGINATOR\_ID attribute if it did not already exist; the ORIGINATOR\_ID indicates the BGP identifier (router ID) of the *non-client* that originated the route.
  - The route reflector prepends the cluster ID of the *client* receiving the route to the CLUSTER\_LIST attribute, creating the attribute if it did not previously exist.



- If the best and valid path for an NLRI is learned from an EBGP peer then advertise that route to all *clients*, *non-clients* and other EBGP peers (as allowed by policy). The ORIGINATOR\_ID and CLUSTER\_LIST attributes are not added to the route.
- If the best and valid path for an NLRI is locally originated (by the RR) — i.e. it was learned through means other than BGP — then advertise that route to all *clients*, *non-clients* and EBGP peers (as allowed by policy). The ORIGINATOR\_ID and CLUSTER\_LIST attributes are not added to the route.

The ORIGINATOR\_ID and CLUSTER\_LIST attributes allow BGP to detect the looping of a route within the AS. If any router receives a BGP route with an ORIGINATOR\_ID attribute containing its own BGP identifier the route is considered *invalid*. In addition if a route reflector receives a BGP route with a CLUSTER\_LIST attribute containing a locally configured cluster ID the route is considered *invalid*. Invalid routes are not installed in the route table and not advertised to other BGP peers.

### 5.3.2 BGP Confederations

BGP confederations are another alternative for avoiding a full mesh of BGP sessions inside an Autonomous System. A BGP confederation is a group of Autonomous Systems managed by a single technical administration that appear as a single AS to BGP routers outside the confederation; the single externally visible AS is called the confederation ID. Each AS in the group is called a *member AS* and the ASN of each member AS is visible only within the confederation. For this reason member ASNs are often private ASNs.

Within a confederation EBGP-type sessions can be setup between BGP routers in different member AS. These confederation-EBGP sessions avoid the need for a full mesh between routers in different member ASes. Within each member AS the BGP routers must be fully-meshed with IBGP sessions or route reflectors must be used to ensure routing consistency.

In SR OS, a confederation EBGP session is formed when the ASN of the peer is different from the local ASN and the peer ASN appears as a member AS in the **confederation** command. The confederation command specifies the confederation ID and up to 15 member AS that are part of the confederation.

When a route is advertised to a confederation-EBGP peer the advertising router prepends its local ASN, which is its member ASN, to a confederation-specific sub-element in the AS\_PATH that is created if it does not already exist. The extensions to the AS\_PATH are used for loop detection but they do not influence best path selection (i.e. they do not increase the AS Path length used in the BGP decision process). The MED, NEXT\_HOP and LOCAL\_PREF attributes in the received route are propagated unchanged by default. The ORIGINATOR\_ID and CLUSTER\_LIST attributes are not included in routes to confed-EBGP peers.

When a route is advertised to an EBGP peer outside the confederation the advertising router removes all member AS elements from the AS\_PATH and prepends its confederation ID rather than its local/member ASN.

---

## 5.4 BGP Messages

BGP protocol operation relies on the exchange of BGP messages between peers. 7450, 7750, 7950 routers, and most other routers, support the following five message types: Open, Update, Notification, Keepalive and Route Refresh. Details about each one are described in the following sections.

The minimum length of a BGP message is 19 bytes and the maximum length is 4096 bytes. BGP messages appear as a stream of bytes to the underlying TCP transport layer so there is no direct association between a BGP message and a TCP segment. One TCP segment can carry parts of one or more BGP messages. The maximum size of a BGP TCP segment sent by a 7450, 7750, or 7950 router is 1024 bytes (assuming a 40 byte TCP/IP header) if path MTU discovery is not enabled for the BGP session and the interfaces have default **tcp-mss** configurations. When path MTU discovery is enabled (with the **path-mtu-discovery** command) the maximum TCP segment size is discovered from received ICMP messages.

### 5.4.1 Open Message

After a TCP connection is established between two BGP routers the first message sent by each one is an Open message. If the received Open message is acceptable a Keepalive message confirming the Open is sent back. (See [BGP Session States](#) for more details.) An Open message contains the following information:

- Version — The current BGP version number is 4.
- Autonomous System number — The 2-byte AS of the sending router. If the sending router has an ASN greater than 65535, this field has the special value 23456 (AS\_TRANS). On a 7450, 7750, or 7950 router, the ASN in the Open message is based on the confederation ID (if the peer is external to the confederation), the global AS (configured using the **autonomous-system** command) or a session-level override of the global AS called the local AS (configured using the **local-as** command). More details about the use of local-AS are described in the section titled [Using Local AS for ASN Migration](#). More details about 4-byte AS numbers are described in the section titled [4-Octet Autonomous System Numbers](#).
- Hold Time — The proposed maximum time BGP will wait between successive messages (Keepalive and/or Update) from its peer before closing the connection. The actual hold time is the minimum of the configured **hold-time** for the session and the hold-time in the peer's Open message. If this minimum is below a configured threshold (**min** hold-time), the connection attempt is rejected.



**Note:** Changes to the configured **hold-time** trigger a session reset.

- BGP Identifier — The router ID of the BGP speaker. In Open messages, the BGP Identifier comes from the **router-id** configured under **bgp**; if that is not configured, then the **router-id** configured under **config>router** (or **config>service>vprn**) is used and if that too is not configured then the system interface IPv4 address is used.



**Note:** A change of the router ID in the **config>router>bgp** context causes all BGP sessions to be reset immediately while other changes resulting in a new BGP identifier only take effect after BGP is shutdown and re-enabled.

- Optional Parameters — A list of optional parameters, each encoded as a TLV. The only optional parameter that has been defined is the optional parameter. The optional parameter supports the process of BGP advertisement (see [5.4.1.3](#) for more information). When a BGP router receives an Open message with an unsupported optional parameter type it terminates the session. Unless **disable-capability-negotiation** is configured, the router always sends an optional parameter in its Open message.

### 5.4.1.1 Changing the Autonomous System Number

If the AS number is changed at the router level (**config>router**) the new AS number is not used until the BGP instance is restarted either by administratively disabling and enabling the BGP instance or by rebooting the system with the new configuration.

On the other hand if the AS number is changed in the BGP configuration (**config>router>bgp**) the effects are as follows:

- A change of the local-AS at the global level causes the BGP instance to restart with the new local AS number.
- A change of the local-AS at the **group** level causes BGP to re-establish sessions with all peers in the group using the new local AS number.
- A change of the local-AS at the **neighbor** level causes BGP to re-establish the session with the new local AS number.

---

### 5.4.1.2 Changing a Confederation Number

Changing the a confederation value on an active BGP instance will not restart the protocol. The change will take affect when the BGP protocol is (re) initialized.

### 5.4.1.3 BGP Advertisement

BGP advertisement allows a BGP router to indicate to a peer, using the optional parameter, the features that it supports so that they can coordinate and use only the features that both support. Each capability in the optional parameter is TLV-encoded with a unique type code. SR OS supports the following capability codes:

- Multi-protocol BGP (code 1)
- Route refresh (code 2)
- Outbound route filtering (code 3)
- Graceful restart (code 64)
- 4-octet AS number (code 65)
- Add-path (code 69)

## 5.4.2 Update Message

Update messages are used to advertise and withdraw routes. An Update message provides the following information:

- Withdrawn routes length — The length of the withdrawn routes field that is described next (may be 0).
- Withdrawn routes — IPv4 prefixes that are no longer considered reachable by the advertising router.
- Total path attribute length — The length of the path attributes field that is discussed next (may be 0).
- Path attributes — The path attributes presented in variable length TLV format. The path attributes apply to all the NLRI in the UPDATE message.
- Network layer reachability information (NLRI) — IPv4 prefixes that are considered reachable by the advertising router.

For fast routing convergence, as many NLRI as possible are packed into a single Update message as possible. This requires identifying all the routes that share the same path attribute values.

---

### 5.4.3 Keepalive Message

After a session is established, each router sends periodic Keepalive messages to its peer to test that the peer is still alive and reachable. If no Keepalive or Update message is received from the peer for the negotiated hold-time duration, the session is terminated. The period between one Keepalive message and the next is 1/3 of the negotiated hold-time duration or the value configured with the **keepalive** command, whichever is less. If the active hold-time or keepalive interval is zero, Keepalive messages are not sent. The default hold-time is 90 seconds and the default keepalive interval is 30 seconds.

Many times a peer (reachability) failure is detected through faster mechanisms than hold-timer expiry, as explained in the section titled [Detecting BGP Session Failures](#).

### 5.4.4 Notification Message

When a non-recoverable error related to a particular session occurs a Notification message is sent to the peer and the session is terminated (or restarted if GR is enabled for this scenario; see the section titled [BGP Graceful Restart](#) for more details). The Notification message provides the following information:

- Error code — Indicates the type of error: message header error, Open message error, Update message error, Hold timer expired, Finite State Machine error, or Cease.
- Error subcode — Provides more specific information about the error. The meaning of the subcode is specific to the error code.

#### 5.4.4.1 UPDATE Message Error Handling

The approach to handling Update message errors has evolved in the past couple of years. The original BGP protocol specification called for all UPDATE message errors to be handled the same way — send a NOTIFICATION to the peer and immediately close the BGP session. This error handling approach was motivated by the goal to ensure protocol “correctness” above all else. But it ignored several important points:

- Not all UPDATE message errors truly have the same severity. If the NLRI cannot be extracted and parsed from an UPDATE message then this is indeed a “critical” error. But other errors such as incorrect attribute flag settings, missing mandatory path attributes, incorrect next-hop length/format, etc. can be considered “non-critical” and handled differently.

- Session resets are extremely costly in terms of their impact on the stability and performance of the network. For many types of UPDATE message errors a session reset does not solve the problem because the root cause remains (e.g. software error, hardware error or misconfiguration). If a session reset is absolutely necessary then the operator should have some control over the timing.
- Some degree of protocol “incorrectness” is tolerable for a short period of time as long as the network operator is fully aware of the issue. In this context “incorrectness” typically means a BGP RIB inconsistency between routers in the same AS. Such inconsistency has become less and less of an issue over time as edge-to-edge tunneling of IP traffic (e.g. BGP shortcuts, IP VPN) has reduced the number of deployments where IP traffic is forwarded hop-by-hop.

In recognition of these points and the general trend towards more flexibility in BGP error handling, SR OS supports a BGP configuration option called **update-fault-tolerance** that allows the operator to decide whether the router should apply new or legacy error handling procedures to UPDATE message errors. If **update-fault-tolerance** is configured, then non-critical errors as described above are handled using the “treat-as-withdraw” or “attribute-discard” approaches to error handling; these approaches do not cause a session reset. If **update-fault-tolerance** is not configured then legacy procedures continue to apply and all errors (critical and non-critical) trigger a session reset.

If the update-fault-tolerance command was previously configured and a non-critical error was already triggered, the BGP session will still be reset when the operator configures no update-fault-tolerance.

## 5.4.5 Route Refresh Message

A BGP router can send a Route Refresh message to its peer only if both have advertised the route refresh capability (code 2). The Route Refresh message is a request for the peer to re-send all or some of its routes associated with a particular pair of AFI/SAFI values. AFI/SAFI values are the same ones used in the MP-BGP capability (see the section titled [Multi-Protocol BGP Attributes](#)).

7450, 7750, and 7950 routers only send Route Refresh messages for AFI/SAFI associated with VPN routes that carry Route Target extended communities, such as VPN-IPv4, VPN-IPv6, L2-VPN, MVPN-IPv4 and MVPN-IPv6 routes. By default routes of these types are discarded if, at the time they are received, there is no VPN that imports any of the route targets they carry. If at a later time a VPN is added or reconfigured (in terms of the route targets that it imports) a Route Refresh message is sent to all relevant peers so that previously discarded routes can be relearned.



**Note:** Route Refresh messages are not sent for VPN-IPv4 and VPN-IPv6 routes if **mp-bgp-keep** is configured; in this situation received VPN-IP routes are kept in the RIB-IN regardless of whether or not they match a VRF import policy.



## 5.5 BGP Path Attributes

Path attributes are fundamental to BGP. A BGP route for a particular NLRI is distinguished from other BGP routes for the same NLRI by its set of path attributes. Each path attribute describes some property of the path and is encoded as a TLV in the Path Attributes field of the Update message. The type field of the TLV identifies the path attribute and the value field carries data specific to the attribute type. There are 4 different categories of path attributes:

- **Well-known mandatory.** These attributes must be recognized by all BGP routers and must be present in every Update message that advertises reachable NLRI towards a certain type of neighbor (EBGP or IBGP).
- **Well-known discretionary.** These attributes must be recognized by all BGP routers but are not required in every Update message.
- **Optional transitive.** These attributes are allowed to be unrecognized by some BGP routers. If a BGP router does not recognize one of these attributes it accepts it, passes it on to other BGP peers, and sets the Partial bit to 1 in the attribute flags byte.
- **Optional non-transitive.** These attributes are allowed to be unrecognized by some BGP routers. If a BGP router does not recognize one of these attributes it is quietly ignored and not passed on to other BGP peers.

SR OS supports the following path attributes, which are described in detail in upcoming sections:

- ORIGIN (well-known mandatory)
- AS\_PATH (well-known mandatory)
- NEXT\_HOP (well-known, required only in Update messages with IPv4 prefixes in the NLRI field)
- MED (optional non-transitive)
- LOCAL\_PREF (well-known, required only in Update messages sent to IBGP peers)
- ATOMIC\_AGGR (well-known discretionary)
- AGGREGATOR (optional transitive)
- COMMUNITY (optional transitive)
- ORIGINATOR\_ID (optional non-transitive)
- CLUSTER\_LIST (optional non-transitive)
- MP\_REACH\_NLRI (optional non-transitive)
- MP\_UNREACH\_NLRI (optional non-transitive)
- EXT\_COMMUNITY (optional transitive)

- AS4\_PATH (optional transitive)
- AS4\_AGGREGATOR (optional transitive)
- CONNECTOR (optional transitive)
- PMSI\_TUNNEL (optional transitive)
- TUNNEL\_ENCAPSULATION (optional transitive)
- AIGP (optional non-transitive)
- BGP-LS (optional non-transitive)
- LARGE\_COMMUNITY (optional transitive)

### 5.5.1 Origin

The ORIGIN path attribute indicates the origin of the path information. There are three supported values:

- IGP (0)
- EGP (1)
- Incomplete (2)

When a router originates a VPN-IP prefix (from a non-BGP route), it sets the value of the Origin attribute to IGP. When a router originates a BGP route for an IP prefix by exporting a non-BGP route from the routing table, it sets the value of the Origin attribute to Incomplete. Route policies (BGP import and export) can be used to change the Origin value.

### 5.5.2 AS Path

The AS\_PATH attribute provides the list of Autonomous Systems through which the routing information has passed. The AS\_PATH attribute is composed of segments. There can be up to 4 different types of segments in an AS\_PATH attribute: AS\_SET, AS\_SEQUENCE, AS\_CONFED\_SET and AS\_CONFED\_SEQUENCE. The AS\_SET and AS\_CONFED\_SET segment types result from route aggregation. AS\_CONFED\_SEQUENCE contains an ordered list of member AS through which the route has passed inside a confederation. AS\_SEQUENCE contains an ordered list of AS (including confederation IDs) through which the route has passed on its way to the local AS/confederation.

The AS numbers in the AS\_PATH attribute are all 2-byte values or all 4-byte values (if the 4-octet ASN capability was announced by both peers).

A BGP router always prepends its AS number to the AS\_PATH attribute when advertising a route to an EBGp peer. The specific details for a 7450, 7750, or 7950 router are described below.

- When a route is advertised to an EBGp peer and the advertising router is not part of a confederation:
  - The global AS (configured using the **autonomous-system** command) is prepended to the AS\_PATH if **local-as** is not configured
  - The local AS followed by the global AS are prepended to the AS\_PATH if **local-as** is configured.
  - Only the local AS is prepended to the AS\_PATH if **local-as no-prepend-global-as** is configured
  - Private AS numbers (64512 - 65534 inclusive) are removed from the AS\_PATH if **remove-private** is configured.
- When a route is advertised to an EBGp peer outside a confederation:
  - The confederation ID is prepended to the AS\_PATH if **local-as** is not configured
  - The local AS followed by the confederation ID are prepended to the AS\_PATH if **local-as** is configured. (The **no-prepend-global-as** option has no effect in this scenario.)
  - Member AS numbers are removed from the AS\_PATH as described in the section titled [BGP Confederations](#).
  - Private AS numbers (64512 - 65534 inclusive) are removed from the AS\_PATH if **remove-private** is configured.
- When a route is advertised to a confederation-EBGP peer:
  - If the route came from an EBGp peer and **local-as** was configured on this session (*without* the **private** option) this local AS number is prepended to the AS\_PATH in a regular AS\_SEQUENCE segment
  - The global AS (configured using the **autonomous-system** command) is prepended, as a member AS, to the AS\_PATH if **local-as** is not configured
  - The local AS followed by the global AS are prepended, as member AS, to the AS\_PATH if **local-as** is configured
  - Only the local AS is prepended, as a member AS, to the AS\_PATH if **local-as no-prepend-global-as** is configured
  - Private AS numbers (64512 - 65534 inclusive) are removed from the AS\_PATH if **remove-private** is configured (except for the local AS added as a member AS).
- When a route is advertised to an IBGP peer:
  - No information is added to the AS\_PATH if the route is locally originated or if it came from an IBGP peer.

- The local AS number is prepended to the AS\_PATH if the route came from an EBGP peer and **local-as** is configured *without* the **private** option.
- The local AS number is prepended, as a member AS, to the AS\_PATH if the route came from a confederation-EBGP peer and **local-as** is configured *without* the **private** option.
- Private AS numbers (64512 - 65534 inclusive) are removed from the AS\_PATH if **remove-private** is configured.

BGP import policies can be used to prepend an AS number multiple times to the AS\_PATH, whether the route is received from an IBGP, EBGP or confederation EBGP peer. The AS path prepend action is also supported in BGP export policies applied to these types of peers, regardless of whether the route is locally originated or not. AS path prepending in export policies occurs before the global and/or local ASes (if applicable) are added to the AS\_PATH.

When a BGP router receives a route containing one of its own Autonomous System numbers (local or global or confederation ID) in the AS\_PATH the route is normally considered *invalid* for reason of an AS path loop. However SR OS provides a **loop-detect** command that allows this check to be bypassed. If it known that advertising certain routes to an EBGP peer will result in an AS path loop condition and yet there is no loop (assured by other mechanisms, such as the Site of Origin (SOO) extended community) then **as-override** can be configured on the advertising router instead of disabling loop detection on the receiving router. The **as-override** command replaces all occurrences of the peer AS in the AS\_PATH with the advertising router's local AS.

### 5.5.2.1 AS Override

The AS Override feature can be used in VPRN scenarios where a customer is running BGP as the PE-CE protocol and some or all of the CE locations are in the same Autonomous System (AS). With normal BGP, two sites in the same AS would not be able to reach each other directly since there is an apparent loop in the AS Path.

When **as-override** is configured on a PE-CE EBGP session the PE rewrites the customer ASN in the AS Path with the VPRN AS number as the route is advertised to the CE.

---

### 5.5.2.2 Using Local AS for ASN Migration

The description in the previous section does fully explain the reasons for using **local-as**. This BGP feature facilitates the process of changing the ASN of all the routers in a network from one number to another. This may be necessary if one network operator merges with or acquires another network operator and the two BGP networks must be consolidated into one Autonomous System.

For example suppose the operator of the ASN 64500 network merges with the operator of the ASN 64501 network and the new merged entity decides to renumber ASN 64501 routers as ASN 64500 routers so that they the entire network can be managed as one Autonomous System. The migration can be carried out using the following sequence of steps:

1. Change the global AS of the route reflectors that used to be part of ASN 64501 to the new value 64500.
2. Change the global AS of the RR clients that used to be part of ASN 64501 to the new value 64500.
3. Configure **local-as 64501 private no-prepend-global-as** on every EBGP session of each RR client migrated in step 2.

This migration procedure has several advantages. First, customers, settlement-free peers and transit providers of the previous ASN 64501 network still perceive that they are peering with ASN 64501 and can delay switching to ASN 64500 until the time is convenient for them. Second, the AS path lengths of the routes exchanged with the EBGP peers are unchanged from before so that best path selections are preserved.

### 5.5.2.3 4-Octet Autonomous System Numbers

When BGP was developed, it was assumed that 16-bit (2-octet) ASNs would be sufficient for global Internet routing. In theory a 16-bit ASN allows for 65536 unique autonomous systems but some of the values are reserved (0 and 64000-65535). Of the assignable space less than 10% remains available. When a new AS number is needed it is now simpler to obtain a 4-octet AS number. 4-octet AS numbers have been available since 2006. A 32-bit (4-octet) ASN allows for 4,294,967,296 unique values (some of which are again, reserved).

When 4-octet AS numbers became available it was recognized that not all routers would immediately support the ability to parse 4-octet AS numbers in BGP messages so two optional transitive attributes called AS4\_PATH and AS4\_AGGREGATOR were introduced to allow a gradual migration.

A BGP router that supports 4-octet AS numbers advertises this capability in its OPEN message; the capability information includes the AS number of the sending BGP router, encoded using 4 bytes (recall the ASN field in the OPEN message is limited to 2 bytes). By default, OPEN messages sent by 7450, 7750, or 7950 routers always include the 4-octet ASN capability but this can be changed using the **disable-4byte-asn** command.

If a BGP router and its peer have both announced the 4-octet ASN capability then the AS numbers in the AS\_PATH and AGGREGATOR attributes are always encoded as 4-byte values in the UPDATE messages they send to each other. These UPDATE messages should not contain the AS4\_PATH and AS4\_AGGREGATOR path attributes.

If one of the routers involved in a session announces the 4-octet ASN capability and the other one does not then the AS numbers in the AS\_PATH and AGGREGATOR attributes are encoded as 2-byte values in the UPDATE messages they send to each other.

When a 7450, 7750, or 7950 router advertises a route to a peer that did not announce the 4-octet ASN capability:

- If there are any AS numbers in the AS\_PATH attribute that cannot be represented using 2 bytes (because they have a value greater than 65535) they are substituted with the special value 23456 (AS\_TRANS) and an AS4\_PATH attribute is added to the route if it is not already present. The AS4\_PATH attribute has the same encoding as the AS\_PATH attribute that would be sent to a 4-octet ASN capable router (i.e. each AS number is encoded using 4 octets) but it does not carry segments of type AS\_CONFED\_SEQUENCE or AS\_CONFED\_SET.
- If the AS number in the AGGREGATOR attribute cannot be represented using 2 bytes (because its value is greater than 65535) it is substituted with the special value 23456 and an AS4\_AGGREGATOR attribute is added to the route if it is not already present. The AS4\_AGGREGATOR is the same as the AGGREGATOR attribute that would be sent to a 4-octet ASN capable router (i.e. the AS number is encoded using 4 octets).

When a 7450, 7750, or 7950 router receives a route with an AS4\_PATH attribute it attempts to reconstruct the full AS path from the AS4\_PATH and AS\_PATH attributes, regardless of whether **disable-4byte-asn** is configured or not. The reconstructed path is the AS path displayed in BGP show commands. If the length of the received AS4\_PATH is N and the length of the received AS\_PATH is N+t then the reconstructed AS path contains the t leading elements of the AS\_PATH followed by all the elements in the AS4\_PATH.

## 5.5.3 Next-Hop

The NEXT\_HOP attribute indicates the IPv4 address of the BGP router that is the next-hop to reach the IPv4 prefixes in the NLRI field. If the Update message is advertising routes other than IPv4 unicast routes the next-hop of these routes is encoded in the MP\_REACH\_NLRI attribute; see the section titled [Multi-Protocol BGP Attributes](#) for more details.

The rules for deciding what next-hop address types to accept in a received BGP route and what next-hop address types to advertise as a BGP next-hop are address family dependent. The following sections summarize the key details.

### 5.5.3.1 Unlabeled IPv4 Unicast Routes

SR OS routers do not send or receive IPv4 routes with IPv6 next-hop addresses and do not signal such a capability even when configured for **extended-nh-encoding**.

When an IPv4 BGP route is advertised to an EBGp peer **next-hop-self** always applies (except if the **third-party-nexthop** command is applied, as described in the following note). Next-hop-self results in one of the following outcomes:

- If the EBGp session uses IPv4 transport, the BGP next-hop is the local-address used to setup the session and this cannot be overridden, even by BGP export policy.
- If the EBGp session uses IPv6 transport, the BGP next-hop is set to the IPv4 address of the system interface. (If the system interface does not have an IPv4 address the route is not advertised unless an export policy sets a valid IPv4 next-hop.)



**Note:** Configuring **third-party-nexthop** allows an IPv4 route received from one EBGp peer to be advertised to another EBGp that is in the same IP subnet with an unchanged BGP next-hop.

When an IPv4 BGP route is re-advertised to an IBGP or confederation-EBGP peer, the advertising router does not modify the BGP next-hop by default; however, this can be changed as follows:

- If the BGP **next-hop-self** command is applied to the IBGP or confederation-EBGP peer, then this changes the BGP next-hop to the local-address used to setup the session (if the transport to the peer is IPv4) or to the IPv4 address of the system interface (if the transport to the peer is IPv6). This command applies to all routes advertised to the peer, regardless of the peer type from which they were received (IBGP, confed-EBGP or EBGP).
- If IPv4 routes are matched and accepted by an export policy applied to an IBGP or confederation-EBGP session, and the matching policy entry has a next-hop-self action, this changes the BGP next-hop of only the matched routes to the local-address used to setup the session (if the transport to the peer is IPv4) or to the IPv4 address of the system interface (if the transport to the peer is IPv6).
- If IPv4 routes are matched and accepted by an export policy applied to an IBGP or confederation-EBGP session, and the matching policy entry has a next-hop <ip-address> action, this changes the BGP next-hop of only the matched routes to <ip-address>, but only if <ip-address> is an IPv4 address. If <ip-address> is an IPv6 address the matched routes are treated as though they were rejected by the policy entry.

When an IPv4 BGP route is locally originated and advertised to an IBGP or confederation-EBGP peer, the BGP next-hop is, by default, copied from the next-hop of the route that was imported into BGP, with certain exceptions (for example, black-hole next-hop). When a static route with indirect next-hop is re-advertised as a BGP route, the BGP next-hop is a copy of the indirect address, however with route-table-import policies BGP can be instructed to take the resolved next-hop of the static route as the BGP next-hop address.

### 5.5.3.2 Unlabeled IPv6 Unicast Routes

SR OS routers never send or receive IPv6 routes with 32-bit IPv4 next-hop addresses.

When an IPv6 BGP route is advertised to an EBGP peer, next-hop-self always applies (except if the third-party-nexthop command is applied, as described in the following note). Next-hop-self results in one of the following outcomes:

- If the EBGP session uses IPv4 transport, the BGP next-hop encodes the local-address used for setup of the session as an IPv4-compatible IPv6 address (all zeros in the first 96 bits followed by the 32 bit IPv4 local-address).
- If the EBGP session uses IPv6 transport, the BGP next-hop is the local-address used to setup the session and this cannot be overridden, even by BGP export policy.





**Note:** Configuring **third-party-nextthop** allows an IPv6 route received from one EBGP peer to be advertised to another EBGP that is in the same IP subnet with an unchanged BGP next-hop.

When an IPv6 BGP route is re-advertised to an IBGP or confederation-EBGP peer, the advertising router does not modify the BGP next-hop by default; however, this can be changed as follows:

- If the BGP **next-hop-self** command is applied to the IBGP peer or confederation-EBGP peer, then this changes the BGP next-hop to the local-address used to setup the session (if the transport to the peer is IPv6) or to an IPv4-compatible IPv6 address derived from the IPv4 local-address used to setup the session (if the transport to the peer is IPv4). This command applies to all routes advertised to the peer, regardless of the peer type from which they were received (IBGP, confed-EBGP, or EBGP).
- If IPv6 routes are matched and accepted by an export policy applied to an IBGP or confederation-EBGP session, and the matching policy entry has a **next-hop-self** action, this changes the BGP next-hop of only the matched routes to the local-address used to setup the session (if the transport to the peer is IPv6) or to an IPv4-compatible IPv6 address derived from the IPv4 local-address used to setup the session (if the transport to the peer is IPv4).
- If IPv6 routes are matched and accepted by an export policy applied to an IBGP or confederation-EBGP session, and the matching policy entry has a **next-hop <ip-address>** action, this changes the BGP next-hop of only the matched routes to <ip-address>, but only if <ip-address> is an IPv6 address. If <ip-address> is an IPv4 address the matched routes are treated as though they were rejected by the policy entry.

When an IPv6 BGP route is locally originated and advertised to an IBGP or confederation-EBGP peer, the BGP next-hop is, by default, copied from the next-hop of the route that was imported into BGP, with certain exceptions (for example, black-hole next-hop). When a static route with indirect next-hop is re-advertised as a BGP route, the BGP next-hop is a copy of the indirect address, however with route-table-import policies BGP can be instructed to take the resolved next-hop of the static route as the BGP next-hop address.

### 5.5.3.3 VPN-IPv4 Routes

SR OS routers can always send and receive VPN-IPv4 routes with IPv4 next-hops. They can also be configured (using the **extended-nh-encoding** command) to receive VPN-IPv4 routes with IPv6 next-hop addresses from selected BGP peers by signaling the corresponding Extended NH Encoding BGP capability to those peers during session setup. If the capability is not advertised to a peer, then such routes will not be accepted from that peer. Also, if the SR OS router does not receive an Extended NH Encoding capability advertisement for [NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2] from a peer, then it will not advertise VPN-IPv4 routes with IPv6 next-hops to that peer.

When a VPN-IPv4 BGP route is advertised to an EBGp peer, next-hop-self applies if **enable-inter-as-vpn** is configured; otherwise there is no change to the next-hop. Next-hop-self results in one of the following outcomes:

- If the EBGp session uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.
- If the EBGp peer opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops vpn-ipv4** command, then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface.

If **enable-inter-as-vpn** is configured, a VPN-IPv4 BGP route is received from an EBGp peer and re-advertised to an IBGP or confederation-EBGP peer, next-hop-self applies. (If **enable-inter-as-vpn** is not configured, the next-hop may be changed with the **next-hop-self** command, however, this is strongly discouraged because it can result in a change of the next-hop without a change in the VPN label.) Next-hop-self results in one of the following outcomes:

- If the IBGP or confederation-EBGP session uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.
- If the IBGP or confederation-EBGP peer opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops vpn-ipv4** command then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface.

When a VPN-IPv4 BGP route is reflected from one IBGP peer to another IBGP peer, the RR does not modify the next-hop by default. However, if the **next-hop-self** command is applied to the IBGP peer receiving the route and **enable-rr-vpn-forwarding** is configured then this combination of commands has one of the following outcomes:

- If the IBGP session receiving the reflected route uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.
- If the IBGP session receiving the reflected route opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops vpn-ipv4** command then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface.

When a VPN-IPv4 BGP route is reflected from one IBGP peer to another IBGP peer and **enable-rr-vpn-forwarding** is configured and the VPN-IPv4 route is matched and accepted by an export policy entry with a **next-hop** <ip-address> action, this changes the BGP next-hop of the matched routes to <ip-address>, except if <ip-address> is an IPv6 address and the receiving IBGP peer did not advertise an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2) or, in the configuration of the local router, the session is not associated with an **advertise-ipv6-next-hops vpn-ipv4** command. In this case, the route is treated as though it was rejected by the policy entry.

#### 5.5.3.4 VPN-IPv6 Routes

SR OS routers never send or receive VPN-IPv6 routes with 32-bit IPv4 next-hop addresses.

When a VPN-IPv6 BGP route is advertised to an EBGp peer, next-hop-self applies if **enable-inter-as-vpn** is configured. Otherwise, there is no change to the next-hop. Next-hop-self results in one of the following outcomes:

- If the EBGp session uses IPv4 transport, then the BGP next-hop is set to the IPv4 local-address used to setup the session but encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).

- If the EBGP session uses IPv6 transport and it is associated with an **advertise-ipv6-next-hops vpn-ipv6** command, then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).

If **enable-inter-as-vpn** is configured, a VPN-IPv6 BGP route is received from an EBGP peer and re-advertised to an IBGP or confederation-EBGP peer next-hop-self applies. (If **enable-inter-as-vpn** is not configured, the next-hop may be changed with the **next-hop-self** command, however, this is strongly discouraged because it can result in a change of the next-hop without a change in the VPN label.) Next-hop-self results in one of the following outcomes:

- If the IBGP or confederation-EBGP session uses IPv4 transport, then the BGP next-hop is set to the IPv4 local-address used to setup the session but encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).
- If the IBGP or confederation-EBGP session uses IPv6 transport and it is associated with an **advertise-ipv6-next-hops vpn-ipv6** command, then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).

When a VPN-IPv6 BGP route is reflected from one IBGP peer to another IBGP peer, the RR does not modify the next-hop by default. However, if the **next-hop-self** command is applied to the IBGP peer receiving the route and **enable-rr-vpn-forwarding** is configured, then this combination of commands has one of the following outcomes:

- If the IBGP session receiving the reflected route uses IPv4 transport then the BGP next-hop is set to the IPv4 local-address used to setup the session but encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).
- If the IBGP session receiving the reflected route uses IPv6 transport and it is associated with an **advertise-ipv6-next-hops vpn-ipv6** command, then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a ::FFFF/96 prefix).

When a VPN-IPv6 BGP route is reflected from one IBGP peer to another IBGP peer and **enable-rr-vpn-forwarding** is configured and the VPN-IPv6 route is matched and accepted by an export policy entry with a **next-hop** <ip-address> action, this changes the BGP next-hop of the matched routes to <ip-address> if it is specified as a 128-bit IPv6 address or to an IPv4-mapped IPv6 address encoding <ip-address> if it is specified as a 32-bit IPv4 address.

### 5.5.3.5 Label-IPv4 Routes

SR OS routers can always send and receive label-IPv4 routes with IPv4 next-hops. They can also be configured (using the **extended-nh-encoding** command) to receive label-IPv4 routes with IPv6 next-hop addresses from selected BGP peers by signaling the corresponding Extended NH Encoding BGP capability to those peers during session setup. If the capability is not advertised to a peer then such routes will not be accepted from that peer. Also, if the SR OS router does not receive an Extended NH Encoding capability advertisement for [NLRI AFI=1, NLRI SAFI=4, next-hop AFI=2] from a peer then it will not advertise label-IPv4 routes with IPv6 next-hops to that peer.

When a label-IPv4 BGP route is advertised to an EBGp peer, next-hop-self applies unless the EBGp session has **next-hop-unchanged** enabled for the label-ipv4 address family. Next-hop-self results in one of the following outcomes:

- If the EBGp session uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.
- If the EBGp peer opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=4, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops label-ipv4** command, then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface.

When a label-IPv4 BGP route is received from an EBGp peer and re-advertised to an IBGP or confederation-EBGP peer, next-hop-self applies unless the IBGP or confederation-EBGP session has **next-hop-unchanged** enabled for the label-ipv4 address family. Next-hop-self results in one of the following outcomes:

- If the IBGP or confederation-EBGP session uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.

- If the IBGP or confederation-EBGP peer opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=4, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops label-ipv4** command then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface

When a label-IPv4 BGP route is reflected from one IBGP peer to another IBGP peer, the RR does not modify the next-hop by default. However, if the **next-hop-self** command is applied to the IBGP peer receiving the route, then this results in one of the following outcomes:

- If the IBGP session receiving the reflected route uses IPv4 transport, then the BGP next-hop is taken from the value of the local-address used to setup the session.
- If the IBGP session receiving the reflected route opened an IPv6-transport session by advertising an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=4, next-hop AFI=2) AND, in the configuration of the local router, the session is associated with an **advertise-ipv6-next-hops label-ipv4** command then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface.

When a label-IPv4 BGP route is reflected from one IBGP peer to another IBGP peer and the label-IPv4 route is matched and accepted by an export policy entry with a **next-hop <ip-address>** action, this changes the BGP next-hop of the matched routes to <ip-address>, except if <ip-address> is an IPv6 address and the receiving IBGP peer did not advertise an extended NH encoding capability with (NLRI AFI=1, NLRI SAFI=128, next-hop AFI=2) or, in the configuration of the local router, the session is not associated with an **advertise-ipv6-next-hops label-ipv4** command. In this case, the route is treated as though it was rejected by the policy entry.

### 5.5.3.6 Label-IPv6 Routes

SR OS routers never send or receive label-IPv6 routes with 32-bit IPv4 next-hop addresses.

In the 16.0R1 software release, there is no support for sending and receiving label-IPv6 routes to or from EBGp peers, although SR OS does not block the family from being configured and advertised as a capability towards EBGp peers.

When a label-IPv6 BGP route is reflected from one IBGP peer to another IBGP peer, the RR does not modify the next-hop by default. However, if the **next-hop-self** command is applied to the IBGP peer receiving the route, then this results in one of the following outcomes:

- If the IBGP session receiving the reflected route uses IPv4 transport, then the BGP next-hop is set to the IPv4 local-address used to setup the session but encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a `::FFFF/96` prefix).
- If the IBGP session receiving the reflected route uses IPv6 transport and it is associated with an **advertise-ipv6-next-hops label-ipv6** command then the BGP next-hop is set to the value of the IPv6 local-address used for setup of the session. Otherwise, the BGP next-hop is set to the IPv4 address of the system interface encoded as an IPv4-mapped IPv6 address (for example, with the IPv4 address in the least significant 32 bits of a `::FFFF/96` prefix).

When a label-IPv6 BGP route is reflected from one IBGP peer to another IBGP peer and the label-IPv6 route is matched and accepted by an export policy entry with a **next-hop** <ip-address> action, this changes the BGP next-hop of the matched routes to <ip-address> if it is specified as a 128-bit IPv6 address or to an IPv4-mapped IPv6 address encoding <ip-address> if it is specified as a 32-bit IPv4 address.

### 5.5.3.7 Next-Hop Resolution

To use a BGP route for forwarding, a BGP router must know how to reach the BGP next-hop of the route. The process of determining the local interface or tunnel used to reach the BGP next-hop is called next-hop resolution. The BGP next-hop resolution process depends on the type of route (the AFI/SAFI) and various configuration settings.

- The following procedures apply to the next-hop resolution of IPv4 and IPv6 (unlabeled) unicast routes:
  - When an IPv6 route is received with an IPv4-mapped IPv6 address as next-hop, BGP next-hop resolution is based on the embedded IPv4 address and this IPv4 address is matched to IPv4 routes and/or IPv4 tunnels, depending on configuration. The IPv4-mapped IPv6 address is never interpreted as an IPv6 address to be compared to IPv6 routes or IPv6 tunnels.
  - BGP routes are eligible to resolve a BGP next-hop only if the **use-bgp-routes** command is configured. A maximum of 4 levels of recursion are supported.
  - If there are multiple eligible routes that match the BGP next-hop, the longest prefix match (LPM) route is selected.

- 
- If the LPM route is rejected by the user-configured next-hop-resolution policy or if there are no eligible matching routes, the BGP next-hop is unresolved and all the routes with that next-hop are considered invalid and not advertised to peers.
  - BGP shortcuts are described in Next-hop Resolution of BGP Routes Using Tunnels.
  - The following procedures apply to the next-hop resolution of VPN-IPv4 and VPN-IPv6 routes:
    - When an VPN-IPv4 or VPN-IPv6 route is received with an IPv4-mapped IPv6 address as next-hop, BGP next-hop resolution is based on the embedded IPv4 address and this IPv4 address is matched to IPv4 routes and/or IPv4 tunnels, depending on configuration. The IPv4-mapped IPv6 address is never interpreted as an IPv6 address to be compared to IPv6 routes or IPv6 tunnels.
    - If the VPN-IP route is imported into a VPRN, the next-hop is resolved to a tunnel based on the auto-bind-tunnel configuration of the importing VPRN. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information.
    - If the VPN-IP route is received from an IBGP or EBGP peer, and the router is a next-hop-self RR or a model-B ASBR, the order of resolution is as follows:
      - local (interface) route
      - longest prefix-match static route, excluding default static routes, but only if allow-static is configured
      - tunnel, based on the transport-tunnel resolution-filter options for family VPN. See [Next-Hop Resolution of BGP Labeled Routes to Tunnels](#) for more information
  - The following procedures apply to the next-hop resolution of label-unicast IPv4 and label-unicast IPv6 routes:
    - When a BGP-LU route is received with an IPv4-mapped IPv6 address as next-hop, BGP next-hop resolution is based on the embedded IPv4 address and this IPv4 address is matched to IPv4 routes and/or IPv4 tunnels, depending on configuration. The IPv4-mapped IPv6 address is never interpreted as an IPv6 address to be compared to IPv6 routes or IPv6 tunnels.
    - If the BGP-LU route is received by a control-plane-only RR with the **disable-route-table-install** and **rr-use-route-table** commands configured, the order of resolution is as follows:
      - local route



- longest prefix-match static route, excluding default static routes. Use this route to resolve the BGP next-hop if it is a static route with blackhole next-hop. For other types of static routes, use this route to resolve the next-hop only if **allow-static** is configured
- longest prefix-match IGP route
- If the BGP-LU route is received from an IBGP or EBGP peer and the router is not an RR with the **rr-use-route-table** command configured, the order of resolution is as follows:
  - local route
  - longest prefix-match static route, excluding default static routes. Use this route to resolve the BGP next-hop if it is a static route with blackhole next-hop. If it is another type of static route, use it to resolve the next-hop only if **allow-static** is configured
  - tunnel, based on the transport-tunnel resolution-filter options for family label-ipv4 or family label-ipv6, depending on the case. See [Next-Hop Resolution of BGP Labeled Routes to Tunnels](#) for more information

#### 5.5.3.7.1 Next-Hop Resolution of BGP Unlabeled IPv4 Unicast Routes to Tunnels

To enable the next-hop resolution of unlabeled IPv4 routes using tunnels in the tunnel-table of the router, it is necessary to add **family ipv4** under the **configure>router>bgp>next-hop-resolution>shortcut-tunnel** context. In the **family ipv4** context, there are commands to specify the resolution mode (**any**, **disabled** or **filter**) and the set of tunnel types that are eligible to be used if filter mode (**resolution filter**) is selected.

If the resolution mode is set to disabled, the next-hops of unlabeled IPv4 routes can only be resolved by route table lookup.

If there are multiple tunnels in tunnel-table that are allowed by **resolution any** or **resolution filter** mode and that can resolve the BGP next-hop, the selection of the resolving tunnel comes down to factors such as route color, admin-tag-policy, tunnel-table preference, and LDP FEC prefix length.

If **disallow-igp** is enabled, then there is no attempt to resolve the IPv4 BGP route using route table lookup if no resolving tunnel can be found in the tunnel-table.

The available tunneling options for IPv4 shortcuts are:

- **bgp** — This refers to IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32 IPv4 prefixes

- **ldp** — This refers to /32 and shorter length LDP FEC prefixes imported into the tunnel table. For resolution purposes, BGP selects the LDP FEC that is the longest-prefix-match (LPM) of the BGP next-hop address
- **rsvp** — This refers to RSVP tunnels in tunnel-table. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id
- **sr-isis** — This refers to segment routing tunnels (shortest path) to destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in tunnel-table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance
- **sr-ospf** — This refers to segment routing tunnels (shortest path) to destinations reachable by the OSPF protocol. This option allows BGP to use the segment routing tunnel in tunnel-table submitted by the lowest preference OSPF instance or (in case of a tie) the lowest numbered OSPF instance
- **sr-policy** — This refers to segment routing policies that are statically configured in the local router or learned via BGP routes (AFI 1/SAFI 73). For BGP to resolve the next-hop of an IPv4 route using an sr-policy the highest numbered color extended community attached to the IPv4 route must match the color of the SR policy and if the CO bits of this color extended community have the value 00 the BGP next-hop of the route must exactly match the endpoint of the SR policy.
- **sr-te** — This refers to traffic engineered (TE) segment routing tunnels. This option allows BGP to use the best metric SR-TE tunnel to the address of the BGP next-hop. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel-id

#### 5.5.3.7.2 Next-Hop Resolution of BGP Unlabeled IPv6 Unicast Routes to Tunnels

To enable the next-hop resolution of unlabeled IPv6 routes using tunnels in the tunnel-table of the router, it is necessary to add **family ipv6** under the **configure>router>bgp>next-hop-resolution>shortcut-tunnel** context. If the next-hop of the IPv6 BGP route contains an IPv4-mapped IPv6 address, the shortcut-tunnel configuration applies to the use of IPv4 tunnels and IPv4 routes that match the embedded IPv4 address in the BGP next-hop. If the BGP next-hop is any other IPv6 address the shortcut-tunnel configuration applies to the use of IPv6 tunnels and IPv6 routes that match the full address of the BGP next-hop.

The **family ipv6** context has commands to specify the resolution mode (**any**, **disabled** or **filter**) and the set of tunnel types that are eligible to be used if filter mode (**resolution filter**) is selected.

If there are multiple tunnels in tunnel-table that are allowed by **resolution any** or **resolution filter** mode and that can resolve the BGP next-hop, the selection of the resolving tunnel comes down to factors such as route color, admin-tag-policy, tunnel-table preference, and LDP FEC prefix length.

If the resolution mode is set to disabled, the next-hops of unlabeled IPv6 routes can only be resolved by route table lookup.

If there are multiple tunnels in tunnel-table that are allowed by **resolution any** or **resolution filter** mode and that can resolve the BGP next-hop, the selection of the resolving tunnel comes down to factors such as route color, admin-tag-policy, tunnel-table preference, and LDP FEC prefix length.

If **disallow-igp** is enabled, then there is no attempt to resolve the IPv6 BGP route using route table lookup if no resolving tunnel can be found in the tunnel-table.

The available tunneling options for IPv6 BGP routes with IPv4-mapped IPv6 next-hops are:

- **bgp** — This refers to IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32 IPv4 prefixes
- **ldp** — This refers to /32 and shorter length LDP FEC prefixes imported into the tunnel table. For resolution purposes, BGP selects the LDP FEC that is the longest-prefix-match (LPM) of the BGP next-hop address
- **rsvp** — This refers to RSVP tunnels in tunnel-table. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id
- **sr-isis** — This refers to segment routing tunnels (shortest path) to destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in tunnel-table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance
- **sr-ospf** — This refers to segment routing tunnels (shortest path) to destinations reachable by the OSPF protocol. This option allows BGP to use the segment routing tunnel in tunnel-table submitted by the lowest preference OSPF instance or (in case of a tie) the lowest numbered OSPF instance
- **sr-policy** — This refers to segment routing policies that are statically configured in the local router or learned via BGP routes (AFI 1/SAFI 73). For BGP to resolve the next-hop of an IPv4 route using an sr-policy the highest numbered color extended community attached to the IPv4 route must match the color of the SR policy and if the CO bits of this color extended community have the value 00 the BGP next-hop of the route must exactly match the endpoint of the SR policy.

- **sr-te** — This refers to traffic engineered (TE) segment routing tunnels. This option allows BGP to use the best metric SR-TE tunnel to the address of the BGP next-hop. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel-id

### 5.5.3.7.3 Next-Hop Resolution of BGP Labeled Routes to Tunnels

Use the following CLI syntax to configure next-hop resolution of BGP labeled routes.

**CLI Syntax:**

```
config>router>bgp>next-hop-res
 labeled-routes
 transport-tunnel
 [no] family {label-ipv4|label-ipv6|vpn}
 resolution {any | disabled | filter}
 resolution-filter
 [no] bgp
 [no] ldp
 [no] rsvp
 [no] sr-isis
 [no] sr-ospf
 [no] sr-te
 [no] udp
```

The **transport-tunnel** context provides separate control for the different types of BGP label routes: label-IPv4, label-IPv6, and VPN routes (which includes both VPN-IPv4 and VPN-IPv6 routes). By default, all labeled routes resolve to LDP (even if the preceding CLI commands are not configured in the system).

If the **resolution** option is set to **disabled**, the default binding to LDP tunnels resumes. If **resolution** is set to **any**, the supported tunnel type selection is based on TTM preference. The order of preference of TTM tunnels is: RSVP, SR-TE, LDP, segment routing OSPF, segment routing IS-IS, and UDP.

The **rsvp** option instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. The address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

The **ldp** option instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **bgp** option instructs BGP to search for a BGP LSP with a RFC 3107 label route prefix matching the address of the BGP next-hop.

When the **sr-isis** or **sr-ospf** option is enabled, an SR tunnel to the BGP next-hop is selected in the TTM from the lowest preference IS-IS or OSPF instance. If many instances have the same lowest preference, the lowest numbered IS-IS or OSPF instance is chosen.

The **sr-te** value launches a search for the best metric SR-TE LSP to the address of the BGP next-hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

The **udp** value instructs BGP to look for an MPLSoUDP tunnel to the address of the BGP next-hop.

If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types are selected again following the TTM preference. The **resolution** command must be set to **filter** to activate the list of tunnel-types configured in **resolution-filter**.

### Entropy Label for RFC 3107 BGP Labeled Routes

The router supports the MPLS entropy label, as specified in RFC 6790, on RFC 3107 BGP labeled routes. LSR nodes in a network can load-balance labeled packets in a more granular way than by hashing on the standard label stack. Refer to the *MPLS Guide* for more information.

### 5.5.3.8 Next-Hop Tracking

In SR OS next-hop resolution is not a one-time event. If the IP route or tunnel that was used to resolve a BGP next-hop is withdrawn due to a failure or configuration change an attempt is made to re-resolve the BGP next-hop using the next-best route or tunnel. If there are no more eligible routes or tunnels to resolve the BGP next-hop then the BGP next-hop becomes unresolved. The continual process of monitoring and reacting to resolving route/tunnel changes is called next-hop tracking. In SR OS next-hop tracking is completely event driven as opposed to timer driven; this provides the best possible convergence performance.

### 5.5.3.9 Next-Hop Indirection

SR OS supports next-hop indirection for most types of BGP routes. Next-hop indirection means BGP next-hops are logically separated from resolved next-hops in the forwarding plane (IOMs). This separation allows routes that share the same BGP next-hops to be grouped so that when there is a change to the way a BGP next-hop is resolved only one forwarding plane update is needed, as opposed to one update for every route in the group. The convergence time after the next-hop resolution change is uniform and not linear with the number of prefixes; in other words the next-hop indirection is a technology that supports *prefix independent convergence* (PIC). SR OS uses next-hop indirection whenever possible; there is no option to disable the functionality.

### 5.5.4 MED

The Multi-Exit Discriminator (MED) attribute is an optional attribute that can be added to routes advertised to an EBGP peer to influence the flow of inbound traffic to the AS. The MED attribute carries a 32-bit metric value. A lower metric is better than a higher metric when MED is compared by the BGP decision process. Unless the **always-compare-med** command is configured MED is compared only if the routes come from the same neighbor AS. By default if a route is received without a MED attribute it is evaluated by the BGP decision process as though it had a MED containing the value 0, but this can be changed so that a missing MED attribute is handled the same as a MED with the maximum value. SR OS always removes the received MED attribute when advertising the route to an EBGP peer.

#### 5.5.4.1 Deterministic MED

Deterministic MED is an optional enhancement to the BGP decision process that causes BGP to group paths that are equal up to the MED comparison step based on the neighbor AS. BGP compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without **deterministic-med**, the overall best path selection is sometimes dependent on the order of route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.



**Note:** When BGP routes are leaked into a target BGP RIB, they are not grouped (in a deterministic MED context) with routes learned by that target RIB, even if the neighbor AS happens to be the same.

---

## 5.5.5 Local Preference

The LOCAL\_PREF attribute is a well-known attribute that should be included in every route advertised to an IBGP or confederation-EBGP peer. It is used to influence the flow of outbound traffic from the AS. The local preference is a 32-bit value and higher values are more preferred by the BGP decision process. The LOCAL\_PREF attribute is not included in routes advertised to EBGP peers. (If the attribute is received from an EBGP peer it is ignored.)

In SR OS the default local preference is 100 but this can be changed with the **local-preference** command or using route policies. When a LOCAL\_PREF attribute needs to be added to a route because it does not have one (e.g. because it was received from an EBGP peer) the value is the configured or default **local-preference** unless overridden by policy.

## 5.5.6 Route Aggregation Path Attributes

An aggregate route is a configured IP route that is activated and installed in the routing table when it has at least one *contributing* route. A route R contributes to an aggregate route S1 if:

- The prefix length of R is greater than the prefix length of S1
- The prefix bits of R match the prefix bits of S1 up to the prefix length of S1
- There is no other aggregate route S2 with a longer prefix length than S1 that meets the previous two conditions
- R is actively used for forwarding and is not an aggregate route

When an aggregate route is activated by a router, it is not installed in the forwarding table by default. In general though it is advisable to specify the **black-hole** next-hop option for an aggregate route so that when it is activated it is installed in the forwarding table with a black-hole next-hop; this avoids the possibility of creating a routing loop. SR OS also supports the option to program an aggregate route into the forwarding table with an **indirect** next-hop; in this case packets matching the aggregate route but not a more-specific contributing route are forwarded towards the indirect next-hop rather than discarded.

An active aggregate route can be advertised to a BGP peer (by exporting it into BGP) and this can avoid the need to advertise the more-specific contributing routes to the peer, reducing the number of routes in the peer AS and improving overall scalability. When a router advertises an aggregate route to a BGP peer the attributes in the route are set as follows:

- The ATOMIC\_AGGREGATE attribute is included in the route if at least one contributing route has the ATOMIC\_AGGREGATE attribute or the aggregate route was formed without the **as-set** option and at least one contributing route has a non-empty AS\_PATH. The ATOMIC\_AGGREGATE attribute indicates that some of the AS numbers present in the AS paths of the contributing routes are missing from the advertised AS\_PATH.
- The AGGREGATOR attribute is added to the route. This attribute encodes, by default, the global AS number (or confederation ID) and router ID (BGP identifier) of the router that formed the aggregate, but these values can be changed on a per aggregate route basis using the **aggregator** command option. The AS number in the AGGREGATOR attribute is either 2 bytes or 4 bytes (if the 4-octet ASN capability was announced by both peers). The router ID in the aggregate routes advertised to a particular set of peers can be set to 0.0.0.0 using the **aggregator-id-zero** command.
- The BGP next-hop is set to the local-address used with the peer receiving the route regardless of the BGP next-hops of the contributing routes.
- The ORIGIN attribute is based on the ORIGIN attributes of the contributing routes as described in RFC 4271.
- The information in the AS\_PATH attribute depends on the **as-set** option of the aggregate route.
  - If the **as-set** option is not specified the AS\_PATH of the aggregate route starts as an empty AS path and has elements added per the description in the section titled [AS Path](#).
  - If the **as-set** option is specified and all the contributing routes have the same AS\_PATH then the AS\_PATH of the aggregate route starts with that common AS\_PATH and has elements added per the description in the section titled [AS Path](#).
  - If the **as-set** option is specified and some of the contributing routes have different AS paths the AS\_PATH of the aggregate route starts with an AS\_SET and/or an AS\_CONFED\_SET and then adds elements per the description in the section titled [AS Path](#).
- The COMMUNITY attribute contains all the communities from all the contributing routes.
- No MED attribute is included by default.



**Note:** SR OS does not require all the contributing routes to have the same MED value.



---

## 5.5.7 Community Attributes

A BGP route can be associated with one or more communities. There are three kinds of BGP communities:

- standard communities (each 4 bytes in length, all packed into a path attribute with type code 8)
- extended communities (each 8 bytes in length, potentially many possible subtypes, all packed into a path attribute with type code 16)
- large communities (each 12 bytes in length, all packed into a path attribute with type code 32)

### 5.5.7.1 Standard Communities

In a standard community, the first two bytes usually encode the AS number of the administrative entity that assigned the value in the last two bytes. In SR OS, a standard community value is configured using the format `<asnum:comm-value>`; the colon is a required separator character. In route policy applications, multiple standard community values can be matched with a regular expression in the format `<regex1>:<regex2>`, where *regex1* and *regex2* are two regular expressions that are evaluated one numerical digit at a time.

The following well-known standard communities are understood and acted upon accordingly by SR OS routers.

- **NO\_EXPORT**: When a route carries this community, it must not be advertised outside a confederation boundary (for example, to EBGp peers).
- **NO\_ADVERTISE**: When a route carries this community, it must not be advertised to any other BGP peer.
- **NO\_EXPORT\_SUBCONFED**: When a route carries this community, it must not be advertised outside a member AS boundary (for example, to confed-EBGP peers or EBGp peers).
- **LLGR\_STALE**: When a route carries this community, it indicates that the route was propagated by a router that is a long-lived graceful restart helper and normally (in the absence of LLGR) the route would have been withdrawn.
- **NO\_LLGR**: When a route carries this community, it indicates that the route should not be retained and used past the normal graceful restart window of time.
- **BLACKHOLE**: When a route carries this community, it indicates that the route should be installed into the FIB with a blackhole next-hop.

Standard communities can be added to or removed from BGP routes using BGP import and export policies. When a BGP route is locally originated by exporting a static or aggregate route into BGP, and the static or aggregate route has one or more standard communities, these community values are automatically added to the BGP route. This may affect the advertisement of the locally originated route if one of the well-known communities is associated with the static or aggregate route.

To remove all the standard communities from all routes advertised to a BGP peer, use the **disable-communities standard** command.

### 5.5.7.2 Extended Communities

Extended communities serve specialized roles. Each extended community is eight bytes. The first one or two bytes identifies the type or sub-type and the remaining six or seven bytes identify a value. Some of the more common extended communities supported by SR OS include:

- Transitive 2-octet AS-specific
  - Route target (type 0x0002)
  - Route origin (type 0x0003)
  - OSPF domain ID (type 0x0005)
  - Source AS (type 0x0009)
  - L2VPN identifier (type 0x000A)
- Non-transitive 2-octet AS-specific
  - Link bandwidth (0x4004)
- Transitive 4-octet AS-specific
  - Route target (type 0x0202)
  - Route origin (type 0x0203)
  - OSPF domain ID (type 0x0205)
  - Source AS (type 0x0209)
- Transitive IPv4-address-specific
  - Route target (type 0x0102)
  - Route origin (type 0x0103)
  - OSPF domain ID (type 0x0105)
  - L2VPN identifier (type 0x010A)
  - VRF route import (type 0x010B)
- Transitive opaque
  - OSPF route type (type 0x0306)

- Color extended community (type 0x030B)
- Non-transitive opaque
  - BGP origin validation state (type 0x4300)
- Transitive experimental
  - FlowSpec traffic rate (type 0x8006)
  - FlowSpec traffic action (type 0x8007)
  - FlowSpec redirect (type 0x8008)
  - FlowSpec traffic-remarking (0x8009)
  - Layer 2 info (type 0x800A)
- Transitive FlowSpec
  - Flowspec interface-set (type 0x0702)
- Non-transitive FlowSpec
  - Flowspec interface-set (type 0x4702)

Extended communities can be added to or removed from BGP routes using BGP import and export policies. When a BGP route is locally originated by exporting a static or aggregate route into BGP, and the static or aggregate route has one or more extended communities, these community values are automatically added to the BGP route.



**Note:** While it may not make sense to add certain types of extended communities to routes of certain address families, SR OS allows such actions.

To remove all the extended communities from all routes advertised to a BGP peer, use the **disable-communities extended** command.

### 5.5.7.3 Large Communities

Each large community is a 12-byte value, formed from the logical concatenation of three 4-octet values: a Global Administrator part, a Local Data part 1, and Local Data part 2. The Global Administrator is a four-octet namespace identifier, which should be an Autonomous System Number assigned by IANA. The Global Administrator field is intended to allow different Autonomous Systems to define large communities without collision. Local Data Part 1 is a four-octet operator-defined value and Local Data Part 2 is another four-octet operator-defined value.

In SR OS, a large community value is configured using the format `<ext-asnum>:<ext-comm-val>:<ext-comm-val>`; the colon is a required separator character between each of the 4-byte values. In route policy applications, it is possible to match multiple large community values with a regular expression in the format `<regex1>&<regex2>&<regex3>`, where *regex1*, *regex2* and *regex3* are three regular expressions, each evaluated one numerical (decimal) digit at a time.

Large communities can be added to or removed from BGP routes using BGP import and export policies. When a BGP route is locally originated by exporting a static or aggregate route into BGP, and the static or aggregate route has one or more large communities, these community values are automatically added to the BGP route.

To remove all the standard communities from all routes advertised to a BGP peer, use the **disable-communities large** command.

## 5.5.8 Route Reflection Attributes

The ORIGINATOR\_ID and CLUSTER\_LIST are optional non-transitive attributes that play a role in route reflection, as described in the section titled [Route Reflection](#).

## 5.5.9 Multi-Protocol BGP Attributes

As discussed in the BGP chapter overview the uses of BGP have increased well beyond Internet IPv4 routing due to its support for multi-protocol extensions, or more simply MP-BGP. MP-BGP allows BGP peers to exchange routes for NLRI other than IPv4 prefixes - for example IPv6 prefixes, Layer 3 VPN routes, Layer 2 VPN routes, flow-spec rules, etc. A BGP router that supports MP-BGP indicates the types of routes it wants to exchange with a peer by including the corresponding AFI (Address Family Identifier) and SAFI (Subsequent Address Family Identifier) values in the MP-BGP capability of its OPEN message. The two peers forming a session do not need indicate support for the same address families; as long as there is one AFI/SAFI in common the session will establish and routes associated with all the common AFI/SAFI can be exchanged between the peers.

The list of AFI/SAFI advertised in the MP-BGP capability is controlled entirely by the **family** commands. The AFI/SAFI supported by the SR OS and the method of configuring the AFI/SAFI support is summarized in [Table 47](#).

**Table 47 Multi-Protocol BGP support in SR OS**

| Name                 | AFI | SAFI | Configuration Commands       |
|----------------------|-----|------|------------------------------|
| IPv4 unicast         | 1   | 1    | <b>family ipv4</b>           |
| IPv4 multicast       | 1   | 2    | <b>family mcast-ipv4</b>     |
| IPv4 labeled unicast | 1   | 4    | <b>family label-ipv4</b>     |
| NG-MVPN IPv4         | 1   | 5    | <b>family mvpn-ipv4</b>      |
| MDT-SAFI             | 1   | 66   | <b>family mdt-safi</b>       |
| VPN-IPv4             | 1   | 128  | <b>family vpn-ipv4</b>       |
| VPN-IPv4 multicast   | 1   | 129  | <b>family mcast-vpn-ipv4</b> |
| RT constrain         | 1   | 132  | <b>family route-target</b>   |
| IPv4 flow-spec       | 1   | 133  | <b>family flow-ipv4</b>      |
| IPv6 unicast         | 2   | 1    | <b>family ipv6</b>           |
| IPv6 multicast       | 2   | 2    | <b>family mcast-ipv6</b>     |
| IPv6 labeled unicast | 2   | 4    | <b>family label-ipv6</b>     |
| NG-MVPN IPv6         | 2   | 5    | <b>family mvpn-ipv6</b>      |
| VPN-IPv6             | 2   | 128  | <b>family vpn-ipv6</b>       |
| IPv6 flow-spec       | 2   | 133  | <b>family flow-ipv6</b>      |
| Multi-segment PW     | 25  | 6    | <b>family ms-pw</b>          |
| L2 VPN               | 25  | 65   | <b>family l2-vpn</b>         |
| EVPN                 | 25  | 70   | <b>family evpn</b>           |

To advertise reachable routes of a particular AFI/SAFI a BGP router includes a single MP\_REACH\_NLRI attribute in the UPDATE message. The MP\_REACH\_NLRI attribute encodes the AFI, the SAFI, the BGP next-hop and all the reachable NLRI. To withdraw routes of a particular AFI/SAFI a BGP router includes a single MP\_UNREACH\_NLRI attribute in the UPDATE message. The MP\_UNREACH\_NLRI attribute encodes the AFI, the SAFI and all the withdrawn NLRI. While it is valid to advertise and withdraw IPv4 unicast routes using the MP\_REACH\_NLRI and MP\_UNREACH\_NLRI attributes, SR OS always uses the IPv4 fields of the UPDATE message to convey reachable and unreachable IPv4 unicast routes.

---

## 5.5.10 4-Octet AS Attributes

The AS4\_PATH and AS4\_AGGREGATOR path attributes are optional transitive attributes that support the gradual migration of routers that can understand and parse 4-octet ASN numbers. The use of these attributes is discussed in the section titled [4-Octet Autonomous System Numbers](#).

## 5.5.11 AIGP Metric

The accumulated IGP (AIGP) metric is an optional non-transitive attribute that can be attached to selected routes (using route policies) to influence the BGP decision process to prefer BGP paths with a lower end-to-end IGP cost, even when the compared paths span more than one AS or IGP instance. AIGP is different from MED in several important ways:

- AIGP is not intended to be transitive between completely distinct autonomous systems (only across internal AS boundaries)
- AIGP is always compared in paths that have the attribute, regardless of whether or not they come from different neighbor AS
- AIGP is more important than MED in the BGP decision process (see the section titled [BGP Decision Process](#))
- AIGP is automatically incremented every time there is a BGP next-hop change so that it can track the end-to-end IGP cost. All arithmetic operations on MED attributes must be done manually (for example, using route policies).

In the SR OS implementation, AIGP is supported only in the base router BGP instance and only for the following types of routes: IPv4, label-IPv4, IPv6 and label-IPv6. The AIGP attribute is only sent to peers configured with the **aigp** command. If the attribute is received from a peer that is not configured for **aigp** or if the attribute is received in a non-supported route type the attribute is discarded and not propagated to other peers (but it is still displayed in BGP show commands).

When a 7450, 7750, or 7950 router receives a route with an AIGP attribute and it re-advertises the route to an AIGP-enabled peer without any change to the BGP next-hop the AIGP metric value is unchanged by the advertisement (RIB-OUT) process. But if the route is re-advertised with a new BGP next-hop the AIGP metric value is automatically incremented by the route table (or tunnel table) cost to reach the received BGP next-hop and/or by a statically configured value (using route policies).

---

## 5.6 BGP Routing Information Base (RIB)

The entire set of BGP routes learned and advertised by a BGP router make up its BGP Routing Information Base (RIB). Conceptually the BGP RIB can be divided into 3 parts:

- RIB-IN
- LOC-RIB
- RIB-OUT

The RIB-IN (or Adj-RIBs-In as defined in RFC 4271) holds the BGP routes that were received from peers and that the router decided to keep (store in its memory).

The LOC-RIB contains modified versions of the BGP routes in the RIB-IN. The path attributes of a RIB-IN route can be modified using BGP import policies. All of the LOC-RIB routes for the same NLRI are compared in a procedure called the BGP decision process that results in the selection of the best path for each NLRI. The best paths in the LOC-RIB are the ones that are actually 'usable' by the local router for forwarding, filtering, auto-discovery, etc.

The RIB-OUT (or Adj-RIBs-Out as defined in RFC 4271) holds the BGP routes that were advertised to peers. Normally a BGP route is not advertised to a peer (in the RIB-OUT) unless it is 'used' locally but there are exceptions. BGP export policies modify the path attributes of a LOC-RIB route to create the path attributes of the RIB-OUT route. A particular LOC-RIB route can be advertised with different path attribute values to different peers so there can exist a 1:N relationship between LOC-RIB and RIB-OUT routes.

The following sections describe many important BGP features in the context of the RIB architecture outlined above.

### 5.6.1 RIB-IN Features

SR OS implements the following features related to RIB-IN processing:

- UPDATE message fault tolerance. This is described in the section titled [UPDATE Message Error Handling](#).
- BGP import policies

### 5.6.1.1 BGP Import Policies

The **import** command is used to apply one or more policies (up to 15) to a neighbor, group or to the entire BGP context. The **import** command that is most-specific to a peer is the one that is applied. An **import** policy command applied at the **neighbor** level takes precedence over the same command applied at the **group** or global level. An **import** policy command applied at the **group** level takes precedence over the same command specified on the global level. The **import** policies applied at different levels are not cumulative. The policies listed in an **import** command are evaluated in the order in which they are specified.



**Note:** The **import** command can reference a policy before it has been created (as a **policy-statement**).

When an IP route is rejected by an import policy it is still maintained in the RIB-IN so that a policy change can be made later on without requiring the peer to re-send all its RIB-OUT routes. This is sometimes called soft reconfiguration inbound and requires no special configuration in SR OS.

When a VPN route is rejected by an import policy or not imported by any services it is deleted from the RIB-IN. For VPN-IPv4 and VPN-IPv6 routes this behavior can be changed by configuring the **mp-bgp-keep** command; this option maintains rejected VPN-IP routes in the RIB-IN so that a Route Refresh message does not have to be issued when there is an import policy change.

### 5.6.2 LOC-RIB Features

SR OS implements the following features related to LOC-RIB processing.

- BGP decision process
- BGP route installation in the route table
- BGP route installation in the tunnel table
- BGP fast reroute
- QoS Policy Propagation via BGP (QPPB)
- Policy accounting
- Route flap damping (RFD)

These features are discussed in the following sections.



---

### 5.6.2.1 BGP Decision Process

When a BGP router has multiple paths in its LOC-RIB for the same NLRI, its BGP decision process is responsible for deciding which path is the best. The best path can be used by the local router and advertised to other BGP peers.

On 7450 ESS, 7750 SR, and 7950 XRS routers, the BGP decision process orders received paths based on the following sequence of comparisons. If there is a tie between paths at any step, BGP proceeds to the next step.

1. Select a valid route over an invalid route.
2. Select the route with the lowest origin validation state, where Valid < Not-Found < Invalid.
3. Select the route with the numerically lowest route-table preference.
4. Select the route with the highest Local Preference.
5. Select the route with an AIGP metric. If they both have an AIGP metric, select the route with the lowest sum of:
  - i. AIGP metric value stored with the LOC-RIB copy of the route
  - ii. route-table or tunnel-table cost between the calculating router and the BGP next-hop in the received route
6. Select the route with the shortest AS Path. AS numbers in AS\_CONFED\_SEQ and AS\_CONFED\_SET elements do not count towards the AS path length. This step is skipped if **as-path-ignore** is configured for the address family.
7. Select the route with the lowest Origin (IGP < EGP < Incomplete).
8. Select the route with the lowest MED. Only compare MED for non-imported routes that have the same neighbor AS by default. A missing MED attribute is considered equivalent to a MED value by 0 of default. Defaults can be changed by using the **always-compare-med** command.
9. Select the route with the lowest owner type (BGP < BGP-LABEL < BGP-VPN). This information is used when BGP routes are imported from other RIBs.
10. Prefer routes learned from EBGP peers over routes learned from IBGP and confed-EBGP peers.
11. Select the route with the lowest route-table or tunnel-table cost to the NEXT\_HOP. This step is skipped if **ignore-nh-metric** is configured, or if the routes are associated with different RIBs.
12. Select the route with the lowest next-hop type. Resolutions made in the tunnel-table are preferred to resolutions made in the route-table. This step is skipped if **ignore-nh-metric** is configured, or if the routes are associated with different RIBs.

13. Select the route received from the peer with the lowest router ID; this comes from the ORIGINATOR\_ID attribute (if present) or the BGP identifier of the peer (received in its OPEN message). If **ignore-router-id** is configured, keep the current best path and skip the remaining steps.
14. Select the route with the shortest CLUSTER\_LIST length.
15. Select the route received from the peer with the lowest IP address.

### 5.6.2.1.1 Always Compare MED

By default, the MED path attribute is used in the decision process only if the routes being compared come from the same neighbor AS; if one of the paths lacks a MED attribute it is considered equal to a route with a MED of 0. These default rules can be modified using the **always-compare-med** command.

The **always-compare-med** command without the **strict-as** keyword allows MED to be compared in paths from different neighbor autonomous systems and from different route owners; in this case, if neither **zero** or **infinity** is part of the command, **zero** is inferred, meaning that a route without a MED attribute is handled as though it had a MED with value 0. When the **strict-as** keyword is present MED is only compared between paths from the same neighbor AS and in this case **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

[Table 48](#) shows how the MED comparison of two paths is influenced by different forms of the **always-compare-med** command.

**Table 48 MED Comparison with always-compare-med**

| Command                                                    | MED comparison step in decision process                                                                                                                      |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| no always-compare-med<br>always-compare-med strict-as zero | Only compare the MED of two paths if they come from the same neighbor AS. If one path is missing a MED attribute treat it the same as MED=0.                 |
| always-compare-med<br>always-compare-med zero              | Always compare the MED of two paths, even if they come from different neighbor AS. If one path is missing a MED attribute treat it the same as MED=0.        |
| always-compare-med infinity                                | Always compare the MED of two paths, even if they come from different neighbor AS. If one path is missing a MED attribute treat it the same as MED=infinity. |
| always-compare-med strict-as infinity                      | Only compare the MED of two paths if they come from the same neighbor AS. If one path is missing a MED attribute treat it the same as MED=infinity.          |

### 5.6.2.1.2 Ignore Next-Hop Metric

The **ignore-nh-metric** command allows the step comparing the distance to the BGP next-hop to be skipped. When this command is present in the **config>service>vprn** context it applies to the comparison of two imported BGP-VPN routes. When this command is present in the **config>router>bgp** context it applies to the comparison of any two BGP routes received by that instance. And when this command is present in the **config>service>vprn>bgp** context it applies to the comparison of two BGP routes learned from VPRN BGP peers (that is, CE peers). In all cases, this option is useful when there are multiple paths for a prefix that are equally preferred up to (but not including) the IGP cost comparison step of the BGP decision process and the network administrator wants all of them to be used for forwarding (*BGP-Multipath*).

### 5.6.2.2 BGP Route Installation in the Route Table

Each BGP RIB holding routes (unlabeled IPv4, labeled-unicast IPv4, unlabeled IPv6, labeled-unicast IPv6) submits its best path for each prefix to the common IP route table, unless **disable-route-table-install** is configured. It is up to the route table to choose the single best of these paths for forwarding to each IP prefix destination. The route table chooses the route by using the BGP decision process. The default preference for BGP routes submitted by the label-IPv4 and label-IPv6 RIBs (which appear in the route table and FIB as having a BGP-LABEL protocol type) can be modified using the **label-preference** command. The default preference for BGP routes submitted by the unlabeled IPv4 and IPv6 RIBs can be modified by using the **preference** command.



**Note:** Consider configuring the **disable-route-table-install** command on control-plane route reflectors that are not involved in packet forwarding (i.e. that do not modify the BGP NEXT\_HOP); this improves the performance and scalability of such route reflectors.

If a BGP RIB has multiple BGP paths (LOC-RIB routes) for the same IPv4 or IPv6 prefix that qualify as the best path up to a certain point in the comparison process, then a certain number of these multipaths can be submitted to the common IP route table. This is called BGP Multipath and it must be explicitly enabled using the **multipath** command. The **multipath** command specifies the maximum number (64) of BGP paths, including the overall best path, that each BGP RIB can submit to the route table for any particular IPv4 or IPv6 prefix. If ECMP, with a limit of  $n$ , is enabled in the base router instance, then up to  $n$  paths are selected for installation in the IP FIB. In the datapath, traffic matching the IP route is load-shared across the ECMP next-hops based on a per-packet hash calculation.

By default the hashing is not sticky, meaning that when one or more of the equal-cost BGP next-hops fail, all traffic flows matching the route are potentially moved to new BGP next-hops. If required, a BGP route can be marked (using the **sticky-ecmp** action in route policies) for sticky ECMP behavior so that BGP next-hop failures are handled by moving only the affected traffic flows to the remaining next-hops as evenly as possible.

In the route table, a BGP path to an IPv4 or IPv6 prefix is a candidate for installation as an ECMP next-hop (subject to the path limits of the **multipath** and **ecmp** commands) only if it meets all of the following criteria:

1. It is the overall best BGP path or it is tied with the overall best path up to and including step 12 of the decision process summarized in the [BGP Decision Process](#) section.
2. Compared to other paths with the same BGP NEXT\_HOP it is the best path, based on evaluation of all steps of the BGP decision process.



**Note:** VPRN routing instances support a special mode of BGP multipath called *EIBGP-Multipath*. In *EIBGP-Multipath* BGP routes learned from CE devices that are typically EBGPeers are combined with imported VPN-IP routes that typically come from IBGP peers to form an IP ECMP route. When *EIBGP-Multipath* is enabled a route is a candidate for installation as an ECMP next-hop if it is the overall best route or else it is tied with the overall best route up to and including the MED step of the BGP decision process.

3. It is the same type of route as the best path and all other multipaths, meaning it came from the same BGP RIB (labeled unicast or unlabeled).

SR OS also supports a feature called *IBGP-Multipath*. In some topologies a BGP next-hop is resolved by an IP route (for example a static, OSPF or IS-IS route) that itself has multiple ECMP next-hops. When **ibgp-multipath** is not configured only one of these ECMP next-hops is programmed as a next-hop of the BGP route in the IOM. But when **ibgp-multipath** is configured the IOM attempts to use all of the ECMP next-hops of the resolving route in forwarding.

Although the name of the **ibgp-multipath** command implies that it is specific to IBGP-learned routes this is not the case; it applies to routes learned from any multi-hop BGP session including routes learned from multi-hop EBGPeers.



**Note:** *BGP-Multipath* and *IBGP-Multipath* are not mutually exclusive and work together. *BGP-Multipath* enables ECMP load-sharing across different BGP next-hops (corresponding to different LOC-RIB routes) and *IBGP-Multipath* enables ECMP load-sharing across different IP next-hops of IP routes that resolve the BGP next-hops.

The final point about *IBGP-Multipath* is that it does not control load-sharing of traffic towards a BGP next-hop that is resolved by a tunnel, such as the case when dealing with BGP shortcuts or labeled routes (VPN-IP, label-IPv4, label-IPv6). When a BGP next-hop is resolved by a tunnel that supports ECMP the load-sharing of traffic across the ECMP next-hops of the tunnel is automatic.



**Note:** At the current time SR OS does not support direct resolution of a BGP next-hop to multiple RSVP-TE tunnels. However a BGP next-hop can be resolved by multiple LDP ECMP next-hops that each correspond to a separate LDP-over-RSVP tunnel. It is also possible for a BGP next-hop to be resolved by an IGP shortcut route that has multiple RSVP-TE tunnels as its ECMP next-hops.

### 5.6.2.3 Weighted ECMP for BGP Routes

In some cases, the ECMP BGP next-hops of an IP route correspond to paths with very different bandwidths and it makes sense for the ECMP load-balancing algorithm to distribute traffic across the BGP next-hops in proportion to their relative bandwidths. The bandwidth associated with a path can be signaled to other BGP routers by including a link-bandwidth extended community in the BGP route. The link-bandwidth extended community is optional and non-transitive and encodes an autonomous system (AS) number and a bandwidth.

The SR OS implementation supports the link-bandwidth extended community in routes associated with the following address families: IPv4, IPv6, label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6. The router automatically performs weighted ECMP for an IP BGP route when all of the ECMP BGP next-hops of the route include a link-bandwidth extended community. The relative weight of traffic sent to each BGP next-hop is visible in the output of the **show router route-table extensive** and **show router fib extensive** commands.

A route with a link-bandwidth extended community can be received from any IBGP peer. If such a route is received from an EBGp peer, the link-bandwidth extended community is stripped from the route unless an **accept-from-ebgp** command applies to that EBGp peer. However, a link-bandwidth extended community can be added to routes received from a directly connected (single hop) EBGp peer, potentially replacing the received Extended Community. This is accomplished using the **add-to-received-ebgp** command, which is available in group and neighbor configuration contexts.

When a route with a link-bandwidth extended community is advertised to an EBGp peer, the link-bandwidth extended community is removed by default; however, transitivity across an AS boundary can be allowed by configuring the **send-to-ebgp** command.

When a route with a link-bandwidth extended community is advertised to a peer using next-hop-self, the Extended Community is usually removed if it was not added locally (that is, by policy or **add-to-received-ebgp** command). However, in the special case that a route is readvertised (with next-hop-self) toward a peer covered by the scope of an **aggregate-used-paths** command, and the re-advertising router has installed multiple ECMP paths toward the destination each associated with a link-bandwidth extended community, the route is readvertised with a link-bandwidth extended community encoding the total bandwidth of all the used multipaths.

The link-bandwidth extended community associated with a BGP route can be displayed using the **show router bgp routes** command. For the bandwidth value, the system automatically converts the binary value in the extended community to a decimal number in units of Mb/s (1 000 000 b/s).

Weighted ECMP across the BGP next-hops of an IP BGP route is supported in combination with ECMP at the level of the route or tunnel that resolves one or more of the ECMP BGP next-hops. This ECMP at the resolving level can also be weighted ECMP when the following conditions all apply:

- The BGP next-hop is resolved by an IP route (OSPF, IS-IS, or static) with MPLS LSP ECMP next-hops
- **ibgp-multipath** is configured under BGP
- **config>router>weighted-ecmp** is configured

#### 5.6.2.4 BGP Route Installation in the Tunnel Table

If the best BGP path for a /32 IPv4 prefix is a label-IPv4 route (AFI 1, SAFI 4), and if it has the numerically lowest **preference** value among all routes (regardless of protocol) for the /32 IPv4 prefix, and if **disable-route-table-install** is *not* configured, the label-IPv4 route is automatically added, as a *BGP tunnel* entry, to the tunnel table. In SR OS the tunnel-table is used to resolve a BGP next-hop to a tunnel when required by the configuration or the type of route (see the section titled [Next-Hop Resolution](#) for many of these details). BGP tunnels play a key role in the following solutions:

- Inter-AS IP VPN model C
- Inter-AS L2 VPN model C
- Carrier Supporting Carrier (CSC)
- Intra-AS seamless MPLS

BGP tunnels have a preference of 10 in the tunnel table, compared to 9 for LDP tunnels and 7 for RSVP tunnels, so if the router configuration allows all types of tunnels to resolve a BGP next-hop an RSVP LSP is preferred over an LDP tunnel and an LDP tunnel is preferred over a BGP tunnel.

If **multipath** and **ecmp** are configured appropriately a BGP tunnel can be installed in the tunnel table with multiple ECMP next-hops, each one corresponding to a path through a different BGP next-hop; the multipath selection process outlined in the previous section ([BGP Route Installation in the Route Table](#)) also applies to this case.

### 5.6.2.5 BGP Fast Reroute

BGP fast reroute is a feature that brings together indirection techniques in the forwarding plane and pre-computation of BGP backup paths in the control plane to support fast reroute of BGP traffic around unreachable/failed BGP next-hops. BGP fast reroute is supported with IPv4, label-IPv4, IPv6, label-IPv6, VPN-IPv4 and VPN-IPv6 routes. The scenarios supported by the base router BGP context are outlined in [Table 49](#).

Refer to the VPRN section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information about BGP fast reroute information specific to IP VPNs.

**Table 49 BGP Fast Reroute Scenarios (Base Context)**

| Ingress Packet | Primary Route                                                               | Backup Route                                                                | Prefix Independent Convergence                                                                          |
|----------------|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| IPv4           | IPv4 route with next-hop A resolved by an IPv4 route or any shortcut tunnel | IPv4 route with next-hop B resolved by an IPv4 route or any shortcut tunnel | Yes                                                                                                     |
| IPv4           | Label-IPv4 route with next-hop A resolved by any transport tunnel           | Label-IPv4 route with next-hop B resolved by any transport tunnel           | Yes, but if the label-IPv4 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |
| IPv4           | Label-IPv4 route with next-hop A resolved by a local route                  | Label-IPv4 route with next-hop B resolved by a local route                  | Yes, but if the label-IPv4 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |

**Table 49 BGP Fast Reroute Scenarios (Base Context) (Continued)**

| Ingress Packet | Primary Route                                                     | Backup Route                                                      | Prefix Independent Convergence                                                                          |
|----------------|-------------------------------------------------------------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| IPv4           | Label-IPv4 route with next-hop A resolved by a static route       | Label-IPv4 route with next-hop B resolved by a static route       | Yes, but if the label-IPv4 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |
| IPv6           | IPv6 route with next-hop A resolved by an IPv6 route              | IPv6 route with next-hop B resolved by an IPv6 route              | Yes                                                                                                     |
| IPv6           | Label-IPv6 route with next-hop A resolved by any transport tunnel | Label-IPv6 route with next-hop B resolved by any transport tunnel | Yes, but if the label-IPv6 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |
| IPv6           | Label-IPv6 route with next-hop A resolved by a local route        | Label-IPv6 route with next-hop B resolved by a local route        | Yes, but if the label-IPv6 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |
| IPv6           | Label-IPv6 route with next-hop A resolved by a static route       | Label-IPv6 route with next-hop B resolved by a static route       | Yes, but if the label-IPv6 routes are label-per-prefix, the ingress card must be IOM3 or better for PIC |

#### 5.6.2.5.1 Calculating Backup Paths

In SR OS, fast reroute is optional and must be enabled by using either the BGP **backup-path** command or the route-policy **install-backup-path** command. Typically only one approach is used.

The **backup-path** command in the base router context is used to control fast reroute on a per-RIB basis (IPv4, label-IPv4, IPv6, and label-IPv6). When the command specifies a particular family, BGP attempts to find a backup path for every prefix learned by the associated BGP RIB.

The **install-backup-path** command, available in route-policy-action contexts, marks a BGP route as requesting a backup path. It only takes effect in BGP import and VRF import policies. If only some prefixes should have backup paths, then the **backup-path** command should not be used, and instead the **install-backup-path** command should be used to mark only those prefixes that require extra protection.



In general, a prefix supports ECMP paths or a backup path, but not both. The backup path is the best path after the primary path and any paths with the same BGP next-hop as the primary path have been removed.

#### **5.6.2.5.2 Failure Detection and Switchover to the Backup Path**

When BGP fast reroute is enabled the IOM reroutes traffic onto a backup path based on input from BGP. When BGP decides that a primary path is no longer usable it notifies the IOM and affected traffic is immediately switched to the backup path.

The following events trigger failure notifications to the IOM and reroute of traffic to backup paths:

- Peer IP address unreachable and peer-tracking is enabled
- BFD session associated with BGP peer goes down
- BGP session terminated with peer (for example, send/receive NOTIFICATION)
- There is no longer any route (allowed by the next-hop resolution policy, if configured) that can resolve the BGP next-hop address
- The LDP tunnel that resolves the next-hop goes down. This could happen because there is no longer any IP route that can resolve the FEC, or the LDP session goes down, or the LDP peer withdraws its label mapping.
- The RSVP tunnel that resolves the next-hop goes down. This could happen because a ResvTear message is received, or the RESV state times out, or the outgoing interface fails and is not protected by FRR or a secondary path.
- The BGP tunnel that resolves the next-hop goes down. This could happen because the BGP label-IPv4 route is withdrawn by the peer or else becomes invalid due to an unresolved next-hop.

#### **5.6.2.6 QoS Policy Propagation via BGP (QPPB)**

QPPB is a feature that allows different QoS values (forwarding class and optionally priority) to be associated with different IPv4 and IPv6 BGP LOC-RIB routes based on BGP import policy processing. This is done so that when traffic arrives on a QPPB-enabled IP interface and its source or destination IP address matches a BGP route with QoS information the packet is handled according to the QoS of the matching route. SR OS supports QPPB on the following types of interfaces:

- Base router network interfaces
- IES and VPRN SAP interfaces
- IES and VPRN spoke-SDP interfaces

- IES and VPRN subscriber interfaces

QPPB is enabled on an interface using the **qos-route-lookup** command. There are separate commands for IPv4 and IPv6 so that QPPB can be enabled in one mode (source or destination or none) for IPv4 packets arriving on the interface and a different mode (source or destination or none) for IPv6 packets arriving on the interface.



**Note:** Source-based QPPB is not supported on subscriber interfaces.

Different LOC-RIB routes for the same IP prefix may be associated with different QPPB information. If these LOC-RIB routes are combined in support of ECMP or BGP fast reroute then the QPPB information becomes next-hop specific. This means that in destination QPPB mode the QoS assigned to a packet depends on the BGP next-hop that is selected for that particular packet by the ECMP hash or fast reroute algorithm. In source QPPB mode the QoS assigned to a packet comes from the first BGP next-hop of the IP route matching the source address.

### 5.6.2.7 BGP Policy Accounting

Policy accounting is a feature that allows different *accounting classes* to be associated with IPv4 and IPv6 BGP LOC-RIB routes based on BGP import policy processing. This is done so that per-accounting-class traffic statistics can be collected on policy accounting-enabled interfaces of the router. Policy accounting interfaces are only supported on IOM3 or better cards. The following types of interfaces are supported:

- Base router network interfaces
- IES and VPRN SAP interfaces
- IES and VPRN spoke-SDP interfaces
- IES and VPRN subscriber interfaces

Policy accounting is enabled on an interface using the **policy-accounting** command. The name of a policy accounting template must be specified. Each policy accounting template contains a list of *source classes* and *destination classes*. Routers support up to 255 different source classes and up to 255 different destination classes. Each source class is identified by an index number (1-255) and each destination class is identified by an index number (1-255). The policy accounting template tells the IOM what accounting classes to collect stats for on a policy accounting interface. SR OS supports up to 1024 different templates, depending on the chassis type.



**Note:** Policy accounting templates containing one or more source class identifiers cannot be applied to subscriber interfaces.

Through policy mechanisms a LOC-RIB route for an IP prefix can have a source class index (1-255 iso 25), a destination class index (1-255) or both. When an ingress packet on a policy-accounting enabled interface [I1] is forwarded by the IOM and its destination address matches a BGP route with a destination class index [D], and [D] is listed in the relevant policy accounting template, packets-forwarded and IP-bytes-forwarded counters for [D] on interface [I1] are incremented accordingly. Similarly, when an ingress packet on a policy-accounting enabled interface [I2] is forwarded by the IOM and its source address matches a BGP route with a source class index [S], and [S] is listed in the relevant policy accounting template, the packets-forwarded and IP-bytes-forwarded counters for [S] on interface [I2] are incremented accordingly.

It is possible that different LOC-RIB routes for the same IP prefix are associated with different accounting class information. If these LOC-RIB routes are combined in support of ECMP or BGP fast reroute then the destination-class of a packet depends on the BGP next-hop that is selected for that particular packet by the ECMP hash or fast reroute algorithm. If the source address of a packet matches a route with multiple BGP next-hops its source-class is derived from the first BGP next-hop of the matching route.

### 5.6.2.8 Route Flap Damping (RFD)

Route flap damping is a mechanism supported by 7450, 7750, and 7950 routers, as well as other BGP routers, that was designed to help improve the stability of Internet routing by mitigating the impact of route flaps. Route flaps describe a situation where a router alternately advertises a route as reachable and then unreachable or as reachable through one path and then another path in rapid succession. Route flaps can result from hardware errors, software errors, configuration errors, unreliable links, etc. However not all perceived route flaps represent a true problem; when a

best path is withdrawn the next-best path may not be immediately known and may trigger a number of intermediate best path selections (and corresponding advertisements) before it is found. These intermediate best path selections may travel at different speeds through different routers due to the effect of the min-route-advertisement interval (MRAI) and other factors. RFD does not handle this type of situation particularly well and for this and other reasons many Internet service providers do not use RFD.

In SR OS route flap damping is configurable; by default it is disabled. It can be enabled on EBGP and confed-EBGP sessions by including the **damping** command in their group or neighbor configuration. The **damping** command has no effect on IBGP sessions. When a route of any type (any AFI/SAFI) is received on a non-IBGP session that has **damping** enabled:

- If the route changes from reachable to unreachable due to a withdrawal by the peer then damping history is created for the route (if it does not already exist) and in that history the Figure of Merit (FOM), an accumulated penalty value, is incremented by 1024.
- If a reachable route is updated by the peer with new path attribute values then the FOM is incremented by 1024.
- In SR OS the FOM has a hard upper limit of 21540 (not configurable).
- The FOM value is decayed exponentially as described in RFC 2439. The **half-life** of the decay is 15 minutes by default, however a BGP import policy can be used to apply a non-default damping profile to the route, and the **half-life** in the non-default damping profile can have any value between 1 and 45 minutes.
- The FOM value at the last time of update can be displayed using the **show router bgp damping detail** command. The time of last update can be up to 640 seconds ago; SR OS does not calculate the current FOM every time the show command is entered.
- When the FOM reaches the suppress limit, which is 3000 by default but can be changed to any value between 1 and 20000 in a non-default damping profile, the route is suppressed, meaning it is not used locally and not advertised to peers. The route remains suppressed until either the FOM exponentially decays to a value less than or equal to the **reuse** threshold or the **max-suppress** time is reached. By default the **reuse** threshold is 750 and the **max-suppress** time is 60 minutes, but these can be changed in a non-default damping profile: **reuse** can have a value between 1 and 20000 and **max-suppress** can have a value between 1 and 720 minutes.

### 5.6.3 RIB-OUT Features

SR OS implements the following features related to RIB-OUT processing.

- BGP export policies
- Outbound route filtering (ORF)
- RT constrained route distribution
- Configurable min-route-advertisement (MRAI)
- Advertise-inactive
- Best-external
- Add-path
- Split-horizon

These features are discussed in the following sections.

### 5.6.3.1 BGP Export Policies

The **export** command is used to apply one or more policies (up to 15) to a neighbor, group or to the entire BGP context. The **export** command that is most-specific to a peer is the one that is applied. An **export** policy command applied at the **neighbor** level takes precedence over the same command applied at the **group** or global level. An **export** policy command applied at the **group** level takes precedence over the same command specified on the global level. The **export** policies applied at different levels are not cumulative. The policies listed in an **export** command are evaluated in the order in which they are specified.



**Note:** The **export** command can reference a policy before it has been created (as a **policy-statement**).

The most common uses for BGP export policies are as follows:

- To locally originate a BGP route by exporting (or redistributing) a non-BGP route that is installed in the route table and actively used for forwarding. The non-BGP route is most frequently a direct, static or aggregate route (exporting IGP routes into BGP is generally not recommended).
- To block the advertisement of certain BGP routes towards specific BGP peers. The routes may be blocked on the basis of IP prefix, communities, etc.
- To modify the attributes of BGP routes advertised to specific BGP peers. The following path attribute modifications are possible using BGP export policies:
  - Change the ORIGIN value
  - Add a sequence of AS numbers to the start of the AS\_PATH. When a route is advertised to an EBGP peer the addition of the local-AS/global-AS numbers to the AS\_PATH is always the final step (done after export policy).

- Replace the AS\_PATH with a new AS\_PATH. When a route is advertised to an EBGp peer the addition of the local-AS/global-AS numbers to the AS\_PATH is always the final step (done after export policy).
- Prepend an AS number multiple times to the start of the AS\_PATH. When a route is advertised to an EBGp peer the addition of the local-AS/global-AS numbers to the AS\_PATH is always the final step (done after export policy). The add/replace action on the AS\_PATH supersedes the prepend action if both are specified in the same policy entry.
- Change the NEXT\_HOP to a specific IP address. When a route is advertised to an EBGp peer the next-hop cannot be changed from the local-address.
- Change the NEXT\_HOP to the local-address used with the peer (next-hop-self).
- Add a value to the MED. If the MED attribute does not exist it is added.
- Subtract a value from the MED. If the MED attribute does not exist it is added with a value of 0. If the result of the subtraction is a negative number the MED metric is set to 0.
- Set the MED to a particular value.
- Set the MED to the cost of the IP route (or tunnel) used to resolve the BGP next-hop.
- Set LOCAL\_PREF to a particular value when advertising to an IBGP peer.
- Add, remove and/or replace standard communities
- Add, remove and/or replace extended communities
- Add a static value to the AIGP metric when advertising the route to an AIGP-enabled peer with a modified BGP next-hop. The static value is incremental to the automatic adjustment of the LOC-RIB AIGP metric to reflect the distance between the local router and the received BGP next-hop.
- Increment the AIGP metric by a fixed amount when advertising the route to an AIGP-enabled peer with a modified BGP next-hop. The static value is a substitute for the dynamic value of the distance between the local router and the received BGP next-hop.

### 5.6.3.2 Outbound Route Filtering (ORF)

Outbound route filtering (ORF) is a mechanism that allows one router, the ORF-sending router to signal to a peer, the ORF-receiving router, a set of route filtering rules (ORF entries) that the ORF-receiving router should apply to its route advertisements towards the ORF-sending router. The ORF entries are encoded in Route Refresh messages.

The use of ORF on a session must be negotiated —i.e. both routers must advertise the ORF capability in their Open messages. The ORF capability describes the address families that support ORF, and for each address family, the ORF types that are supported and the ability to send/receive each type. 7450, 7750, and 7950 routers support ORF type 3, which is ORF based on Extended Communities. It is supported for only the following address families:

- VPN-IPv4
- VPN-IPv6
- MVPN-IPv4
- MVPN-IPv6

In SR OS the send/receive capability for ORF type 3 is configurable (with the **send-orf** and **accept-orf** commands) but the setting applies to all supported address families.

SR OS support for ORF type 3 allows a PE router that imports VPN routes with a particular set of Route Target Extended Communities to indicate to a peer (for example a route reflector) that it only wants to receive VPN routes that contain one or more of these Extended Communities. When the PE router wants to inform its peer about a new RT Extended Community it sends a Route Refresh message to the peer containing an ORF type 3 entry instructing the peer to *add a permit* entry for the 8-byte extended community value. When the PE router wants to inform its peer about a RT Extended Community that is no longer needed it sends a Route Refresh message to the peer containing an ORF type 3 entry instructing the peer to *remove* the *permit* entry for the 8-byte extended community value.

In SR OS the type-3 ORF entries that are sent to a peer can be generated dynamically (if no Route Target Extended Communities are specified with the **send-orf** command) or else specified statically. Dynamically generated ORF entries are based on the route targets that are imported by all locally-configured VPRNs.

A router that has installed ORF entries received from a peer can still apply BGP export policies to the session. If the evaluation of a BGP export policy results in a reject action for a VPN route that matches a permit ORF entry the route is not advertised — i.e. the export policy has the final word.



**Note:** The SR OS implementation of ORF filtering is very efficient. It takes less time to filter a large number of VPN routes with ORF than it does to reject non-matching VPN routes using a conventional BGP export policy.

Despite the advantages of ORF compared to manually configured BGP export policies a better technology, when it comes to dynamic filtering based on Route Target Extended Communities, is RT Constraint. RT Constraint is discussed further in the next section.

### 5.6.3.3 RT Constrained Route Distribution

RT constrained route distribution, or RT-constrain for short, is a mechanism that allows a router to advertise to certain peers a special type of MP-BGP route called an RTC route; the associated AFI is 1 and the SAFI is 132. The NLRI of an RTC route encodes an Origin AS and a Route Target Extended Community with prefix-type encoding (i.e. there is a prefix-length and “host” bits after the prefix-length are set to zero). A peer receiving RTC routes does not advertise VPN routes to the RTC-sending router unless they contain a Route Target Extended Community that matches one of the received RTC routes. As with any other type of BGP route RTC routes are propagated loop-free throughout and between Autonomous Systems. If there are multiple RTC routes for the same NLRI the BGP decision process selects one as the best path. The propagation of the best path installs RIB-OUT filter rules as it travels from one router to the next and this process creates an optimal VPN route distribution tree rooted at the source of the RTC route.



**Note:** RT-constrain and Extended Community-based ORF are similar to the extent that they both allow a router to signal to a peer the Route Target Extended Communities they want to receive in VPN routes from that peer. But RT-constrain has distinct advantages over Extended Community-based ORF: it is more widely supported, it is simpler to configure, and its distribution scope is not limited to a direct peer.

In SR OS the capability to exchange RTC routes is advertised when the **route-target** keyword is added to the relevant **family** command. RT-constrain is supported on EBGp and IBGP sessions of the base router instance. On any particular session either ORF or RT-constrain may be used but not both; if RT-constrain is configured the ORF capability is not announced to the peer.

When RT-constrain has been negotiated with one or more peers SR OS automatically originates and advertises to these peers one /96 RTC route (the origin AS and Route Target Extended Community are fully specified) for every route target imported by a locally-configured VPRN or BGP-based L2 VPN; this includes MVPN-specific route targets.



SR OS also supports a group/neighbor level **default-route-target** command that causes routers to generate and send a 0:0:0/0 default RTC route to one or more peers. Sending the default RTC route to a peer conveys a request to receive all VPN routes from that peer. The **default-route-target** command is typically configured on sessions that a route reflector has with its PE clients. A received default RTC route is never propagated to other routers.

The advertisement of RTC routes by a route reflector follows special rules that are described in RFC 4684. These rules are needed to ensure that RTC routes for the same NLRI that are originated by different PE routers in the same Autonomous System are properly distributed within the AS.

When a BGP session comes up, and RT-constrain is enabled on the session (both peers advertised the MP-BGP capability), routers delay sending any VPN-IPv4 and VPN-IPv6 routes until either the session has been up for 60 seconds or the End-of-RIB marker is received for the RT-constrain address family. When the VPN-IPv4 and VPN-IPv6 routes are sent they are filtered to include only those with a Route Target Extended Community that matches an RTC route from the peer. VPN-IP routes matching an RTC route originated in the local AS are advertised to any IBGP peer that advertises a valid path for the RTC NLRI — i.e. route distribution is not constrained to only the IBGP peer advertising the best path. On the other hand VPN-IP routes matching an RTC route originated outside the local AS are only advertised to the EBGP or IBGP peer that advertises the best path.



**Note:** SR OS does not support an equivalent of *BGP-Multipath* for RT-Constrain routes. There is no way to distribute VPN routes across more than one ‘almost’ equal set of inter-AS paths.

#### 5.6.3.4 Min Route Advertisement Interval (MRAI)

According to the BGP standard (RFC 4271), a BGP router should not send updated reachability information for an NLRI to a BGP peer until a certain period of time (*Min Route Advertisement Interval*) has elapsed since the last update. The RFC suggests the MRAI should be configurable per peer but does not propose a specific algorithm, and therefore, MRAI implementation details vary from one router operating system to another.

In SR OS, the MRAI is configurable, on a per-session basis, using the **min-route-advertisement** command. The **min-route-advertisement** command can be configured with any value between 1 and 255 seconds and the setting applies to all address families. The default value is 30 seconds, regardless of the session type (EBGP or IBGP). The MRAI timer is started at the configured value when the session is established and counts down continuously, resetting to the configured value whenever it reaches zero. Every time it reaches zero, all pending RIB-OUT routes are sent to the peer.

To send UPDATE messages that advertise new NLRI reachability information more frequently for some address families than others, SR OS offers a **rapid-update** command that overrides the remaining time on a peer's MRAI timer and immediately sends routes belonging to specified address families (and all other pending updates) to the peers receiving these routes. The address families that can be configured with **rapid-update** support are:

- EVPN
- L2-VPN
- Label-IPv4
- Label-IPv6
- MCAST-VPN-IPv4
- MCAST-VPN-IPv6
- MDT-SAFI
- MVPN-IPv4
- MVPN-IPv6
- VPN-IPv4
- VPN-IPv6

In many cases, the default MRAI is appropriate for all address families (or at least those not included in the preceding list) when it applies to UPDATE messages that advertise reachable NLRI, but it is not the best option for UPDATE messages that advertise unreachable NLRI (route withdrawals). Fast re-convergence after some types of failures requires route withdrawals to propagate to other routers as quickly as possible so that they can calculate and start using new best paths, which would be impeded by the effect of the MRAI timer at each router hop. This is facilitated by the **rapid-withdrawal** configuration command.

When **rapid-withdrawal** is configured, UPDATE messages containing withdrawn NLRI are sent immediately to a peer without waiting for the MRAI timer to expire. UPDATE messages containing reachable NLRI continue to wait for the MRAI timer to expire, or for a **rapid-update** trigger, if it applies. When **rapid-withdrawal** is enabled, it applies to all address families.

---

### 5.6.3.5 Advertise-Inactive

BGP does not allow a route to be advertised unless it is the best path in the RIB and an export policy allows the advertisement.

In some cases it may be useful to advertise the best BGP path to peers despite the fact that is *inactive* —for example, because there are one or more preferred non-BGP routes to the same destination and one of these other routes is the active route. One way SR OS supports this flexibility is using the **advertise-inactive** command; other methods include *Best-External* and *Add-Paths*.

When the BGP **advertise-inactive** command is configured so that it applies to a BGP session it has the following effect on the IPv4, IPv6, mcast-ipv4, mcast-ipv6, label-IPv4 and label-IPv6 routes advertised to that peer:

- If the active route for the IP prefix is a BGP route then that route is advertised.
- If the active route for the IP prefix is a non-BGP route and there is at least one valid but inactive BGP route for the same destination then the best of the inactive and valid BGP routes is advertised unless the non-BGP active route is matched and accepted by an export policy applied to the session.
- If the active route for the IP prefix is a non-BGP route and there are no (valid) BGP routes for the same destination then no route is advertised for the prefix unless the non-BGP active route is matched and accepted by an export policy applied to the session.

### 5.6.3.6 Best-External

*Best-External* is a BGP enhancement that allows a BGP speaker to advertise to its IBGP peers its best “external” route for a prefix/NLRI when its best overall route for the prefix/NLRI is an “internal” route. This is not possible in a normal BGP configuration because the base BGP specification prevents a BGP speaker from advertising a non-best route for a destination.

In certain topologies *Best-External* can improve convergence times, reduce route oscillation and allow better loadsharing. This is achieved because routers internal to the AS have knowledge of more exit paths from the AS. Enabling *Add-Paths* on border routers of the AS can achieve a similar result but *Add-Paths* introduces NLRI format changes that must be supported by BGP peers of the border router and therefore has more interoperability constraints than *Best-External* (which requires no messaging changes).

*Best-External* is supported in the base router BGP context. (A related feature is also supported in VPRNs; consult the Services Guide for more details.) It is configured using the **advertise-external** command, which provides IPv4, label-IPv4, IPv6, and label-IPv6 as options.

The advertisement rules when **advertise-external** is enabled can be summarized as follows:

- If a router has **advertise-external** enabled and its best overall route is a route from an IBGP peer then this best route is advertised to EBGp and confed-EBGP peers, and the “best external” route is advertised to IBGP peers. The “best external” route is the one found by running the BGP path selection algorithm on all LOC-RIB paths except for those learned from the IBGP peers.



**Note:** A route reflector with **advertise-external** enabled does not include IBGP routes learned from other clusters in its definition of ‘external’.

- If a router has **advertise-external** enabled and its best overall route is a route from an EBGp peer then this best route is advertised to EBGp, confed-EBGP, and IBGP peers.
- If a router has **advertise-external** enabled and its best overall route is a route from a confed-EBGP peer in member AS X then this best route is advertised to EBGp, IBGP peers and confed-EBGP peers in all member AS except X and the “best external” route is advertised to confed-EBGP peers in member AS X. In this case the “best external” route is the one found by running the BGP path selection algorithm on all RIB-IN paths except for those learned from member AS X.



**Note:** If the best-external route is not the best overall route it is not installed in the forwarding table and in some cases this can lead to a short-duration traffic loop after failure of the overall best path.


### 5.6.3.7 Add-Paths

Add-Paths is a BGP enhancement that allows a BGP router to advertise multiple distinct paths for the same prefix/NLRI. Add-Paths provides a number of potential benefits, including reduced routing churn, faster convergence, and better loadsharing.

For a router to receive multiple paths per NLRI from a peer, for a particular address family, the peer must announce the BGP capability to send multiple paths for the address family and the local router must announce the BGP capability to receive multiple paths for the address family. When the Add-Path capability has been negotiated this way, all advertisements and withdrawals of NLRI by the peer must include a path identifier. The path identifier has no significance to the receiving router. If the combination of NLRI and path identifier in an advertisement from a peer is unique (does not match an existing route in the RIB-IN from that peer) then the route is added to the RIB-IN. If the combination of NLRI and path identifier in a received advertisement is the same as an existing route in the RIB-IN from the peer then the new route replaces the existing one. If the combination of NLRI and path identifier in a received withdrawal matches an existing route in the RIB-IN from the peer, then that route is removed from the RIB-IN.

An UPDATE message carrying an IPv4 NLRI with a path identifier is shown in [Figure 33](#).

**Figure 33** BGP Update Message with Path Identifier for IPv4 NLRI



*Add-Paths* is only supported by the base router BGP instance and the EBGp and IBGP sessions it forms with other peers capable of *Add-Paths*. The ability to send and receive multiple paths per prefix is configurable per family, with the supported options being:

- IPv4
- Label-IPv4
- VPN-IPv4
- IPv6
- Label-IPv6
- VPN-IPv6
- MCAST-VPN-IPv4

- MCAST-VPN-IPv6
- MVPN-IPv4
- MVPN-IPv6

#### 5.6.3.7.1 Path Selection with Add-Paths

The LOC-RIB may have multiple paths for a prefix. The path selection mode refers to the algorithm used to decide which of these paths to advertise to an Add-Paths peer. SR OS supports the Add-N path selection algorithm described in *draft-ietf-idr-add-paths-guidelines*. The Add-N algorithm selects, as candidates for advertisement, the N best paths with unique BGP next-hops. In the SR OS implementation, the default value of N is configurable, per address-family, at the BGP instance, group and neighbor levels, however, this default value can be overridden, for specific prefixes, using route policies. The maximum number of paths to advertise for a prefix to an Add-Paths neighbor is the value N assigned by a BGP import policy to the best path for P, otherwise it defaults to the neighbor, group or instance level configuration of N for the address family to which P belongs.

Add-Paths allows non-best paths to be advertised to a peer, but it still complies with basic BGP advertisement rules such as the IBGP split horizon rule: a route learned from an IBGP neighbor cannot be re-advertised to another IBGP neighbor unless the router is configured as a route reflector.

#### 5.6.3.8 Split-Horizon

Split-horizon refers to the action taken by a router to avoid advertising a route back to the peer from which it was received. By default SR OS applies split-horizon behavior only to routes received from IBGP non-client peers, and split-horizon only works for routes to non-imported routes within a RIB. This split-horizon functionality, which can never be disabled, prevents a route learned from a non-client IBGP peer to be advertised to the sending peer or any other non-client peer.

To apply split-horizon behavior to routes learned from RR clients, confed-EBGP peers or (non-confed) EBGp peers the **split-horizon** command must be configured in the appropriate contexts; it is supported at the global BGP, **group** and **neighbor** levels. When **split-horizon** is enabled on these types of sessions it only prevents the advertisement of a route back to its originating peer; for example SR OS does not prevent the advertisement of a route learned from one EBGp peer back to different EBGp peer in the same neighbor AS.

## 5.7 BGP Monitoring Protocol

The BGP Monitoring Protocol (BMP) provides a monitoring station that obtains route updates and statistics from a BGP router. The BMP protocol is described in detail in RFC 7854, *BGP Monitoring Protocol (BMP)*. A router communicates information about one or more BGP sessions to a BMP station. Specifically, BMP allows a BGP router to advertise the pre-policy or post-policy BGP RIB-In from specific BGP peers to a monitoring station. This allows the monitoring station to monitor the routing table size, identify issues, and monitor trends in the table size and update or withdraw the frequency. The BMP station is also sometimes called a BMP collector. A router sends information in BMP messages to a BMP station.

BMP is a uni-directional protocol. A BMP station never sends back any messages to a router.

BMP allows a router to report different types of information.

- A router can send BMP messages with notifications when neighbors go into or out of the established mode (for example, when the peer goes "up" or "down"). These notifications are called BMP peer-up and peer-down messages.
- A router can periodically send statistical information about one or more neighbors. This information consists of a number of counters, for example, the number of routes received from a particular neighbor or the number of rejected or accepted routes due to ingress policy parameters.  
Other counters report the number of errors that were encountered, for example, AS-path loops, duplicate prefixes, or withdrawals received.
- A router can also report the exact routes received from a particular neighbor. This action is called route monitoring. A router encapsulates a BGP route into the original BGP update message, then encapsulates that BGP update message in a BMP route monitoring message.

BMP on an SR OS router reports information about routes that were received from a neighbor. The SR OS cannot report routes that were sent to a neighbor.

When periodic statistics are enabled, the router will send all the statistics as described in RFC 7854, section 4.8, with the exception of statistic number 13, "Number of duplicate update messages received". The supported statistics are listed in [Table 50](#).

**Table 50**      **Supported Statistics**

| Statistic                                     | Type |
|-----------------------------------------------|------|
| Number of Prefixes rejected by inbound policy | 0    |

**Table 50**      **Supported Statistics (Continued)**

| Statistic                                                         | Type |
|-------------------------------------------------------------------|------|
| Number of duplicate prefix advertisements received                | 1    |
| Number of duplicate withdraws received                            | 2    |
| Number of invalidated prefixes due to Cluster_List loop detection | 3    |
| Number of invalidated prefixes due to AS_PATH loop detection      | 4    |
| Number of invalidated prefixes due to Originator ID validation    | 5    |
| Number of invalidated prefixes due to AS-Confed loop detection    | 6    |
| Total number of routes in adj-rib-in (all families)               | 7    |
| Total number of routes in Local-RIB (all families)                | 8    |
| Number of routes per address-family in adj-rib-in                 | 9    |
| Number of routes per address-family in loc-rib                    | 10   |
| Number of updates subjected to treat-as-withdraw                  | 11   |
| Number of prefixes subjected to treat-as-withdraw                 | 12   |



**Note:** Statistics 9 and 10 are per address family. The address family is specified as an AFI/SAFI pair. Regardless of which families are configured for route-monitoring, a router will report the statistics of all address families that were negotiated with the neighbor.

The values in these counters are the same values that can be seen in the **show>router>bgp>neighbor ip-address [detail]** command in the CLI.



---

## 5.8 BGP Applications

SR OS implements the following BGP applications:

- [BGP Flow-Spec](#)

### 5.8.1 BGP Flow-Spec

Flow-spec is a standardized method for using BGP to distribute traffic flow specifications (flow routes) throughout a network. A flow route carries a description of a flow in terms of packet header fields such as source IP address, destination IP address, or TCP/UDP port number and indicates (through a community attribute) an action to take on packets matching the flow. The primary application for Flow-spec is DDoS mitigation.

Flow-spec is supported for both IPv4 and IPv6. To exchange IPv4 Flow-spec routes with a BGP peer the **flow-ipv4** keyword must be part of the **family** command that applies to the session and to exchange IPv6 Flow-spec routes with a BGP peer **flow-ipv6** must be present in the **family** configuration.

The NLRI of an IPv4 flow route can contain one or more of the subcomponents shown in [Table 51](#).

**Table 51 Subcomponents of IPv4 Flow Route NLRI**

| Subcomponent Name [Type]    | Value Encoding                        | SR OS Support                                                             |
|-----------------------------|---------------------------------------|---------------------------------------------------------------------------|
| Destination IPv4 Prefix [1] | Prefix length, prefix                 | Yes                                                                       |
| Source IPv4 Prefix [2]      | Prefix length, prefix                 | Yes                                                                       |
| IP Protocol [3]             | One or more (operator, value) pairs   | Partial. No support for multiple values other than "TCP or UDP".          |
| Port [4] <sup>1</sup>       | One or more (operator, value) pairs   | Yes                                                                       |
| Destination Port [5]        | One or more (operator, value) pairs   | Yes                                                                       |
| Source Port [6]             | One or more (operator, value) pairs   | Yes                                                                       |
| ICMP Type [7]               | One or more (operator, value) pairs   | Partial. Only a single value is supported.                                |
| ICMP Code [8]               | One or more (operator, value) pairs   | Partial. Only a single value is supported.                                |
| TCP Flags [9]               | One or more (operator, bitmask) pairs | Partial. Only SYN and ACK flags can be matched.                           |
| Packet Length [10]          | One or more (operator, value) pairs   | Yes                                                                       |
| DSCP [11]                   | One or more (operator, value) pairs   | Partial. Only a single value is supported.                                |
| Fragment [12]               | One or more (operator, bitmask) pairs | Partial. No support for matching DF bit, first-fragment or last-fragment. |

**Note:**

1. The Port [4] subcomponent specifies both source and destination ports.

The NLRI of an IPv6 flow route can contain one or more of the subcomponents shown in [Table 52](#).

**Table 52** Subcomponents of IPv6 Flow Route NLRI

| Subcomponent Name<br>[Type] | Value Encoding                        | SR OS Support                                   |
|-----------------------------|---------------------------------------|-------------------------------------------------|
| Destination IPv6 Prefix [1] | Prefix length, prefix offset, prefix  | Partial. No support for prefix offset.          |
| Source IPv6 Prefix [2]      | Prefix length, prefix offset, prefix  | Partial. No support for prefix offset.          |
| Next Header [3]             | One or more (operator, value) pairs   | Partial. Only a single value supported.         |
| Port [4] <sup>1</sup>       | One or more (operator, value) pairs   | Yes                                             |
| Destination Port [5]        | One or more (operator, value) pairs   | Yes                                             |
| Source Port [6]             | One or more (operator, value) pairs   | Yes                                             |
| ICMP Type [7]               | One or more (operator, value) pairs   | Partial. Only a single value is supported.      |
| ICMP Code [8]               | One or more (operator, value) pairs   | Partial. Only a single value is supported.      |
| TCP Flags [9]               | One or more (operator, bitmask) pairs | Partial. Only SYN and ACK flags can be matched. |
| Packet Length [10]          | One or more (operator, value) pairs   | Yes                                             |
| Traffic Class [11]          | One or more (operator, value) pairs   | Partial. Only a single value is supported.      |
| Fragment [11]               | One or more (operator, bitmask) pairs | Partial. No support for matching Last Fragment. |
| Flow Label [13]             | One or more (operator, value) pairs   | Partial. Only a single value is supported.      |

**Note:**

1. The Port [4] subcomponent specifies both source and destination ports.

[Table 53](#) summarizes the actions that may be associated with an IPv4 or IPv6 flow route and how each type of action is encoded.

**Table 53** IPv4 Flowspec Actions

| Action     | Encoding                       | SR OS Support |
|------------|--------------------------------|---------------|
| Rate Limit | Extended community type 0x8006 | Yes           |

**Table 53 IPv4 Flowspec Actions (Continued)**

| Action             | Encoding                                | SR OS Support                                         |
|--------------------|-----------------------------------------|-------------------------------------------------------|
| Sample/Log         | Extended community type 0x8007<br>S-bit | Yes                                                   |
| Next Entry         | Extended community type 0x8007<br>T-bit | No                                                    |
| Redirect to VRF    | Extended community type 0x8008          | Yes                                                   |
| Mark Traffic Class | Extended community type 0x8009          | Yes                                                   |
| Redirect to IPv4   | Extended community type 0x010c          | Yes                                                   |
| Redirect to LSP    | Extended community type 0x0900          | Partial, only support for ID-type 0x00 (localized ID) |

### 5.8.1.1 Validating Received FlowSpec Routes

Received flow-spec-IPv4 and flow-spec-IPv6 routes are validated following the procedures documented in RFC 5575 and draft-ietf-idr-bgp-flowspec-oid-03. You must configure the **validate-dest-prefix** command in a routing instance for the validation checks based on destination prefix to be applied. By default, no checking is done. When the command is enabled, BGP determines whether a flow-spec route is valid or invalid based on the following logic:

1. If the flow-spec route was originated in the same autonomous system (AS) as the receiving BGP router then it is automatically valid.
2. If rule 1 does not apply and the flow-spec route was originated in an external AS, and it does not contain a destination prefix subcomponent, then it is considered valid.
3. If rule 1 does not apply and the flow-spec route was originated in an external AS, and it does contain a destination prefix subcomponent, then it is considered if all of the following are true.
  - The neighbor AS (last non-confed AS in the AS\_PATH) of the flow-spec route matches the neighbor AS of the unicast IP route that is the best match of the destination prefix. The best match unicast IP route must be a BGP route (that is, not static, IGP, or other routes).
  - The neighbor AS of the flow-spec route matches the neighbor AS of all unicast IP routes that are longer matches of the destination prefix. All longer match unicast IP routes must be BGP routes (that is, not static, IGP, or other routes).

Flowspec-IPv4 routes that are received with a redirect-to-IPv4 extended community action are also be subject to a further set of validation checks. If the **validate-redirect-ip** command is enabled in the receiving BGP instance, then a flowspec-IPv4 route is considered invalid if it is deemed to have originated in a different AS than the IP route that resolves the redirection IPv4 address. The originating AS of a flow-spec route is determined from its AS paths.

A flow-spec route that is determined to be invalid by any of the validation rules described earlier is retained in the BGP RIB, but not used for traffic filtering and not propagated to other BGP speakers.

### 5.8.1.2 Using Flow Routes to Create Dynamic Filter Entries

When the base router BGP instance receives an IPv4 or IPv6 flow route and that route is valid/best, the system attempts to construct an IPv4 or IPv6 filter entry from the NLRI contents and the actions encoded in the UPDATE message. If successful, the filter entry is added to the system-created “fSpec-0” IPv4 embedded filter or to the “fSpec-0” IPv6 embedded filter. These embedded filters can be inserted into configured IPv4 and IPv6 filter policies that are applied to ingress traffic on a selected set of the base router IP interfaces. These interfaces can include network interfaces, IES SAP interfaces, and IES spoke SDP interfaces.

Similarly, filter entries can be added to system-created “fSpec-\$vprnId” embedded filters for use with VPRN interfaces.

When flowspec rules are embedded into a user-defined filter policy, the insertion point of the rules is configurable through the **offset** parameter of the **embed-filter** command. The sum of the *ip-filter-max-size* and *offset* must not exceed the maximum filter *entry-id* range.

## 5.8.2 Configuration of TTL Propagation for BGP Label Routes

This feature allows the separate configuration of TTL propagation for in transit and CPM generated IP packets at the ingress LER within a BGP label route context.

### 5.8.2.1 TTL Propagation for RFC 3107 Label Route at Ingress LER

For IPv4 and IPv6 packets forwarded using a RFC 3107 label route in the global routing instance, including label-IPv6, the following command specified with the **all** value enables TTL propagation from the IP header into all labels in the transport label stack:

- **config router ttl-propagate label-route-local [none | all]**
- **config router ttl-propagate label-route-transit [none | all]**

The **none** value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack.

These commands do not have a **no** version.



**Note:**

- The TTL of the IP packet is always propagated into the RFC 3107 label itself. The commands only control the propagation into the transport labels, for example, the labels of the RSVP or LDP LSP which the BGP label route resolves to and which are pushed on top of the BGP label.
- If the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves:
- RSVP LSP shortcut:
  - **configure router mpls shortcut-transit-ttl-propagate**
  - **configure router mpls shortcut-local-ttl-propagate**
- LDP LSP shortcut:
  - **configure router ldp shortcut-transit-ttl-propagate**
  - **configure router ldp shortcut-local-ttl-propagate**

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for RSVP or LDP LSP shortcut.

### 5.8.2.2 TTL Propagation for RFC 3107 Label Routes at LSR

This feature configures the TTL propagation for transit packets at a router acting as an LSR for a BGP label route.

When an LSR swaps the BGP label for a IPv4 prefix packet, thus acting as a ABR, ASBR, or data-path Route-Reflector (RR) in the base routing instance, or swaps the BGP label for a vpn-IPv4 or vpn-IPv6 prefix packet, thus acting as an inter-AS Option B VPRN ASBR or VPRN data path Route-Reflector (RR), the all value of the following command enables TTL propagation of the decremented TTL of the swapped BGP label into all LDP or RSVP transport labels.

- **config router ttl-propagate lsr-label-route [none | all]**

When an LSR swaps a label or stitches a label, it always writes the decremented TTL value into the outgoing swapped or stitched label. What the above CLI controls is whether this decremented TTL value is also propagated to the transport label stack pushed on top of the swapped or stitched label.

The **none** value reverts to the default mode which disables TTL propagation. This changes the existing default behavior which propagates the TTL to the transport label stack. When a customer upgrades, the new default becomes in effect. The above commands do not have a no version.

The following describes the behavior of LSR TTL propagation in a number of other use cases and indicates if the above CLI command applies or not:

1. When an LSR stitches an LDP label to a BGP label, the decremented TTL of the stitched label is propagated or not to the LDP or RSVP transport labels as per the above configuration.
2. When an LSR stitches a BGP label to an LDP label, the decremented TTL of the stitched label is automatically propagated into the RSVP label if the outgoing LDP LSP is tunneled over RSVP. This behavior is not controlled by the above CLI.
3. When a LSR pops a BGP label and forwards the packet using an IGP route (IGP route to destination of prefix wins over the BGP label route), it pushes an LDP label on the packet and the TTL behavior is like described in (2) when stitching from a BGP label to an LDP label.
4. Carrier Supporting Carrier (CsC) VPRN. The ingress CsC PE swaps the incoming eBGP label into a VPN-IPv4 label. The reverse operation is performed by the egress CsC PE. In both cases, the decremented TTL of the swapped label is propagated or not to the LDP or RSVP transport labels as per the above configuration.
5. SR OS does not support ASBR or data path RR functionality for labeled IPv6 routes in the global routing instance (label-IPv6). As such the CLI command above has no impact on prefix packets forwarded in this context.

---

### 5.8.3 BGP Prefix Origin Validation

BGP prefix origin validation is a solution developed by the IETF SIDR working group for reducing the vulnerability of BGP networks to prefix mis-announcements and certain man-in-the-middle attacks. BGP has traditionally relied on a trust model where it is assumed that when a peer AS originates a route it has the right to announce the associated prefix. BGP prefix origin validation takes extra steps to ensure that the origin AS of a route is valid for the advertised prefix.

7450, 7750, and 7950 routers support BGP prefix origin validation for IPv4 and IPv6 routes received by the base router BGP instance from selected peers. When prefix origin validation is enabled on a session using the **enable-origin-validation** command every received IPv4 and/or IPv6 route received from the peer is checked to determine whether the origin AS is valid for the received prefix. The origin AS is generally the right most AS in the AS\_PATH attribute and indicates the autonomous system that originated the route.

For purposes of determining the origin validation state of received BGP routes, the router maintains an Origin Validation database consisting of static and dynamic entries. Each entry is called a VRP (Validated ROA Payload) and associates a prefix (range) with an origin AS.

Static VRP entries are configured using the **static-entry** command available in the **config>router>origin-validation** context of the base router. In SR OS, a static entry can express that a specific prefix and origin AS combination is either valid or invalid.

Dynamic VRP entries are learned from RPKI local cache servers and express valid origin AS and prefix combinations. The router communicates with RPKI local cache servers using the RPKI-RTR protocol. SR OS supports the RPKI-RTR protocol over TCP/IPv4 or TCP/IPv6 transport; TCP-MD5 and other forms of session security are not supported. 7450, 7750, and 7950 routers can set up an RPKI-RTR session using the base routing table (in-band) or the management router (out-of-band). For more information, refer to the **origin-validation** configuration command and show commands in the *Router Configuration Guide*.

An RPKI local cache server is one element of the larger RPKI system. The RPKI is a distributed database containing cryptographic objects relating to Internet Number resources. Local cache servers are deployed in the service provider network and retrieve digitally signed Route Origin Authorization (ROA) objects from Global RPKI servers. The local cache servers cryptographically validate the ROAs before passing the information along to the routers.

The algorithm used to determine the origin validation states of routes received over a session with **enable-origin-validation** configured uses the following definitions:

- A route is **matched** by a VRP entry if all of the following occurs:



- the prefix bits in the route match the prefix bits in the VRP entry (up to its min prefix length)
- the route prefix length is greater than or equal to the VRP entry min prefix length
- the route prefix length is less than or equal to the VRP entry max prefix length
- the origin AS of the route matches the origin AS of the VRP entry
- A route is **covered** by a VRP entry if all of the follow occurs:
  - the prefix bits in the route match the prefix bits in the VRP entry (up to its min prefix length)
  - the route prefix length is greater than or equal to the VRP entry min prefix length
  - the VRP entry type is static-valid or dynamic

Using the above definitions, the origin validation state of a route is based on the following rules.

1. If a route is matched by at least one VRP entry, and the most specific of these matching entries includes a static-invalid entry then the origin validation state is Invalid (2).
2. If a route is matched by at least one VRP entry, and the most specific of these matching entries does not include a static-invalid entry then the origin validation state is Valid (0).
3. If a route is not matched by any VRP entry, but it is covered by at least one VRP entry then the origin validation state is Invalid (2).
4. If a route is not covered by any VRP entry then the origin validation state is Not-Found (1).

Consider the following example. Suppose the Origin Validation database has the following entries:

*10.1.0.0/16-32, origin AS=5, dynamic*

*10.1.1.0/24-32, origin AS=4, dynamic*

*10.0.0.0/8-32, origin AS=5, static invalid*

*10.1.1.0/24-32, origin AS=4, static invalid*

In this case, the origin validation state of the following routes are as indicated:

10.1.0.0/16 with AS\_PATH {...5}: Valid

10.1.1.0/24 with AS\_PATH {...4}: Invalid

10.2.0.0/16 with AS\_PATH {...5}: Invalid

10.2.0.0/16 with AS\_PATH {...6}: Not-Found

The origin validation state of a route can affect its ranking in the BGP decision process. When **origin-invalid-unusable** is configured, all routes that have an origin validation state of 'Invalid' are considered unusable by the best path selection algorithm, that is, they cannot be used for forwarding and cannot be advertised to peers.

If **origin-invalid-unusable** is not configured then routes with an origin validation state of 'Invalid' are compared to other 'usable' routes for the same prefix according to the BGP decision process.

When **compare-origin-validation-state** is configured a new step is added to the BGP decision process after removal of invalid routes and before the comparison of Local Preference. The new step compares the origin validation state, so that a route with a 'Valid' state is preferred over a route with a 'Not-Found' state, and a route with a 'Not-Found' state is preferred over a route with an 'Invalid' state assuming that these routes are considered 'usable'. The new step is skipped if the **compare-origin-validation-state** command is not configured.

Route policies can be used to attach an Origin Validation State extended community to a route received from an EBGp peer in order to convey its origin validation state to IBGP peers and save them the effort of repeating the Origin Validation database lookup. To add an Origin Validation State extended community encoding the 'Valid' result, the route policy should add a community list that contains a member in the format **ext:4300:0**. To add an Origin Validation State extended community encoding the 'Not-Found' result, the route policy should add a community list that contains a member in the format **ext:4300:1**. To add an Origin Validation State extended community encoding the 'Invalid' result, the route policy should add a community list that contains a member in the format **ext:4300:2**.

## 5.8.4 BGP Route Leaking

It is possible to leak a copy of a BGP route (including all its path attributes) from one routing instance and RIB to another routing instance and RIB of the same type (labeled or unlabeled) in the same router. Leaking is supported from the GRT to a VPRN, from one VPRN to another VPRN, and from a VPRN to the GRT. Any valid BGP route for an IPv4, IPv6, or label-IPv4 prefix can be leaked. A BGP route does not have to be the best path or used for forwarding in the source instance in order to be leaked, but it does have to be valid (that is, the next-hop must be resolved, the AS PATH must not exhibit a loop, etc.).

An IPv4, IPv6, or label-IPv4 BGP route becomes a candidate for leaking to another instance when it is specially marked by a BGP import policy. This special marking is achieved by accepting the route with a `bgp-leak` action in the route policy. Routes that are candidates for leaking to other instances show a leakable flag in the output of various `show router BGP` commands. In order to copy a leakable BGP route received in a source instance S into the BGP RIB of a target instance T the target instance must be configured with a leak-import policy that matches and accepts the leakable route. Different leak-import policies can be specified for each of the following RIBs: IPv4, label-IPv4, and IPv6. Multiple (up to 15) leak-import policies can be chained together for more complex use cases. The leak-import policies are configured under the `rib-management` CLI node.



**Note:** Using a leak-import policy to change the BGP attributes of leaked route (compared to the original source copy) is not supported. The only attribute that can be changed is the RTM preference.

In the target instance leaked BGP routes are compared to other (leaked and non-leaked) BGP routes for the same prefix based on the complete BGP decision process, but leaked routes do not have information about the router ID and peer IP address of the original peer and use all-zero values for these properties.

The BGP next-hop of a leaked BGP route is always resolved in the original (source) routing instance. There is no need to leak resolving routes and tunnels into the target instance. If there is no resolving route/tunnel in the source instance then the unresolved route is not leaked. If the cost to reach the BGP next-hop in the source instance is N then this is next-hop cost used by the BGP decision process in both the source and target instances.

If a target instance has BGP multipath and ECMP enabled and some of the equal-cost best paths for a prefix are leaked routes they can be used along with non-leaked best paths as ECMP next-hops of the route.

If the original (source) routing instance has IBGP multipath and ECMP enabled and the route or tunnel that resolves the BGP next-hop of a leakable route has multiple ECMP next-hops then traffic matching the leaked route in the target instance is load-shared across the ECMP next-hops the same way as traffic matching the original route in the source instance. In this case, the ECMP and IBGP-multipath configurations of the target instance are effectively ignored.

When BGP fast reroute is enabled in a target instance T (for a particular IP prefix) BGP attempts to find a qualifying backup path considering both leaked and non-leaked BGP routes. The backup path criteria are unchanged by this feature – i.e. the backup path is the best path remaining after the primary paths and all paths with the same BGP next-hops as the primary paths have been removed.

A leaked BGP route can be advertised to direct BGP neighbors of the target routing instance. The BGP next-hop of a leaked route is automatically be reset to self whenever it is advertised to a peer of the target instance. Normal route advertisement rules apply, meaning that by default the leaked route is advertised if and only if (in the target instance) it is the overall best path and it is used as the active route to the destination and it is not blocked by the IBGP-to-IBGP split-horizon rule.

A BGP route leaked into a VPRN can be exported from the VPRN as a VPN-IPv4/v6 route if it matches the VRF export policy. Normal VPN export rules apply, meaning that by default the leaked route is exported if and only if (in the VPRN) it is the overall best path and it is used as the active route to the destination.



**Note:** A leaked route cannot be exported as a VPN-IP route and then re-imported into another local VPRN.

## 5.8.5 BGP Optimal Route Reflection

BGP Route Reflectors (RRs) are used in networks to improve network scalability by eliminating or reducing the need for a full-mesh of IBGP sessions. When a BGP RR receives multiple paths for the same IP prefix, it typically selects a single best path to send to all clients. If the RR has multiple nearly-equal best paths and the tie-break is determined by the next-hop cost, the RR advertises the path based on its view of next-hop costs. The advertised route may differ from the path that a client would select if it had visibility of the same set of candidate paths and used its own view of next-hop costs.

Non-optimal advertisements by the RR can be a problem in hot-potato routing designs. Hot-potato routing aims to hand off traffic to the next AS using the closest possible exit point from the local AS. In this context, the closest exit point implies minimum IGP cost to reach the BGP next-hop. SR OS implements the hot-potato routing solution described in *draft-ietf-idr-bgp-optimal-route-reflection*

Optimal Route Reflection (ORR) is supported in the base router BGP instance only. It applies to routes in the following address families: IPv4 unicast, label-IPv4, label-IPv6 (6PE), VPN-IPv4, and VPN-IPv6.



**Note:** For the RR to compare two VPN routes (and therefore for ORR to apply), the routes must contain the same RD and IP prefix information.

ORR locations are created when you configure a **config>router>bgp>orr>location** context. The RR can maintain information for a maximum of 16 ORR locations. A primary IPv4 or primary IPv6 address is required for each location; optionally, you may specify a secondary and tertiary IPv4 and IPv6 addresses for the location. The IP addresses are used to find a node in the network topology that can serve as the root for SPF calculations. The IP addresses must correspond to loopback or system IP addresses of routers that participate in IGP protocols. The secondary and tertiary IP address parameters provide redundancy in case the node selected to be root for the SPF calculations disappears.

The route reflector's TE database, populated with information from local IGP instances or BGP-LS NLRI, is used to compute the SPF cost from each ORR location to IPv4 and IPv6 BGP next-hops in the candidate set of best paths. The use of BGP-LS allows the route reflector to learn IGP topology information for OSPF areas, IS-IS levels, and others in which the route reflector is not a direct participant.

To configure an ORR client, configure the **cluster** command for the BGP session to reference one of the defined ORR locations. The association of a client with an ORR location is not automatic. Choose an ORR location as close as possible to the client you are configuring. The **allow-local-fallback** option of the **cluster** command affects RR behavior when no BGP routes are reachable from the ORR location of the client. When **allow-local-fallback** is configured, the RR is allowed, in this circumstance only, to advertise the best reachable BGP path from its own topology location. If **allow-local-fallback** is not configured and this situation applies, then no route is advertised to the client.



**Note:** ORR is supported with Add-Paths; Add-Paths advertised to an ORR client are based on ORR location.

## 5.8.6 LSP Tagging for BGP Next-Hops or Prefixes and BGP-LU

The tunnels used by the system for resolution of BGP next-hops or prefixes and BGP labeled unicast routes can be constrained using LSP administrative tags. Refer to the “LSP Tagging and Auto-Bind Using Tag Information” section of the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for more information.


## 5.8.7 BGP-LS

BGP-LS is a new BGP address family that is intended to distribute IGP topology information to external servers such as Application Layer Traffic Optimization (ALTO) or Path Computation Engines (PCE) servers. These external traffic engineering databases can then use this information when calculating optimal paths.

BGP-LS provides external ALTO and PCE servers with topology information for a multi-area or multi-level network. Through the use of one or two BGP-LS speakers per area or level, the external ALTO or PCE servers can receive full topology information for the entire network. The BGP-LS information can also be distributed through route reflectors supporting the BGP-LS to minimize the peering requirements.

Figure 34 shows a sample BGP-LS network.

**Figure 34** Sample BGP-LS Network



### 5.8.7.1 Supported BGP-LS components

The following BGP-LS components are currently supported.

---

Protocol-ID:

- IS-IS Level 1
- IS-IS Level 2
- OSPFv2

NLRI Types:

- Node NLRI
- Link NLRI
- IPv4 Topology Prefix NLRI

Node Descriptor TLVs:

- 512 — Autonomous System
- 513 — BGP-LS Identifier
- 514 — OSPF Area-ID
- 515 — IGP Router-ID

Node Attribute TLVs:

- 1024 — Node Flag Bits (O & B bits supported)
- 1028 — IPv4 Router-ID of Local Node (only supported for IS-IS)
- Segment Routing:
  - 1034 — SR Capabilities (only supported for IS-IS)
  - 1035 — SR Algorithm (only supported for IS-IS)

Link Descriptors TLVs:

- 258 — Link Local/Remote Identifiers
- 259 — IPv4 interface address
- 260 — IPv4 neighbor address

Link Attributes TLVs:

- 1028 — IPv4 Router-ID of Local Node (only supported for IS-IS)
- 1088 — Administrative group (color)
- 1089 — Maximum link bandwidth
- 1090 — Max. reservable link bandwidth
- 1091 — Unreserved bandwidth
- 1092 — TE Default Metric

- 1095 — IGP Metric
- 1096 — Shared Risk Link Group
- Segment Routing:
  - 1099 — Adjacency Segment Identifier
  - 1100 — LAN Adjacency Segment Identifier

Prefix Descriptors TLVs:

- 264 — OSPF Route Type (only Intra-Area and Inter-Area)
- 265 — IP Reachability Information

Prefix Attributes TLVs:


- 1152 — IGP Flags (only D flag supported)
- 1155 — Prefix Metric
- Segment Routing:
  - 1158 — Prefix SID
  - 1159 — Range (prefix-SID & sub-TLV Only)
  - 1170 — IGP Prefix Attributes (only supported for OSPF)



## 5.9 BGP Configuration Process Overview

Figure 35 displays the process to provision basic BGP parameters.

**Figure 35 BGP Configuration and Implementation Flow**



Router\_Config\_29

---

## 5.10 Configuration Notes

This section describes BGP configuration caveats.

### 5.10.1 General

- Before BGP can be configured, the router ID and autonomous system should be configured.
- BGP must be added to the router configuration. There are no default BGP instances on a router.

#### 5.10.1.1 BGP Defaults

The following list summarizes the BGP configuration defaults:

- By default, the router is not assigned to an AS.
- A BGP instance is created in the administratively enabled state.
- A BGP group is created in the administratively enabled state.
- A BGP neighbor is created in the administratively enabled state.
- No BGP router ID is specified. If no BGP router ID is specified, BGP uses the router system interface address.
- The router BGP timer defaults are generally the values recommended in IETF drafts and RFCs (see [BGP MIB Notes](#))
- If no *import* route policy statements are specified, then all BGP routes are accepted.
- If no *export* route policy statements specified, then all best and used BGP routes are advertised and non-BGP routes are not advertised.

#### 5.10.1.2 BGP MIB Notes

The router implementation of the RFC 1657 MIB variables listed in [Table 54](#) differs from the IETF MIB specification.

**Table 54 SR OS and IETF MIB Variations**

| MIB Variable                         | Description                                                                                                     | RFC 1657 Allowed Values | SR OS Allowed Values                                                                                                  |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------|
| bgpPeerMinRouteAdvertisementInterval | Time interval in seconds for the MinRouteAdvertisementInterval timer. The suggested value for this timer is 30. | 1 to 65535              | 1 to 255<br>A value of 0 is supported when the rapid-update command is applied to an address family that supports it. |

If SNMP is used to set a value of X to the MIB variable in [Table 54](#), there are three possible results:

**Table 55 MIB Variable with SNMP**

| Condition                                                  | Result                                                                                                                                                                                                    |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| X is within IETF MIB values and X is within SR OS values   | SNMP set operation does not return an error<br>MIB variable set to X                                                                                                                                      |
| X is within IETF MIB values and X is outside SR OS values  | SNMP set operation does not return an error<br>MIB variable set to “nearest” SR OS supported value (e.g., SR OS range is 2 - 255 and X = 65535, MIB variable will be set to 255)<br>Log message generated |
| X is outside IETF MIB values and X is outside SR OS values | SNMP set operation returns an error                                                                                                                                                                       |

When the value set using SNMP is within the IETF allowed values and outside the SR OS values as specified in [Table 54](#) and [Table 55](#), a log message is generated.

The log messages that display are similar to the following log messages:

**Sample Log Message for setting bgpPeerMinRouteAdvertisementInterval to 256**

```
535 2006/11/12 19:40:53 [Snmpd] BGP-4-bgpVariableRangeViolation: Trying to set
bgpPeerMinRouteAdvInt to 256 - valid range is [2-255] - setting to 255
```

**Sample Log Message for setting bgpPeerMinRouteAdvertisementInterval to 1**

```
566 2006/11/12 19:44:41 [Snmpd] BGP-4-bgpVariableRangeViolation: Trying to set
bgpPeerMinRouteAdvInt to 1 - valid range is [2-255] - setting to 2
```



---

## 5.11 Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

### 5.11.1 Configuration Overview

#### 5.11.1.1 Preconfiguration Requirements

Before BGP can be implemented, the following entities must be configured:

- The autonomous system (AS) number for the router.

An AS number is a globally unique value which associates a router to a specific autonomous system. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself. Each router participating in BGP must have an AS number specified.

In order to implement BGP, the AS number must be specified in the `config>router` context.

- Router ID — The router ID is the IP address of the local router. The router ID identifies a packet's origin. The router ID must be a valid host address.

#### 5.11.1.2 BGP Hierarchy

BGP is configured in the `config>router>bgp` context. Three hierarchical levels are included in BGP configurations:

- Global level
- Group level
- Neighbor level

Commands and parameters configured on the global level are inherited to the group and neighbor levels although parameters configured on the group and neighbor levels take precedence over global configurations.

---

### 5.11.1.3 Internal and External BGP Configuration

A BGP system is comprised of ASs which share network reachability information. Network reachability information is shared with adjacent BGP peers. BGP supports two types of routing information exchanges:

- External BGP (EBGP) is used between ASs.

EBGP speakers peer to different ASs and typically share a subnet. In an external group, the next hop is dependent upon the interface shared between the external peer and the specific neighbor. The multihop command must be specified if an EBGP peer is more than one hop away from the local router.

- Internal BGP (IBGP) is used within an AS.

IBGP peers belong to the same AS and typically does not share a subnet. Neighbors do not have to be directly connected to each other. Since IBGP peers are not required to be directly connected, IBGP uses the IGP path (the IP next-hop learned from the IGP) to reach an IBGP peer for its peering connection.

## 5.11.2 Basic BGP Configuration

This section provides information to configure BGP and configuration examples of common configuration tasks. The minimal BGP parameters that need to be configured are:

- An autonomous system number for the router.
- A router ID. If a new or different router ID value is entered in the BGP context, then the new value takes precedence and overwrites the router-level router ID.
- A BGP peer group.
- A BGP neighbor with which to peer.
- A BGP peer-AS that is associated with the above peer.

The BGP configuration commands have three primary configuration levels: **bgp** for global configurations, group **name** for BGP group configuration, and neighbor **ip-address** for BGP neighbor configuration. Within the different levels, many of the configuration commands are repeated. For the repeated commands, the command that is most specific to the neighboring router is in effect, that is, neighbor settings have precedence over group settings which have precedence over BGP global settings.

Following is a sample configuration that includes the above parameters. The other parameters shown below are optional:

```
info
```

```
#-----
echo "IP Configuration"
#-----
...
 autonomous-system 200
 confederation 300 members 200 400 500 600
 router-id 10.10.10.103
#-----
...
#-----
echo "BGP Configuration"
#-----
 bgp
 graceful-restart
 exit
 cluster 0.0.0.100
 export "direct2bgp"
 router-id 10.0.0.12
 group "To_AS_10000"
 connect-retry 20
 hold-time 90
 keepalive 30
 local-preference 100
 remove-private
 peer-as 10000
 neighbor 10.0.0.8
 description "To_Router B - EBGp Peer"
 connect-retry 20
 hold-time 90
 keepalive 30
 local-address 10.0.0.12
 passive
 preference 99
 peer-as 10000
 exit
 exit
 group "To_AS_30000"
 connect-retry 20
 hold-time 90
 keepalive 30
 local-preference 100
 remove-private
 peer-as 30000
 neighbor 10.0.3.10
 description "To_Router C - EBGp Peer"
 connect-retry 20
 hold-time 90
 keepalive 30
 peer-as 30000
 exit
 exit
 group "To_AS_40000"
 connect-retry 20
 hold-time 30
 keepalive 30
 local-preference 100
 peer-as 65206
 neighbor 10.0.0.15
 description "To_Router E - Sub Confederation AS 65205"
```

```
connect-retry 20
hold-time 90
keepalive 30
local-address 10.0.0.12
peer-as 65205
exit
exit
exit
#-----
....
A:ALA-48>config>router#
```

### 5.11.3 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure BGP and provides the CLI commands. In order to enable BGP, one AS must be configured and at least one group must be configured which includes neighbor (system or IP address) and peering information (AS number).

All BGP instances must be explicitly created on each router. Once created, BGP is administratively enabled.

Configuration planning is essential to organize ASs and the SRs within the ASs, and determine the internal and external BGP peering.

To configure a basic autonomous system, perform the following tasks:

- Step 1.** Prepare a plan detailing the autonomous systems, the router belonging to each group, group names, and peering connections.
- Step 2.** Associate each router with an autonomous system number.
- Step 3.** Configure each router with a router ID.
- Step 4.** Associate each router with a peer group name.
- Step 5.** Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
- Step 6.** Specify neighbors.
- Step 7.** Specify the autonomous system number associated with each neighbor.



### 5.11.3.1 Creating an Autonomous System

Before BGP can be configured, the autonomous system must be configured first. In BGP, routing reachability information is exchanged between autonomous systems (ASs). An AS is a group of networks that share routing information. The **autonomous-system** command associates an autonomous system number to the router being configured. The autonomous-system command is configured in the **config>router** context.

Use the following CLI syntax to associate a router to an autonomous system:

**CLI Syntax:**     `config>router# autonomous-system autonomous-system`

The router series supports 4 bytes AS numbers by default. This means autonomous-system can have any value from 1 to 4294967295. The following example displays autonomous system configuration command usage:

**Example:**       `config>router# autonomous-system 100`

The following example displays the autonomous system configuration:

```
ALA-B>config>router# info
#-----
IP Configuration
#-----
 interface "system"
 address 10.10.10.104/32
 exit
 interface "to-103"
 address 10.0.0.104/24
 port 1/1/1
 exit
 autonomous-system 100

#-----
ALA-B>config>router#
```

### 5.11.3.2 Configuring a Router ID

In BGP, routing information is exchanged between autonomous systems. The BGP router ID, expressed like an IPv4 address, uniquely identifies the router. It can be set to be the same as the system interface address.

It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

If a new or different router ID value is entered in the BGP context, then the new router ID value is used instead of the router ID configured on the router level, system interface level, or inherited from the MAC address. The router-level router ID value remains intact. The router ID used by BGP is selected in the following order:

- The routed-id configured under **config>router>bgp**
- The router-id configured under **config>router**
- The system interface IPv4 address
- The last 4 bytes of the system MAC address

When configuring a new router ID outside of the **config>router>bgp** context, BGP is not automatically restarted with the new router ID; the next time BGP is (re) initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the shutdown and no shutdown commands for BGP or restart the entire router. Use the following CLI syntax to configure the router ID for multiple protocols:

**CLI Syntax:**     `config>router# router-id router-id`

The following example displays router ID configuration command usage:

**Example:**       `config>router# router-id 10.10.10.104`

The following example displays the router ID configuration:

```
ALA-B>config>router# info


IP Configuration
#-----
 interface "system"
 address 10.10.10.104/32
 exit
 interface "to-103"
 address 10.0.0.104/24
 port 1/1/1
 exit
 autonomous-system 100
 router-id 10.10.10.104
#-----
...
ALA-B>config>router#
```

### 5.11.3.3 BGP Confederations

Follow these steps to configure a confederation:

- Step 1.** Configure the autonomous system number of the confederation using the confederation command in the **config>router** context.
- Step 2.** Configure the BGP confederation members using the confederation command in the **config>router** context.
- Step 3.** Configure IBGP peering within the (local) sub-confederation.
- Step 4.** Configure one or more confed-EBGP peerings to peers in other neighboring sub-confederations.

**Figure 36 Confederation Network Diagram Example**



OSSG206

The following configuration displays the minimum BGP configuration for routers in sub-confederation AS 65001 outlined in [Figure 36](#).

```
ALA-A
 config router
 autonomous-system 65001
```

```
 confederation 100 members 65001 65002 65003
 bgp
 group confed1
 peer-as 65001
 neighbor 2.2.2.2
 exit
 neighbor 3.3.3.3
 exit
 neighbor 4.4.4.4
 exit
 exit
 group external_confed
 neighbor 5.5.5.5
 peer-as 65002
 exit
 neighbor 9.9.9.9
 peer-as 65003
 exit
 exit
 exit
exit
ALA-D
config router
 autonomous-system 65001
 confederation 100 members 65001 65002 65003
 bgp
 group confed1
 peer-as 65001
 neighbor 1.1.1.1
 exit
 neighbor 2.2.2.2
 exit
 neighbor 3.3.3.3
 exit
 exit
 exit
exit
ROUTER 1
config router
 autonomous-system 65003
 confederation 100 members 65001 65002 65003
 bgp
 group confed1
 peer-as 65001
 neighbor 1.1.1.1
 exit
 neighbor 5.5.5.5
 peer-as 65002
 exit
 exit
 exit
exit
```


---

#### **5.11.3.4 BGP Router Reflectors**

In a standard BGP configuration, all BGP speakers within an AS must have a full BGP mesh to ensure that all externally learned routes are redistributed through the entire AS. IBGP speakers do not re-advertise routes learned from one IBGP peer to another IBGP peer. If a network grows, scaling issues could emerge because of the full mesh configuration requirement. Route reflection circumvents the full mesh requirement but still maintains the full distribution of external routing information within an AS.

Autonomous systems using route reflection arrange BGP routers into groups called clusters. Each cluster contains at least one route reflector which is responsible for redistributing route updates to all clients. Route reflector clients do not need to maintain a full peering mesh between each other. They only require a peering to the route reflector(s) in their cluster. The route reflectors must maintain a full peering mesh between all non-clients within the AS.

Each route reflector must be assigned a cluster ID and specify which neighbors are clients and which are non-clients to determine which neighbors should receive reflected routes and which should be treated as a standard IBGP peer. Additional configuration is not required for the route reflector besides the typical BGP neighbor parameters.

**Figure 37** Route Reflection Network Diagram Example

The following configuration displays the minimum BGP configuration for routers in Cluster 1.1.1.1 outlined in [Figure 37](#).

```
ALA-A
config router bgp
 group cluster1
 peer-as 100
 cluster 1.1.1.1
 neighbor 2.2.2.2
 exit
 neighbor 3.3.3.3
 exit
 neighbor 4.4.4.4
 exit
 exit
 group RRs
 peer-as 100
 neighbor 5.5.5.5
 exit
 neighbor 9.9.9.9
 exit
 exit
exit
```

```
ALA-B
config router bgp
 group cluster1
```

```

 peer-as 100
 neighbor 1.1.1.1
 exit
 exit
exit

ALA-C
config router bgp
 group cluster1
 peer-as 100
 neighbor 1.1.1.1
 exit
 exit
exit

ALA-D
config router bgp
 group cluster1
 peer-as 100
 neighbor 1.1.1.1
 exit
 exit
exit

```

### 5.11.3.5 BGP Components

Use the CLI syntax displayed in the following subsections to configure BGP attributes.

#### 5.11.3.5.1 Configuring Group Attributes

A group is a collection of related BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

The following example displays the BGP group configuration:

```

ALA-B>config>router>bgp# info

...
 group "headquarters1"
 description "HQ execs"
 local-address 10.0.0.104
 disable-communities standard extended
 ttl-security 255
 exit

```

```
exit
...

ALA-B>config>router>bgp#
```

### 5.11.3.5.2 Configuring Neighbor Attributes

After you create a group name and assign options, add neighbors within the same autonomous system to create IBGP connections and/or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

The following example displays neighbors configured in group “headquarters1”.

```
ALA-B>config>router>bgp# info

...
 group "headquarters1"
 description "HQ execs"
 local-address 10.0.0.104
 disable-communities standard extended
 ttl-security 255
 neighbor 10.0.0.5
 passive
 peer-as 300
 exit
 neighbor 10.0.0.106
 peer-as 100
 exit
 neighbor 17.5.0.2
 hold-time 90
 keepalive 30
 local-preference 170
 peer-as 10701
 exit
 neighbor 17.5.1.2
 hold-time 90
 keepalive 30
 local-preference 100
 min-route-advertisement 30
 preference 170
 peer-as 10702
 exit
 exit
 ...

ALA-B>config>router>bgp#
```



### 5.11.3.5.3 Configuring Route Reflection

Route reflection can be implemented in autonomous systems with a large internal BGP mesh to reduce the number of IBGP sessions required. One or more routers can be selected to act as focal points for internal BGP sessions. Several BGP speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors. A router assumes the role as a route reflector by configuring the **cluster cluster-id** command. No other command is required unless you want to disable reflection to specific peers.

If you configure the **cluster** command at the global level, then all subordinate groups and neighbors are members of the cluster. The route reflector cluster ID is expressed in dotted decimal notation. The ID should be a significant topology-specific value. No other command is required unless you want to disable reflection to specific peers.

If a route reflector client is fully meshed, the **disable-client-reflect** command can be enabled to stop the route reflector from reflecting redundant route updates to a client.

The following example displays a route reflection configuration:

```
ALA-B>config>router>bgp# info

 cluster 0.0.0.100
 group "Santa Clara"
 local-address 10.0.0.103
 neighbor 10.0.0.91
 peer-as 100
 exit
 neighbor 10.0.0.92
 peer-as 100
 exit
 neighbor 10.0.0.93
 disable-client-reflect
 peer-as 100
 exit
 exit

ALA-B>config>router>bgp#
```

### 5.11.3.5.4 Configuring a Confederation

Reducing a complicated IBGP mesh can be accomplished by dividing a large autonomous system into smaller autonomous systems. The smaller ASs can be grouped into a confederation. A confederation looks like a single AS to routers outside the confederation. Each confederation is identified by its own (confederation) AS number.

To configure a BGP confederation, you must specify a confederation identifier, an AS number expressed as a decimal integer. The collection of autonomous systems appears as a single autonomous system with the confederation number acting as the “all-inclusive” autonomous system number. Up to 15 members (ASs) can be added to a confederation.

The confederation command is configured in the **config>router** context.

Use the following CLI syntax to configure a confederation:

**CLI Syntax:**     `config>router# confederation confed-as-num members  
                                  member-as-num`

When 4-byte AS number support is not disabled on router, the confederation and any of its members can be assigned an AS number in the range from 1 to 4294967295. The following example displays a confederation configuration command usage:

**Example:**       `config>router># confederation 1000 members 100 200 300`

The following example displays the confederation configuration:

```
ALA-B>config>router# info
#-----
IP Configuration
#-----
 interface "system"
 address 10.10.10.103/32
 exit
 interface "to-104"
 shutdown
 address 10.0.0.103/24
 port 1/1/1
 exit
 autonomous-system 100
 confederation 1000 members 100 200 300
 router-id 10.10.10.103
#-----
ALA-B>config>router#
```


### 5.11.3.5.5 Configuring BMP

Refer to [BGP Monitoring Protocol](#) for more information about the protocol. To enable BMP, follow these steps.

1. Configure a BMP station where BMP information will be sent.
2. Configure one or more neighbors to be monitored by that station.

[Figure 38](#) depicts a BMP configuration example.

**Figure 38** Configuring BMP



### Configuring a BMP Station

BMP stations are configured at a top-level configuration context (**config>bmp**). A BMP station can then monitor BGP in the base router, or in one or more virtual routers for VPRN services. For each instance of BGP, a separate BMP TCP session will be set up between the router and that BMP station. A total of eight different BMP stations can be configured on an SR OS router.

The following command creates a BMP station.

```
config>bmp>station station-name [create]
```

For example, **config>bmp>station *antwerp* [create]**

This creates a BMP station named *antwerp*. The station name must be used when configuring BGP peers to be monitored by this station. The name can also be used in the **show>router>bmp** CLI commands to display associated information.

The next step is to configure how this station can be reached. To do this, the IP address of the Linux station, or container, on which the BMP collector is run is required. Also, the TCP port number on which the BMP station is listening is needed. BMP does not use a well-known port number.

- A network operator can pick any appropriate TCP port number.
- BMP sessions from an SR OS router can run over either TCP/IPv4 or TCP/IPv6.

The following is an IP address and port number configuration for a BMP station:

```
configure bmp
 station antwerp create
 connection
 station-address 1.2.3.4 port 5000
 exit
exit
```

This creates a BMP station that can monitor one or more BGP peers.

Next, assign the BGP peers to be monitored:

- In BGP configuration mode, navigate to the monitor command, **config>router>bgp>monitor**. The **monitor** command can be used at the BGP-instance level the group or neighbor levels.
- Configure up to eight station name that will monitor these peers, **config>router>bgp>monitor>station *name***
- Enable monitoring, **config>router>bgp>monitor>no shutdown**. On SR OS routers, both the BMP protocol and each individual station are in a **shutdown** state by default. To allow BMP to start BMP sessions, BMP and the station must be administratively enabled.

```
configure router bgp
 monitor
 station antwerp
 no shutdown
 exit
exit
```

---

All peers in the BGP instance of the base router are now monitored by station *antwerp*. The router will only send BMP peer-up and peer-down messages to the BMP station. Sending periodic statistics messages, or reporting incoming BGP routes must be explicitly configured.

---

## 5.12 BGP Configuration Management Tasks

This section discusses BGP configuration management tasks.

### 5.12.1 Modifying an AS Number

You can modify an AS number on a router but the new AS number will not be used until the BGP instance is restarted either by administratively disabling or enabling the BGP instance or by rebooting the system with the new configuration.

Since the AS number is defined in the **config>router** context, not in the BGP configuration context, the BGP instance is not aware of the change. Re-examine the plan detailing the autonomous systems, the SRs belonging to each group, group names, and peering connections. Changing an AS number on a router could cause configuration inconsistencies if associated **peer-as** values are not also modified as required. At the group and neighbor levels, BGP will re-establish the peer relationships with all peers in the group with the new AS number.

Use the following CLI syntax to change an autonomous system number:

**CLI Syntax:**     `config>router# autonomous-system autonomous-system`

**CLI Syntax:**     `config>router# bgp  
                  group name  
                  neighbor ip-addr  
                  peer-as asn`

**Example:**       `config>router# autonomous-system 400  
config>router# bgp  
config>router>bgp# group headquarters1  
config>router>bgp>group# neighbor 10.10.10.103  
config>router>bgp>group# peer-as 400  
config>router>bgp>group# exit`

### 5.12.2 Modifying a Confederation Number

Modifying a confederation number will cause BGP to restart automatically. Changes immediately take effect.

### 5.12.3 Modifying the BGP Router ID

Changing the router ID number in the BGP context causes the new value to overwrite the router ID configured on the router level, system interface level, or the value inherited from the MAC address. It triggers an immediate reset of all peering sessions.

**Example:**      `config>router>bgp# router-id 10.0.0.123`

This example displays the BGP configuration with the BGP router ID specified:

```
ALA-B>config>router>bgp# info detail

no shutdown
no description
no always-compare-med
ibgp-multipath
. . .
router-id 10.0.0.123

ALA-B>config>router>bgp#
```

### 5.12.4 Modifying the Router-Level Router ID

Changing the router ID number in the **config>router** context causes the new value to overwrite the router ID derive from the system interface address, or the value inherited from the MAC address.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID or restart the entire router.

Use the following CLI syntax to change a router ID:

**CLI Syntax:**    `config>router# router-id router-id`

**Example:**      `config>router# router-id 10.10.10.104`  
                 `config>router# no shutdown`  
                 `config>router>bgp# shutdown`  
                 `config>router>bgp# no shutdown`

The following example displays the router ID configuration:

```
ALA-A>config>router# info
```

```
#-----
IP Configuration
#-----
 interface "system"
 address 10.10.10.104/32
 exit
 interface "to-103"
 address 10.0.0.104/24
 port 1/1/1
 exit
 autonomous-system 100
 router-id 10.10.10.104

#-----
ALA-B>config>router#
```

## 5.12.5 Deleting a Neighbor

In order to delete a neighbor, you must shut down the neighbor before issuing the `no neighbor ip-addr` command.

Use the following CLI syntax to delete a neighbor:

**CLI Syntax:**

```
config>router# bgp
group name
 no neighbor ip-address
 shutdown
 no peer-as asn
 shutdown
```

**Example:**

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# no neighbor 10.0.0.103
```

The following example displays the “headquarters1” configuration with the neighbor 10.0.0.103 removed.

```
ALA-B>config>router>bgp# info

 group "headquarters1"
description "HQ execs"
 local-address 10.0.0.104
 neighbor 10.0.0.5
 passive
 peer-as 300
 exit
exit
```



---

```
ALA-B>config>router>bgp#
```

## 5.12.6 Deleting Groups

In order to delete a group, the neighbor configurations must be shut down first. After each neighbor is shut down, you must shut down the group before issuing the **no group name** command.

Use the following CLI syntax to shut down a peer and neighbor and then delete a group:

**CLI Syntax:**

```
config>router# bgp
no group name
shutdown
no neighbor ip-address
shutdown
shutdown
```

**Example:**

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.105
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group# shutdown
config>router>bgp>group# exit
config>router>bgp# no group headquarters1
```

If you try to delete the group without shutting down the peer-group, the following message appears:

```
ALA-B>config>router>bgp# no group headquarters1
MINOR: CLI BGP Peer Group should be shutdown before deleted. BGP Peer Group not
deleted.
```



## 5.13 BGP Command Reference

### 5.13.1 Command Hierarchies

- [Global BGP Commands](#)
- [Group BGP Commands](#)
- [Neighbor BGP Commands](#)
- [BGP BMP Commands](#)
- [Other BGP-Related Commands](#)

#### 5.13.1.1 Global BGP Commands

```

config
 — router [router-name]
 — confederation confed-as-num [members as-number [as-number]]
 — no confederation
 — no confederation confederation confed-as-num members as-number [as-number]
 — router-id id-address
 — no router-id
 — [no] bgp
 — [no] add-paths
 — ipv4 send send-limit
 — ipv4 send send-limit receive [none]
 — no ipv4
 — ipv6 send send-limit
 — ipv6 send send-limit receive [none]
 — no ipv6
 — label-ipv4 send send-limit
 — label-ipv4 send send-limit receive [none]
 — no label-ipv4
 — label-ipv6 send send-limit
 — label-ipv6 send send-limit receive [none]
 — no label-ipv6
 — mcast-vpn-ipv4 send send-limit
 — mcast-vpn-ipv4 send send-limit receive [none]
 — no mcast-vpn-ipv4
 — mcast-vpn-ipv6 send send-limit
 — mcast-vpn-ipv6 send send-limit receive [none]
 — no mcast-vpn-ipv6
 — mvpn-ipv4 send send-limit
 — mvpn-ipv4 send send-limit receive [none]
 — no mvpn-ipv4
 — mvpn-ipv6 send send-limit
 — mvpn-ipv6 send send-limit receive [none]

```

---

```

— no mvpn-ipv6
— vpn-ipv4 end send-limit
— vpn-ipv4 send send-limit receive [none]
— no vpn-ipv4
— vpn-ipv6 end send-limit
— vpn-ipv6 send send-limit receive [none]
— no vpn-ipv6
— [no] advertise-external [ipv4] [label-ipv4] [ipv6] [label-ipv6]
— [no] advertise-inactive
— advertise-ipv6-next-hops [vpn-ipv6] [label-ipv6] [vpn-ipv4] [label-ipv4]
— no advertise-ipv6-next-hops
— [no] aggregator-id-zero
— auth-keychain name
— authentication-key [authentication-key | hash-key] [hash | hash2]
— no authentication-key
— [no] backup-path [ipv4] [label-ipv4] [ipv6] [label-ipv6]
— best-path-selection
— always-compare-med {zero | infinity}
— always-compare-med strict-as {zero | infinity}
— no always-compare-med
— as-path-ignore [ipv4] [label-ipv4] [vpn-ipv4] [label-ipv6] [ipv6] [vpn-
 ipv6] [mcast-ipv4] [mvpn-ipv4] [l2-vpn] [mvpn-ipv6] [mcast-ipv6]
— no as-path-ignore
— [no] compare-origin-validation-state
— [no] deterministic-med
— ebgp-ibgp-equal [ipv4] [ipv6] [label-ipv4] [label-ipv6] [vpn-ipv4]
 [vpn-ipv6]
— no ebgp-ibgp-equal
— ignore-nh-metric
— ignore-router-id include-internal family [family]
— [no] ignore-router-id
— [no] origin-invalid-unusable
— [no] bfd-enable
— cluster ip-address [orr-location location-id [allow-local-fallback]]
— no cluster
— connect-retry seconds
— no connect-retry
— damp-peer-oscillations [idle-hold-time initial-wait second-wait max-wait]
 [error-interval minutes]
— [no] damping
— description description-string
— no description
— [no] disable-4byte-asn
— [no] disable-client-reflect
— disable-communities [standard] [extended] [large]
— no disable-communities
— [no] disable-fast-external-failover
— [no] disable-route-table-install
— dynamic-neighbor-limit peers
— no dynamic-neighbor-limit
— [no] enable-inter-as-vpn
— [no] enable-peer-tracking
— [no] enable-rr-vpn-forwarding
— [no] enforce-first-as

```

- **error-handling**
  - [no] **update-fault-tolerance**
- **export** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
- **no export**
- **extended-nh-encoding** [label-ipv4] [vpn-ipv4]
- **no extended-nh-encoding**
- **family** [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [ms-pw] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] [bgp-ls] [mcast-ipv6] [mcast-vpn-ipv6] [sr-policy-ipv4]
- **no family**
- **flowspec**
  - [no] **validate-dest-prefix**
  - [no] **validate-redirect-ip**
- [no] **graceful-restart**
  - [no] **enable-notification**
  - [no] **long-lived**
    - **advertise-stale-to-all-neighbors** [without-no-export | no without-no-export]
    - **no advertise-stale-to-all-neighbors**
    - **advertised-stale-time** *seconds*
    - **no advertised-stale-time**
    - [no] **family** {ipv4 | ipv6 | label-ipv4 | label-ipv6 | vpn-ipv4 | vpn-ipv6 | l2-vpn | route-target | flow-ipv4 | flow-ipv6}
      - **advertised-stale-time** *seconds*
      - **no advertised-stale-time**
      - **helper-override-stale-time** *seconds*
      - **no helper-override-stale-time**
    - **forwarding-bits-set** {all | non-fwd}
    - **no forwarding-bits-set**
    - **helper-override-restart-time** *seconds*
    - **no helper-override-restart-time**
    - **helper-override-stale-time** *seconds*
    - **no helper-override-stale-time**
  - **restart-time** *seconds*
  - **no restart-time**
  - **stale-routes-time** *time*
  - **no stale-routes-time**
- **hold-time** *seconds* [min *seconds*]
- **no hold-time**
- [no] **ibgp-multipath**
- **import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
- **no import**
- [no] **initial-send-delay-zero**
- **keepalive** *seconds*
- **no keepalive**
- **label-preference** *value*
- **no label-preference**
- [no] **link-state-import-enable**
- [no] **link-state-import-enable**
- **local-as** *as-number* [private] [no-prepend-global-as]
- **no local-as**
- **local-preference** *local-preference*
- **no local-preference**

---

```

— loop-detect {drop-peer | discard-route | ignore-loop | off}
— no loop-detect
— med-out {number | igp-cost}
— no med-out
— min-route-advertisement seconds
— no min-route-advertisement
— [no] mp-bgp-keep
— multihop ttl-value
— no multihop
— multipath max-paths [ebgp ebgp-max-paths] [ibgp ibgp-max-paths] [restrict
 {same-neighbor-as | exact-as-path}]
— no multipath
— [no] mvpn-vrf-import-subtype-new
— neighbor-trust [vpn-ipv4] [vpn-ipv6]
— no neighbor-trust
— next-hop-resolution
 — labeled-routes
 — [no] allow-static
 — [no] rr-use-route-table
 — transport-tunnel
 — family {label-ipv4 | label-ipv6 | vpn}
 — [no] enforce-strict-tunnel-tagging
 — resolution {any | filter | disabled}
 — resolution-filter
 — [no] bgp
 — [no] ldp
 — [no] rib-api
 — [no] rsvp
 — [no] sr-isis
 — [no] sr-ospf
 — [no] sr-ospf3
 — [no] sr-policy
 — [no] sr-te
 — [no] udp
 — policy policy-name
 — no policy
 — shortcut-tunnel
 — family {ipv4 | ipv6}
 — [no] enforce-strict-tunnel-tagging
 — [no] disallow-igp
 — resolution {any | filter | disabled}
 — resolution-filter
 — [no] bgp
 — [no] ldp
 — [no] rib-api
 — [no] rsvp
 — [no] sr-isis
 — [no] sr-ospf
 — [no] sr-ospf3
 — [no] sr-policy
 — [no] sr-te
 — [no] use-bgp-routes
— optimal-route-reflection

```

- **location** *location-id* [**primary-ip-address** *ipv4-address*] [**secondary-ip-address** *ipv4-address*] [**tertiary-ip-address** *ipv4-address*]
  - **primary-ip-address** *ipv4-address*
  - **no primary-ip-address**
  - **primary-ipv6-address** *ipv6-address*
  - **no primary-ipv6-address**
  - **secondary-ip-address** *ipv4-address*
  - **no secondary-ip-address**
  - **secondary-ipv6-address** *ipv6-address*
  - **no secondary-ipv6-address**
  - **tertiary-ip-address** *ipv4-address*
  - **no tertiary-ip-address**
  - **tertiary-ipv6-address** *ipv6-address*
  - **no tertiary-ipv6-address**
- **spf-wait** *max-wait* [**initial-wait** *initial-wait*] [**second-wait** *second-wait*]
- **no spf-wait**
- **[no] outbound-route-filtering**
  - **[no] extended-community**
    - **[no] accept-orf**
    - **send-orf** *[comm-id]*
    - **no send-orf** *[comm-id]*
- **[no] override-tunnel-elc**
- **[no] path-mtu-discovery**
- **peer-tracking-policy** *policy-name*
- **no peer-tracking-policy**
- **preference** *preference*
- **no preference**
- **purge-timer** *minutes*
- **no purge-timer**
- **rapid-update** [**l2-vpn**] [**mvpn-ipv4**] [**mdt-safi**] [**mvpn-ipv6**] [**evpn**] [**label-ipv4**] [**label-ipv6**] [**vpn-ipv4**] [**vpn-ipv6**] [**mcast-vpn-ipv4**] [**mcast-vpn-ipv6**]
- **no rapid-update**
- **[no] rapid-withdrawal**
- **[no] remove-private** [**limited**] [**skip-peer-as**]
- **rib-management**
  - **ipv4**
    - **leak-import** *plcy-or-long-expr* [*plcy-or-expr*]
    - **no leak-import**
    - **route-table-import** *policy-name*
    - **no route-table-import**
  - **ipv6**
    - **leak-import** *plcy-or-long-expr* [*plcy-or-expr*]
    - **no leak-import**
    - **route-table-import** *policy-name*
    - **no route-table-import**
  - **label-ipv4**
    - **leak-import** *plcy-or-long-expr* [*plcy-or-expr*]
    - **no leak-import**
    - **route-table-import** *policy-name*
    - **no route-table-import**
  - **label-ipv6**
    - **route-table-import** *policy-name*
    - **no route-table-import**
- **route-target-list** *comm-id* [*comm-id*]

- **no route-target-list** [*comm-id*]
- **router-id** *ip-address*
- **no router-id**
- [no] **shutdown**
- [no] **split-horizon**
- [no] **sr-policy-import**
- [no] **third-party-nexthop**
- [no] **vpn-apply-export**
- [no] **vpn-apply-import**

### 5.13.1.2 Group BGP Commands

```

config
— router [router-name]
 — [no] bgp
 — [no] group name
 — [no] add-paths
 — ipv4 send send-limit
 — ipv4 send send-limit receive [none]
 — no ipv4
 — ipv6 send send-limit
 — ipv6 send send-limit receive [none]
 — no ipv6
 — label-ipv4 send send-limit
 — label-ipv4 send send-limit receive [none]
 — no label-ipv4
 — label-ipv6 send send-limit
 — label-ipv6 send send-limit receive [none]
 — no label-ipv6
 — mcast-vpn-ipv4 send send-limit
 — mcast-vpn-ipv4 send send-limit receive [none]
 — no mcast-vpn-ipv4
 — mcast-vpn-ipv6 send send-limit
 — mcast-vpn-ipv6 send send-limit receive [none]
 — no mcast-vpn-ipv6
 — mvpn-ipv4 send send-limit
 — mvpn-ipv4 send send-limit receive [none]
 — no mvpn-ipv4
 — mvpn-ipv6 send send-limit
 — mvpn-ipv6 send send-limit receive [none]
 — no mvpn-ipv6
 — vpn-ipv4 end send-limit
 — vpn-ipv4 send send-limit receive [none]
 — no vpn-ipv4
 — vpn-ipv6 end send-limit
 — vpn-ipv6 send send-limit receive [none]
 — no vpn-ipv6
 — [no] advertise-inactive
 — advertise-ipv6-next-hops [vpn-ipv6] [label-ipv6] [vpn-ipv4] [label-ipv4]
 — no advertise-ipv6-next-hops

```



- [no] **aggregator-id-zero**
- [no] **algp**
- **auth-keychain** *name*
- **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
- **no authentication-key**
- [no] **bfd-enable**
- **cluster** *cluster-id* **orr-location** *orr-location* [**allow-local-fallback**]
- **cluster** *cluster-id*
- **no cluster**
- **connect-retry** *seconds*
- **no connect-retry**
- **damp-peer-oscillations** [**idle-hold-time** *initial-wait* *second-wait* *max-wait*] [**error-interval** *minutes*]
- [no] **damping**
- **def-recv-evpn-encap** *encap-type*
- **no def-recv-evpn-encap**
- [no] **default-route-target**
- **description** *description-string*
- **no description**
- [no] **disable-4byte-asn**
- [no] **disable-capability-negotiation**
- [no] **disable-client-reflect**
- **disable-communities** [**standard**] [**extended**] [**large**]
- **no disable-communities**
- [no] **disable-fast-external-failover**
- **dynamic-neighbor**
  - [no] **prefix** *ip-prefix* *prefix-length*
- **dynamic-neighbor-limit** *peers*
- **no dynamic-neighbor-limit**
- **enable-origin-validation** [**ipv4**] [**label-ipv4**] [**label-ipv6**] [**ipv6**]
- **no enable-origin-validation**
- [no] **enable-peer-tracking**
- [no] **enforce-first-as**
- **error-handling**
  - [no] **update-fault-tolerance**
- **export** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
- **no export**
- **extended-nh-encoding** [**label-ipv4**] [**vpn-ipv4**]
- **no extended-nh-encoding**
- **family** [**ipv4**] [**label-ipv4**] [**vpn-ipv4**] [**ipv6**] [**label-ipv6**] [**vpn-ipv6**] [**mcast-ipv4**] [**I2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**ms-pw**] [**flow-ipv4**] [**flow-ipv6**] [**route-target**] [**mcast-vpn-ipv4**] [**evpn**] [**bgp-ls**] [**mcast-ipv6**] [**mcast-vpn-ipv6**] [**sr-policy-ipv4**]
- **no family**
- [no] **graceful-restart**
  - [no] **enable-notification**
  - [no] **long-lived**
    - **advertise-stale-to-all-neighbors** [**without-no-export** | **no without-no-export**]
    - **no advertise-stale-to-all-neighbors**
    - **advertised-stale-time** *seconds*
    - **no advertised-stale-time**

- 
- [no] **family** {ipv4 | ipv6 | label-ipv4 | label-ipv6 | vpn-ipv4 | vpn-ipv6 | l2-vpn | route-target | flow-ipv4 | flow-ipv6}
    - **advertised-stale-time** *seconds*
    - **no advertised-stale-time**
    - **helper-override-stale-time** *seconds*
    - **no helper-override-stale-time**
  - **forwarding-bits-set** {all | non-fwd}
  - **no forwarding-bits-set**
  - **helper-override-restart-time** *seconds*
  - [no] **helper-override-restart-time**
  - **helper-override-stale-time** *seconds*
  - [no] **helper-override-stale-time**
  - **restart-time** *seconds*
  - **no restart-time**
  - **stale-routes-time** *time*
  - **no stale-routes-time**
  - **hold-time** *seconds* [min *seconds*]
  - **no hold-time**
  - **import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
  - **no import**
  - [no] **initial-send-delay-zero**
  - **keepalive** *seconds*
  - **no keepalive**
  - **label-preference** *value*
  - **no label-preference**
  - **link-bandwidth**
    - **accept-from-ebgp** *family* [*family*]
    - **no accept-from-ebgp**
    - **add-to-received-ebgp** *family* [*family*]
    - **no add-to-received-ebgp**
    - **aggregate-used-paths** *family* [*family*]
    - **no aggregate-used-paths**
    - **send-to-ebgp** *family* [*family*]
    - **no send-to-ebgp**
  - **local-address** [*ip-int-name* | *ip-address*] *ipv6-address*
  - **no local-address**
  - **local-as** *as-number* [private] [no-prepend-global-as]
  - **no local-as**
  - **local-preference** *local preference*
  - **no local-preference**
  - **loop-detect** {drop-peer | discard-route | ignore-loop | off}
  - **no loop-detect**
  - **med-out** {*number* | *igp-cost*}
  - **no med-out**
  - **min-route-advertisement** *seconds*
  - **no min-route-advertisement**
  - [no] **monitor**
    - **route-monitoring** pre-policy [post-policy]
    - **route-monitoring** post-policy
    - **no route-monitoring**
    - [no] **shutdown**
    - **station** *name* [*name...*(*name*)]
    - **station** all

- **no station**
- **multihop** *ttl-value*
- **no multihop**
- **[no] next-hop-self**
- **next-hop-unchanged** [*label-ipv4*] [*label-ipv6*]
- **no next-hop-unchanged**
- **[no] outbound-route-filtering**
  - **[no] extended-community**
    - **[no] accept-orf**
    - **send-orf** [*comm-id*]
    - **no send-orf** [*comm-id*]
- **[no] passive**
- **[no] path-mtu-discovery**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **prefix-limit** *family limit [threshold percentage] [idle-timeout {minutes | forever} | log-only] [post-import]*
- **no prefix-limit** *family*
- **[no] remove-private** {*limited*} {*skip-peer-as*}
- **[no] shutdown**
- **[no] split-horizon**
- **[no] third-party-nexthop**
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- **type** {*internal* | *external*}
- **no type**
- **[no] vpn-apply-export**
- **[no] vpn-apply-import**

### 5.13.1.3 Neighbor BGP Commands

- ```

config
  — router [router-name]
    — [no] bgp
      — [no] group name
        — [no] neighbor ip-address
          — [no] add-paths
            — ipv4 send send-limit receive [none]
            — ipv4 send send-limit
            — no ipv4
            — ipv6 send send-limit receive [none]
            — ipv6 end send-limit
            — no ipv6
            — label-ipv4 send send-limit receive [none]
            — no label-ipv4
            — label-ipv6 send send-limit receive [none]
            — no label-ipv6
            — mcast-vpn-ipv4 send send-limit receive [none]
            — mcast-vpn-ipv4 send send-limit

```

```

— no mcast-vpn-ipv4
— mcast-vpn-ipv6 send send-limit receive [none]
— mcast-vpn-ipv6 send send-limit
— no mcast-vpn-ipv6
— mvpn-ipv4 send send-limit receive [none]
— mvpn-ipv4 send send-limit
— no mvpn-ipv4
— mvpn-ipv6 send send-limit receive [none]
— mvpn-ipv6 send send-limit
— no mvpn-ipv6
— vpn-ipv4 send send-limit receive [none]
— vpn-ipv4 end send-limit
— no vpn-ipv4
— vpn-ipv6 send send-limit receive [none]
— vpn-ipv6 end send-limit
— no vpn-ipv6
— [no] advertise-inactive
— advertise-ipv6-next-hops [vpn-ipv6] [label-ipv6] [vpn-ipv4]
  [label-ipv4]
— no advertise-ipv6-next-hops
— [no] advertise-ldp-prefix
— [no] aggregator-id-zero
— [no] aigp
— auth-keychain name
— authentication-key [authentication-key | hash-key] [hash |
  hash2]
— no authentication-key
— [no] bfd-enable
— cluster cluster-id [orr-location location-id [allow-local-
  fallback]]
— cluster cluster-id
— no cluster
— connect-retry seconds
— no connect-retry
— [no] damp-peer-oscillations [idle-hold-time initial-wait
  second-wait max-wait] [error-interval minutes]
— [no] damping
— [no] default-route-target
— description description-string
— no description
— [no] disable-4byte-asn
— [no] disable-capability-negotiation
— [no] disable-client-reflect
— disable-communities [standard] [extended] [large]
— no disable-communities
— [no] disable-fast-external-failover
— enable-origin-validation [ipv4] [ipv6] [label-ipv4] [label-ipv6]
— no enable-origin-validation
— [no] enable-peer-tracking
— [no] enforce-first-as
— error-handling
  — [no] update-fault-tolerance
— export plcy-or-long-expr [plcy-or-expr [plcy-or-expr]]
— no export

```

-
- **extended-nh-encoding** [label-ipv4] [vpn-ipv4]
 - **no extended-nh-encoding**
 - **family** [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [ms-pw] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] [bgp-ls] [mcast-ipv6] [mcast-vpn-ipv6] [sr-policy-ipv4]
 - **no family**
 - [no] **graceful-restart**
 - [no] **enable-notification**
 - [no] **long-lived**
 - **advertise-stale-to-all-neighbors** [without-no-export | no without-no-export]
 - **no advertise-stale-to-all-neighbors**
 - **advertised-stale-time** *seconds*
 - **no advertised-stale-time**
 - [no] **family** {ipv4 | ipv6 | label-ipv4 | label-ipv6 | vpn-ipv4 | vpn-ipv6 | l2-vpn | route-target | flow-ipv4 | flow-ipv6}
 - **advertised-stale-time** *seconds*
 - **no advertised-stale-time**
 - **helper-override-stale-time** *seconds*
 - **no helper-override-stale-time**
 - **forwarding-bits-set** {all | non-fwd}
 - **no forwarding-bits-set**
 - **helper-override-restart-time** *seconds*
 - **no helper-override-restart-time**
 - **helper-override-stale-time** *seconds*
 - **no helper-override-stale-time**
 - **restart-time** *seconds*
 - **no restart-time**
 - **stale-routes-time** *time*
 - **no stale-routes-time**
 - **hold-time** *seconds* [min *seconds*]
 - **no hold-time**
 - **import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
 - **no import**
 - [no] **initial-send-delay-zero**
 - **keepalive** *seconds*
 - **no keepalive**
 - **label-preference** *value*
 - **no label-preference**
 - **link-bandwidth**
 - **accept-from-ebgp** *family* [*family*]
 - **no accept-from-ebgp**
 - **add-to-received-ebgp** *family* [*family*]
 - **no add-to-received-ebgp**
 - **aggregate-used-paths** *family* [*family*]
 - **no aggregate-used-paths**
 - **send-to-ebgp** *family* [*family*]
 - **no send-to-ebgp**
 - **local-address** {*ip-address* | *ip-int-name* | *ipv6-address*}
 - **no local-address**
 - **local-as** *as-number* [private] [no-prepend-global-as]

```

— no local-as
— local-preference local-preference
— no local-preference
— loop-detect {drop-peer | discard-route | ignore-loop | off}
— no loop-detect
— med-out {number | igp-cost}
— no med-out
— min-route-advertisement seconds
— no min-route-advertisement
— [no] monitor
  — route-monitoring pre-policy [post-policy]
  — route-monitoring post-policy
  — no route-monitoring
  — [no] shutdown
  — station name [name...(name)]
  — station all
  — no station
— multihop ttl-value
— no multihop
— [no] next-hop-self
— next-hop-unchanged [label-ipv4] [label-ipv6]
— no next-hop-unchanged
— [no] outbound-route-filtering
  — [no] extended-community
    — [no] accept-orf
      — send-orf [comm-id]
      — no send-orf [comm-id]
— [no] passive
— [no] path-mtu-discovery
— peer-as as-number
— no peer-as
— preference preference
— no preference
— prefix-limit family limit [log-only] [threshold percentage] [idle-
  timeout {minutes | forever}] [post-import]
— no prefix-limit family
— [no] remove-private {limited} {skip-peer-as}
— [no] shutdown
— [no] split-horizon
— [no] third-party-nexthop
— ttl-security min-ttl-value
— no ttl-security
— type {internal | external}
— no type
— [no] vpn-apply-export
— [no] vpn-apply-import

```

5.13.1.4 BGP BMP Commands

```

config
  — bmp
    — [no] shutdown
    — station station-name [create]
    — no station station-name
      — connection
        — connect-retry seconds
        — no connect-retry
        — local-address {ip-address | ipv6-address}
        — no local-address
        — router service-name service-name
        — router router-instance
        — no router
        — station-address {ip-address | ipv6-address port port}
        — no station-address
        — tcp-keepalive
          — keep-count count
          — no keep-count
          — keep-idle idle
          — no keep-idle
          — keep-interval interval
          — no keep-interval
          — [no] shutdown
        — description description-string
        — no description
        — family [ipv4] [ipv6] [label-ipv4] [label-ipv6] [vpn-ipv4] [vpn-ipv6]
        — no family
        — initiation-message [initiation-message]
        — no initiation-message
        — [no] report-local-routes
        — [no] shutdown
        — stats-report-interval [seconds]
        — no stats-report-interval

config
  — router
    — bgp
      — [no] monitor
        — route-monitoring pre-policy [post-policy]
        — route-monitoring post-policy
        — no route-monitoring
        — [no] shutdown
        — station name [name]
        — station all
        — no station

```

For information about VPRN commands and syntax, refer to 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN. The VPRN BGP BMP commands are listed below.

```

config
  — service
    — vprn service-id [name name] [customer customer-id] [create]
    — no vprn service-id
      — bgp
        — group name [esm-dynamic-peer]
        — no group name
          — [no] monitor
            — route-monitoring pre-policy [post-policy]
            — route-monitoring post-policy
            — no route-monitoring
            — [no] shutdown
            — station name [name]
            — station all
            — no station
          — [no] neighbor ip-address
            — [no] monitor
              — route-monitoring pre-policy [post-policy]
              — route-monitoring post-policy
              — no route-monitoring
              — [no] shutdown
              — station name [name]
              — station all
              — no station
        — [no] monitor
          — route-monitoring pre-policy [post-policy]
          — route-monitoring post-policy
          — no route-monitoring
          — [no] shutdown
          — station name [name]
          — station all
          — no station

```

5.13.1.5 Other BGP-Related Commands

```

config
  — router [router-name]
    — autonomous-system as-number
    — no autonomous-system
    — router-id ip-address
    — no router-id

```


5.13.2 Command Descriptions

5.13.2.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor config>bmp>station
Description	This command creates a text description stored in the configuration file for a configuration context. The no form of the command removes the description string from the context.
Default	no description
Parameters	<i>string</i> — Specifies the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax	[no] shutdown
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor config>router>bgp>monitor config>bmp config>bmp>station config>bmp>station>connection>tcp-keepalive config>router>bgp>monitor config>router>bgp>group>monitor config>router>bgp>group>neighbor>monitor config>service>vprn>bgp>monitor config>service>vprn>bgp>group>monitor config>service>vprn>bgp>group>neighbor>monitor

Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted. The no form of this command administratively enables an entity. Unlike other commands and parameters where the default state is not indicated in the configuration file, the shutdown and no shutdown states are always indicated in system generated configuration files. Default administrative states for services and service entities are described in Special Cases. The no form of the command places an entity in an administratively enabled state.
Special Cases	BGP Global — The BGP protocol is created in the no shutdown state. BGP Group — BGP groups are created in the no shutdown state. BGP Neighbor — BGP neighbors/peers are created in the no shutdown state.

5.13.2.2 BGP Commands

bgp

Syntax	[no] bgp
Context	config>router
Description	This command creates the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation. The no form of the command deletes the BGP protocol instance and removes all configuration parameters for the BGP instance. BGP must be shutdown before deleting the BGP instance. An error occurs if BGP is not shutdown first.

add-paths

Syntax	[no] add-paths
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor

Description This command allows adds the add-paths node to be the configured for one or more families configuration of the BGP instance, a group or a neighbor. The BGP add-paths capability allows the router to send and/or receive multiple paths per prefix to/from a peer. The add-paths command without additional parameters is equivalent to removing Add-Paths support for all address families, which causes sessions that previously negotiated the add-paths capability for one or more address families to go down and come back up without the add-paths capability.

The no form of the command (no add-paths) removes add-paths from the configuration of BGP, the group or the neighbor, causing sessions established using add-paths to go down and come back up without the add-paths capability.

Default no add-paths

ipv4

Syntax **ipv4 send *send-limit* receive [none]**
ipv4 send *send-limit*
no ipv4

Context config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths

Description This command configures the add-paths capability for unlabeled IPv4 unicast routes. By default, add-paths is not enabled for unlabeled IPv4 unicast routes.

The maximum number of unlabeled unicast paths per IPv4 prefix to send is the configured send limit, which is a mandatory parameter. The capability to receive multiple unlabeled IPv4 unicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of the command disables add-paths support for unlabeled IPv4 unicast routes, causing sessions established using add-paths for unlabeled IPv4 unicast to go down and come back up without the add-paths capability.

Default no ipv4

Parameters *send-limit* — Specifies the maximum number of paths per unlabeled IPv4 unicast prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be less.) If the value is none, the router does not negotiate the send capability with respect to IPv4 AFI/SAFI.

Values 1 to 16, none

receive — Specifies that the router negotiates to receive multiple unlabeled unicast routes per IPv4 prefix.

none — Specifies that the router does not negotiate to receive multiple unlabeled unicast routes per IPv4 prefix.

ipv6

Syntax	ipv6 send <i>send-limit</i> receive [none] ipv6 send <i>send-limit</i> no ipv6
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for unlabeled IPv6 unicast routes. By default, add-paths is not enabled for unlabeled IPv6 unicast routes. The maximum number of unlabeled unicast paths per IPv6 prefix to send is the configured send limit, which is a mandatory parameter. The capability to receive multiple unlabeled IPv6 unicast paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command the receive capability is enabled by default. The no form of the command disables add-paths support for unlabeled IPv6 unicast routes, causing sessions established using add-paths for unlabeled IPv6 unicast to go down and come back up without the add-paths capability.
Default	no ipv6
Parameters	send <i>send-limit</i> — Specifies the maximum number of paths per unlabeled IPv6 unicast prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be less.) If the value is none, the router does not negotiate the send capability with respect to IPv6 AFI/SAFI. Values 1 to 16, none receive — Specifies the router negotiates to receive multiple unlabeled unicast routes per IPv6 prefix. none — Specifies the router does not negotiate to receive multiple unlabeled unicast routes per IPv6 prefix.

label-ipv4

Syntax	label-ipv4 send <i>send-limit</i> label-ipv4 send <i>send-limit</i> receive [none] no label-ipv4
Context	config>router>bgp>add-paths

	<pre>config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths</pre>
Description	This command configures the add-paths capability for labeled-unicast IPv4 routes. By default, add-paths is not enabled for labeled-unicast IPv4 routes. The maximum number of labeled-unicast paths per IPv4 prefix to send is the configured <i>send-limit</i>, which is a mandatory parameter. The capability to receive multiple labeled-unicast paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command, receive capability is enabled by default. The no form of the command disables add-paths support for labeled-unicast IPv4 routes, causing sessions established using add-paths for labeled-unicast IPv4 to go down and come back up without the add-paths capability.
Default	no label-ipv4
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per labeled-unicast IPv4 prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be less.) If the value is none, the router does not negotiate the send capability with respect to label-IPv4 AFI/SAFI. Values 1 to 16, none receive — Specifies the router negotiates to receive multiple labeled-unicast routes per IPv4 prefix. none — Specifies that the router does not negotiate to receive multiple labeled-unicast routes per IPv4 prefix.

label-ipv6

Syntax	<pre>label-ipv6 send send-limit label-ipv6 send send-limit receive [none] no label-ipv6</pre>
Context	<pre>config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths</pre>
Description	This command configures the add-paths capability for labeled-unicast IPv6 routes. By default, add-paths is not enabled for labeled-unicast IPv6 routes. The maximum number of labeled-unicast paths per IPv6 prefix to send is the configured <i>send-limit</i>, which is a mandatory parameter. The capability to receive multiple labeled-unicast paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command, receive capability is enabled by default. The no form of the command disables add-paths support for labeled-unicast IPv6 routes, causing sessions established using add-paths for labeled-unicast IPv6 to go down and come back up without the add-paths capability.

Default	no label-ipv6
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per labeled-unicast IPv6 prefix that are allowed to be advertised to add-paths peers. (The actual number of advertised routes may be less.) If the value is none, the router does not negotiate the send capability with respect to label-IPv6 AFI/SAFI. Values 1 to 16, none receive — Specifies that the router negotiates to receive multiple labeled-unicast routes per IPv6 prefix. none — Specifies that the router does not negotiate to receive multiple labeled-unicast routes per IPv6 prefix.

mcast-vpn-ipv4

Syntax	mcast-vpn-ipv4 send <i>send-limit</i> receive [none] mcast-vpn-ipv4 send <i>send-limit</i> no mcast-vpn-ipv4
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for multicast IPv4 VPN routes. By default, add-paths is not enabled for multicast IPv4 VPN routes. The maximum number of multicast paths per IPv4 VPN prefix to send is the configured <i>send-limit</i>, which is a mandatory parameter. The capability to receive multiple multicast paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command, receive capability is enabled by default. The none option disables the receive capability. The no form of the command disables add-paths support for multicast IPv4 VPN routes, causing sessions established using add-paths for multicast IPv4 VPN to go down and come back up without the add-paths capability.
Default	no mcast-vpn-ipv4
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per multicast IPv4 VPN prefix that are allowed to be advertised to add-paths peers. The actual number of advertised routes may be less. If the value is none, the router does not negotiate the send capability with respect to multicast IPv4 VPN AFI/SAFI. Values 1 to 16, none receive — Specifies that the router negotiates to receive multiple multicast routes per IPv4 VPN prefix. none — Specifies that the router does not negotiate to receive multiple multicast routes per IPv4 VPN prefix.

mcast-vpn-ipv6

Syntax	mcast-vpn-ipv6 send <i>send-limit</i> receive [none] mcast-vpn-ipv6 send <i>send-limit</i> no mcast-vpn-ipv6
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for multicast IPv6 VPN routes. By default, add-paths is not enabled for multicast IPv6 VPN routes. The maximum number of multicast paths per IPv6 VPN prefix to send is the configured <i>send-limit</i>, which is a mandatory parameter. The capability to receive multiple multicast paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command, receive capability is enabled by default. The none option disables the receive capability. The no form of the command disables add-paths support for multicast IPv6 VPN routes, causing sessions established using add-paths for multicast IPv6 VPN to go down and come back up without the add-paths capability.
Default	no mcast-vpn-ipv6
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per multicast IPv6 VPN prefix that are allowed to be advertised to add-paths peers. The actual number of advertised routes may be less. If the value is none, the router does not negotiate the send capability with respect to multicast IPv6 VPN AFI/SAFI. Values 1 to 16, none receive — Specifies that the router negotiates to receive multiple multicast routes per IPv6 VPN prefix. none — Specifies that the router does not negotiate to receive multiple multicast routes per IPv6 VPN prefix.

mvpn-ipv4

Syntax	mvpn-ipv4 send <i>send-limit</i> receive [none] mvpn-ipv4 send <i>send-limit</i> no mvpn-ipv4
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for multicast VPN IPv4 routes. By default, add-paths is not enabled for multicast VPN IPv4 routes.

The maximum number of multicast VPN paths per IPv4 prefix to send is the configured *send-limit*, which is a mandatory parameter. The capability to receive multiple multicast paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command, receive capability is enabled by default. The **none** option disables the receive capability.

The **no** form of the command disables add-paths support for multicast VPN IPv4 routes, causing sessions established using add-paths for multicast VPN IPv4 to go down and come back up without the add-paths capability.

Default	no mvpn-ipv4
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per multicast VPN IPv4 prefix that are allowed to be advertised to add-paths peers. The actual number of advertised routes may be less. If the value is none, the router does not negotiate the send capability with respect to multicast VPN IPv4 AFI/SAFI. Default 1 to 16, none receive — Specifies that the router negotiates to receive multiple multicast VPN routes per IPv4 prefix. none — Specifies that the router does not negotiate to receive multiple multicast VPN routes per IPv4 prefix.

mvpn-ipv6

Syntax	mvpn-ipv6 send <i>send-limit</i> receive [none] mvpn-ipv6 send <i>send-limit</i> no mvpn-ipv6
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for multicast VPN IPv6 routes. By default, add-paths is not enabled for multicast VPN IPv6 routes. The maximum number of multicast VPN paths per IPv6 prefix to send is the configured <i>send-limit</i>, which is a mandatory parameter. The capability to receive multiple multicast VPN paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command, receive capability is enabled by default. The none option disables the receive capability. The no form of the command disables add-paths support for multicast VPN IPv6 routes, causing sessions established using add-paths for multicast VPN IPv6 to go down and come back up without the add-paths capability.
Default	no mvpn-ipv6

- Parameters**
- send-limit* — Specifies the maximum number of paths per multicast VPN IPv6 prefix that are allowed to be advertised to add-paths peers. The actual number of advertised routes may be less. If the value is **none**, the router does not negotiate the send capability with respect to multicast VPN IPv6 AFI/SAFI.
 - receive** — Specifies that the router negotiates to receive multiple multicast VPN routes per IPv6 prefix.
 - none** — Specifies that the router does not negotiate to receive multiple multicast VPN routes per IPv6 prefix.

vpn-ipv4

- Syntax** **vpn-ipv4 send** *send-limit* **receive** [**none**]
vpn-ipv4 send *send-limit*
no vpn-ipv4
- Context** config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths
- Description** This command configures the add-paths capability for VPN-IPv4 routes. By default, add-paths is not enabled for VPN-IPv4 routes.
- The maximum number of paths per VPN-IPv4 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.
- The **no** form of the command disables add-paths support for VPN-IPv4 routes, causing sessions established using add-paths for VPN-IPv4 to go down and come back up without the add-paths capability.
- Default** no vpn-ipv4
- Parameters**
- send-limit* — Specifies the maximum number of paths per VPN-IPv4 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).
- Values** 1 to 16, none
- receive** — Specifies that the router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers
 - none** — Specifies that the router does not negotiate the Add-Paths receive capability for VPN-IPv4 routes with its peers.

vpn-ipv6

Syntax	vpn-ipv6 send <i>send-limit</i> receive [none] vpn-ipv6 send <i>send-limit</i> no vpn-ipv6
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	This command configures the add-paths capability for VPN-IPv6 routes. By default, add-paths is not enabled for VPN-IPv6 routes. The maximum number of paths per VPN-IPv6 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the receive keyword, which is optional. If the receive keyword is not included in the command the receive capability is enabled by default. The no form of the command disables add-paths support for VPN-IPv6 routes, causing sessions established using add-paths for VPN-IPv6 to go down and come back up without the add-paths capability.
Default	no vpn-ipv6
Parameters	<i>send-limit</i> — Specifies the maximum number of paths per VPN-IPv6 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules). Values 1 to 16, none receive — Specifies that the router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers none — Specifies that the router does not negotiate the add-paths receive capability for VPN-IPv6 routes with its peers.

advertise-external

Syntax	[no] advertise-external [ipv4] [ipv6] [label-ipv4] [label-ipv6]
Context	config>router>bgp
Description	This command allows BGP to advertise its best external route to a destination even when its best overall route is an internal route. Entering the command (or its no form) with no address family parameters is equivalent to specifying all supported address families. The no form of the command disables Advertise Best External for the BGP family.
Default	no advertise-external

Parameters	ipv4 — Enables or disables the best-external advertisement for unlabeled unicast IPv4 routes.
	ipv6 — Enables or disables the best-external advertisement for unlabeled unicast IPv6 routes.
	label-ipv4 — Enables or disables the best-external advertisement for labeled-unicast IPv4 routes.
	label-ipv6 — Enables or disables the best-external advertisement for labeled-unicast IPv6 routes.

advertise-inactive

Syntax	[no] advertise-inactive
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination. The BGP advertise-inactive command has the following effect on the IPv4, IPv6, multicast IPv4, multicast IPv6, label IPv4 and label IPv6 routes advertised to that peer: • If the active route for the IP prefix is a BGP route, then that route is advertised. • If the active route for the IP prefix is a non-BGP route and there is at least one valid but inactive BGP route for the same destination, then the best of the inactive and valid BGP routes is advertised unless the non-BGP active route is matched and accepted by an export policy applied to the session. • If the active route for the IP prefix is a non-BGP route and there are no valid BGP routes for the same destination, then no route is advertised for the prefix unless the non-BGP active route is matched and accepted by an export policy applied to the session. The no form of the command disables the advertising of inactive BGP routers to other BGP peers.
Default	no advertise-inactive

advertise-ipv6-next-hops

Syntax	advertise-ipv6-next-hops [vpn-ipv6] [label-ipv6] [vpn-ipv4] [label-ipv4] no advertise-ipv6-next-hops

Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command applies to a BGP session established on top of IPv6; transport, BGP routes belonging to the specified families can be advertised with a true IPv6 address when originated or when next-hop-self (configured or automatic) is applied. This command has no effect on routes advertised to IPv4 peers. When this command is not enabled, the following considerations apply: • If a VPN IPv6 or label IPv6 route needs to be originated or advertised with next-hop-self to an IPv6 transport peer the route is advertised with the IPv4 system address as BGP next hop (encoded as an IPv4-mapped IPv6 address). • If a VPN-IPv4 or label IPv4 route needs to be originated or advertised with next-hop-self or if an appropriate extended-NH-encoding capability was not received from the remote peer, the route is advertised with the IPv4 system address as the BGP next hop. • If a VPN IPv4 or label IPv4 route is matched by a BGP export policy entry that tries to change the next-hop to an IPv6 address or if an appropriate extended-NH-encoding capability was not received from the remote peer, the route is handled as though it was rejected by the policy entry.
Default	no advertise-ipv6-next-hops
Parameters	vpn-ipv6 — Allows VPN IPv6 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next hop in cases of route origination or next-hop-self (configured or automatic). label-ipv6 — Allows label IPv6 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next hop in cases of route origination or next-hop-self (configured or automatic). vpn-ipv4 — Allows VPN IPv4 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next hop in cases of route origination or next-hop-self (configured or automatic). It also allows export policies to change the BGP next hop of a VPN IPv4 route to an IPv6 address. All of these cases require the remote peer to advertise the necessary extended NH encoding capability. label-ipv4 — Allows label IPv4 routes to be advertised to IPv6 transport peers with an IPv6 address as the BGP next hop in cases of route origination or next-hop-self (configured or automatic). It also allows export policies to change the BGP next hop of a label IPv4 route to an IPv6 address. All of these cases require the remote peer to advertise the necessary extended NH encoding capability.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>router>bgp

	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures a TCP authentication keychain to use for the session. The keychain allows the rollover of authentication keys during the lifetime of a session.
Default	no auth-keychain
Parameters	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified TCP session or sessions.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the BGP authentication key. Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message based digest. The no form of the command reverts to the default value.
Default	no authentication-key
Parameters	<i>authentication-key</i> — Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). <i>hash-key</i> — Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" "). This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided. hash — Specifies the key is entered in an encrypted form. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

backup-path

Syntax	[no] backup-path [ipv4] [ipv6] [label-ipv4] [label-ipv6]
Context	config>router>bgp
Description	This command enables the computation and use of a backup path for IPv4 and/or IPv6 BGP-learned prefixes belonging to the base router. Multiple paths must be received for a prefix in order to take advantage of this feature. When a prefix has a backup path and its primary paths fail, the affected traffic is rapidly diverted to the backup path without waiting for control plane re-convergence to occur. When many prefixes share the same primary paths, and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes. By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.
Default	no backup-path
Parameters	ipv4 — Enables BGP fast reroute for unlabeled unicast IPv4 routes. ipv6 — Enables BGP fast reroute for unlabeled unicast IPv6 routes. label-ipv4 — Enables BGP fast reroute for labeled-unicast IPv4 routes. label-ipv6 — Enables BGP fast reroute for labeled-unicast IPv6 routes.

best-path-selection

Syntax	best-path-selection
Context	config>router>bgp
Description	This command enables the context to configure path selection parameters.

always-compare-med

Syntax	always-compare-med {zero infinity} no always-compare-med strict-as {zero infinity} no always-compare-med
Context	config>router>bgp>best-path-selection config>service>vprn>bgp>best-path-selection
Description	This command configures the comparison of BGP routes based on the MED attribute. The default behavior of SR-OS (equivalent to the no form of the command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The always-compare-med command without the strict-as keyword allows MED to be compared even if the paths have

a different neighbor AS; in this case, if neither **zero** or **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

Default	no always-compare-med
Parameters	zero — Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred. infinity — Specifies for routes learned without a MED attribute that a value of infinity ($2^{32}-1$) is used in the MED comparison. This in effect makes these routes the least desirable. strict-as — Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

Syntax	as-path-ignore [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [mcast-ipv6] [mvpn-ipv4] [mvpn-ipv6] [l2-vpn] no as-path-ignore
Context	config>router>bgp>best-path-selection
Description	This command configures whether AS path length is considered in the selection of the best BGP route for a prefix. If an address family is listed in this command, then the length of AS paths is not a factor in the route selection process for routes of that address family. The no form of the command removes the parameter from the configuration.
Default	no as-path-ignore
Parameters	ipv4 — Specifies that the AS-path length will be ignored for all unlabeled unicast IPv4 routes. label-ipv4 — Specifies that the AS-path length will be ignored for all labeled-unicast IPv4 routes. vpn-ipv4 — Specifies that the length AS-path will be ignored for all IPv4 VPRN (SAFI 128) routes. ipv6 — Specifies that the AS-path length will be ignored for all unlabeled unicast IPv6 routes. label-ipv6 — Specifies that the AS-path length will be ignored for all labeled-unicast IPv6 routes. vpn-ipv6 — Specifies that the AS-path length will be ignored for all IPv6 VPRN (SAFI 128) routes.

mcast-ipv4 — Specifies that the AS-path length will be ignored for all IPv4 multicast routes.

mcast-ipv6 — Specifies that the AS-path length will be ignored for all IPv6 multicast routes.

mvpn-ipv4 — Specifies that the AS-path length will be ignored for all IPv4 MVPN routes.

mvpn-ipv6 — Specifies that the AS-path length will be ignored for all IPv6 MVPN routes.

l2-vpn — The AS-path length will be ignored for all L2-VPN NLRIs.

compare-origin-validation-state

Syntax	compare-origin-validation-state no compare-origin-validation-state
Context	config>router>bgp>best-path-selection
Description	When this command is configured, a new step is inserted in the BGP decision process after removal of invalid routes and before the comparison of Local Preference. The new step compares the origin validation state so that a BGP route with a 'Valid' state is preferred over a BGP route with a 'Not-Found' state, and a BGP route with a 'Not-Found' state is preferred over a BGP route with an 'Invalid' state assuming that these routes are considered 'usable'. The new step is skipped when no compare-origin-validation-state is configured.
Default	no compare-origin-validation-state

deterministic-med

Syntax	[no] deterministic-med
Context	config>router>bgp>best-path-selection
Description	This command controls how the BGP decision process compares routes on the basis of MED. When deterministic-med is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without deterministic-med , the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.
Default	no deterministic-med

ebgp-ibgp-equal

Syntax **ebgp-ibgp-equal [ipv4] [ipv6] [label-ipv4] [label-ipv6] [vpn-ipv4] [vpn-ipv6]**

no ebgp-ibgp-equal

Context	config>router>bgp>best-path-selection
Description	This command instructs the BGP decision process to ignore the difference between EBGp and IBGP routes in selecting the best path and eligible multipaths (if multipath and ECMP are enabled). The result is a form of EIBGP load-balancing in a multipath scenario. By default (with the no form of the command), the BGP decision process prefers an EBGp learned route over an IBGP learned route. The behavior can be applied selectively to only certain types of routes by specifying one or more address family names in the command. If no families are specified, the command applies to IPv4, IPv6, label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6 routes.
Default	no ebgp-ibgp-equal
Parameters	ipv4 — Specifies that the command should be applied to unlabeled unicast IPv4 routes. ipv6 — Specifies that the command should be applied to unlabeled unicast IPv6 routes. label-ipv4 — Specifies that the command should be applied to labeled unicast IPv4 routes. label-ipv6 — Specifies that the command should be applied to labeled unicast IPv6 routes. vpn-ipv4 — Specifies that the command should be applied to IPv4 VPN routes. vpn-ipv6 — Specifies that the command should be applied to IPv6 VPN routes.

ignore-nh-metric

Syntax	ignore-nh-metric no ignore-nh-metric
Context	config>router>bgp>best-path-selection config>service>vprn config>service>vprn>bgp>best-path-selection
Description	This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers. The no form of the command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default no ignore-nh-metric

ignore-router-id

Syntax **ignore-router-id include-internal** *family* [*family*]
[no] ignore-router-id

Context config>router>bgp>best-path-selection
config>service>vprn>bgp>best-path-selection

Description When the **ignore-router-id** command is present, and the current best path to a destination was learned from EBGp peer X with BGP identifier x and a new path is received from EBGp peer Y with BGP identifier y, the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x.

The **no** form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default no ignore-router-id

Parameters *family* — Specifies up to two internal families to be included in this configuration.

Values mvpn-ipv4, mvpn-ipv6

origin-invalid-unusable

Syntax **[no] origin-invalid-unusable**

Context config>router>bgp>best-path-selection

Description When **origin-invalid-unusable** is configured, all routes that have an origin validation state of 'Invalid' are considered unusable by the best path selection algorithm, meaning they are not used for forwarding and not advertised to BGP peers.

With the default of **no origin-invalid-unusable**, routes with an origin validation state of 'Invalid' are compared to other 'usable' routes for the same prefix according to the BGP decision process.

Default no origin-invalid-unusable

bfd-enable

Syntax **[no] bfd-enable**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

Default no bfd-enable

cluster

Syntax **cluster** *cluster-id* **orr-location** *location-id* [**allow-local-fallback**]
cluster *cluster-id*
no cluster

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the cluster ID for a route reflector server ID and implicitly configures the associated BGP sessions as route reflector clients of the BGP instance. If an ORR location ID is specified with the cluster ID, the clients in that cluster receive routes optimal for that specific location; see *draft-ietf-idr-bgp-optimal-route-reflection* for more information.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives best path from a non-client peer, it sends the route to all clients. When the route reflector receives a best path from a client peer it sends the route to all non-client and all client peers except the originator.

With optimal route reflection, the best path advertised to a client takes location ID into account, which means that if the tie-break for best path (or Add-Paths) comes down to next-hop IGP cost, the IGP costs will be calculated relative to the specified location. In the SR OS implementation, the IGP costs from arbitrary ORR locations are calculated using OSPF/OSPFv3, IS-IS, or BGP-LS information in the TE DB.

The **no** form of the command deletes the cluster ID and effectively disables route reflection for the group.

Default no cluster

Parameters	<i>ip-address</i> — Specifies the route reflector cluster ID is expressed in dot decimal notation.
Values	Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)
	orr-location <i>location-id</i> — Specifies the optimal route reflection location index for this set of route reflector clients.
Values	1 to 16
	allow-local-fallback — Controls the behavior when there are no BGP routes to advertise to the RR clients that are reachable from the perspective of their ORR location. If this option is configured, the RR is allowed (in this circumstance only), to advertise the best reachable BGP path from its own topology location. If this option is not configured and this situation applies, then no route is advertised to the clients.

connect-retry

Syntax	connect-retry <i>seconds</i> no connect-retry
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the BGP connect retry timer value in seconds. When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command used at the global level reverts to the default value. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	connect-retry 120
Parameters	<i>seconds</i> — The BGP Connect Retry timer value in seconds expressed as a decimal integer. Values 1 to 65535

damp-peer-oscillations

Syntax	damp-peer-oscillations [<i>idle-hold-time</i> <i>initial-wait</i> <i>second-wait</i> <i>max-wait</i>] [<i>error-interval</i>]

	<i>minutes]</i>
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, <i>A Border Gateway Protocol 4 (BGP-4)</i>. The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.
Default	no damp-peer-oscillations
Parameters	<i>initial-wait</i> — The amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while. Values 0 to 2048 Default 0 <i>second-wait</i> — A period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time. Values 1 to 2048 Default 5 <i>max-wait</i> — The maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability. Values 1 to 2048 Default 60 <i>minutes</i> — The interval of time, in minutes after a session reset, during which the session must be error-free in order to reset the penalty counter and return to idle-hold-time to initial-wait. Values 0 to 2048 Default 30

damping

Syntax	[no] damping
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor

Description	This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition. The no form of the command used at the global level reverts route damping. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level. When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters: • Half-life: 15 minutes • Max-suppress: 60 minutes • Suppress-threshold: 3000 • Reuse-threshold: 750
Default	no damping

def-recv-evpn-encap

Syntax	def-recv-evpn-encap <i>encap-type</i> no def-recv-evpn-encap
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures how BGP treats a received EVPN route without BGP encapsulation extended community. If no encapsulation is received, BGP will validate the route as MPLS or VXLAN. The no form of the command removes the encapsulation type from the configuration.
Parameters	<i>encap-type</i> — Specifies the default EVPN encapsulation type.
Values	mpls, vxlan

default-route-target

Syntax	[no] default-route-target
Context	config>router>bgp>group config>router>bgp>group>neighbor

Description This command originates the default RTC route (zero prefix length) towards the selected peers.

Default no default-route-target

disable-4byte-asn

Syntax [no] disable-4byte-asn

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peers.

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

Default no disable-4byte-asn

disable-client-reflect

Syntax [no] disable-client-reflect

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the reflection of routes by the route reflector to the clients in a specific group or neighbor.

This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default no disable-client-reflect

disable-communities

Syntax disable-communities [standard] [extended] [large]
no disable-communities

Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures BGP to disable sending standard, extended, or large communities to specific peers. By default, all communities that are attached to a BGP route (any address family) are not stripped from the route when it is advertised to any type of peer: IBGP, EBGP, or confed-EBGP.
Default	no disable-communities
Parameters	standard — Specifies that standard 4-byte communities should be removed. extended — Specifies that 8-byte extended communities (of all types) should be removed. large — Specifies that 12-byte large communities should be removed.

disable-fast-external-failover

Syntax	[no] disable-fast-external-failover
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures BGP fast external failover.
Default	no disable-fast-external-failover

disable-route-table-install

Syntax	[no] disable-route-table-install
Context	config>router>bgp
Description	This command specifies whether to disable the installation of all (labeled and unlabeled) IPv4 and IPv6 BGP routes into RTM (Routing Table Manager) and the FIB (Forwarding Information Base) on the base router instance.
Default	no disable-route-table-install

dynamic-neighbor-limit

Syntax	dynamic-neighbor-limit <i>peers</i> no dynamic-neighbor-limit

Context	config>router>bgp config>router>bgp>group
Description	This command configures the maximum number of dynamic BGP sessions that will be accepted from remote peers associated with the entire BGP instance or a specific peer group. If accepting a new dynamic session would cause either the group limit or the instance limit to be exceeded, then the new session attempt is rejected and a Notification message is sent back to the remote peer. The no form of the command removes the limit on the number of dynamic sessions.
Default	no dynamic-neighbor-limit
Parameters	<i>peers</i> — The maximum number of dynamic BGP sessions Values 1 to 8192

prefix

Syntax	[no] prefix <i>ip-prefix/prefix-length</i>
Context	config>router>bgp>group>dynamic-neighbor
Description	This command configures a prefix from which to accept dynamic BGP sessions; particularly, sessions from source IP addresses not matching any configured neighbor addresses. A dynamic session is associated with the group having the longest match prefix entry for the source IP address of the peer. The group association determines local parameters that apply to the session, including the local AS, the local IP address, the peer AS, the MP-BGP families, the import and export policies, and so on. The no form of the command removes a prefix entry.
Parameters	<i>ip-prefix/prefix-length</i> — Specifies a prefix from which to accept dynamic BGP sessions Values <i>ipv4-prefix</i> — a.b.c.d (host bits must be 0) <i>ipv4-prefix-length</i> — 0 to 32 <i>ipv6-prefix</i> — x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x — [0..FFFF]H d — [0..255]D <i>ipv6-prefix-length</i> — 0 to 128

enable-origin-validation

Syntax	enable-origin-validation [ipv4] [ipv6] [label-ipv4] [label-ipv6] no enable-origin-validation

Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	When the enable-origin-validation command is added to the configuration of a group or neighbor, it causes every inbound IPv4 and/or IPv6 route from that peer to be marked with one of the 3 following origin validation states: • Valid (0) • Not-Found (1) • Invalid (2) By default (when neither the ipv4 or ipv6 option is present in the command) or when both the ipv4 and ipv6 options are specified, all unicast IPv4 (AFI1/SAFI1), label-IPv4 (AFI1/SAFI4), unicast IPv6 (AFI2/SAFI1), and label-IPv6 (AFI2/SAFI4) routes are evaluated to determine their origin validation states. When only the ipv4 or ipv6 option is present, only the corresponding address family routes (unlabeled and labeled) are evaluated. The enable-origin-validation command applies to all types of BGP peers, but as a general rule, it should only be applied to EBGp peers and groups that contain only EBGp peers.
Default	no enable-origin-validation
Parameters	ipv4 — Enables origin validation processing for unlabeled unicast IPv4 routes. ipv6 — Enables origin validation processing for unlabeled unicast IPv6 routes. label-ipv4 — Enables origin validation processing for labeled IPv4 routes. label-ipv6 — Enables origin validation processing for labeled IPv6 routes.

enable-inter-as-vpn

Syntax	[no] enable-inter-as-vpn
Context	config>router>bgp
Description	This command specifies whether VPNs can exchange routes across autonomous system boundaries, providing model B connectivity The no form of the command disallows ASBRs to advertise VPRN routes to their peers in other autonomous systems.
Default	no enable-inter-as-vpn

enable-peer-tracking

Syntax	[no] enable-peer-tracking
Context	config>router>bgp

```
config>router>bgp>group
config>router>bgp>group>neighbor
```

Description This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the holdtimer to expire; therefore, the BGP re-convergence process is accelerated.

The **no** form of the command disables peer tracking.

Default no enable-peer-tracking

enable-rr-vpn-forwarding

Syntax [no] enable-rr-vpn-forwarding

Context config>router>bgp

Description When this command is configured all received VPN-IP routes, regardless of route target, are imported into the dummy VRF, where the BGP next-hops are resolved. The **label-route-transport-tunnel** under **config>router>bgp>next-hop-resolution** determines what types of tunnels are eligible to resolve the next-hops. If a received VPN-IP route from IBGP peer X is resolved and selected as best so that it can be re-advertised to an IBGP peer Y, and the BGP next-hop is modified towards peer Y (by using the next-hop-self command in Y's group or neighbor context or by using a next-hop action in an export policy applied to Y) then BGP allocates a new VPRN service label value for the route, signals that new label value to Y and programs the IOM to do the corresponding label swap operation. The supported combinations of X and Y are outlined below:

- from X (client) to Y (client)
- from X (client) to Y (non-client)
- from X (non-client) to Y (client)

The **no** form of the command causes the re-advertisement of a VPN-IP route between one IBGP peer and another IBGP peer does not cause a new VPRN service label value to be signaled and programmed even if the BGP next-hop is changed through group/neighbor configuration or policy.

Nokia recommends leaving this command disabled for scaling and convergence reasons.

Default no enable-rr-vpn-forwarding

export

Syntax **export** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]
no export [*plcy-or-long-expr*]

Context config>router>bgp

	<pre>config>router>bgp>group config>router>bgp>group>neighbor</pre>
Description	This command specifies route policies that control the handling of outbound routes transmitted to certain peers. Route policies are configured in the config>router>policy-options context. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used. The export command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route. Only one of the 15 objects referenced by the command can be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters; the remaining 14 objects have a maximum length of 64 characters each. When multiple export commands are issued, the last command entered overrides the previous command. When an export policy is not specified, BGP-learned routes are advertised by default and non-BGP routes are not advertised. The no form of the command removes the policy association.
Default	no export
Parameters	<i>plcy-or-long-expr</i> — Specifies the route policy name (up to 64 characters) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. <i>plcy-or-expr</i> — Specifies up to 14 route policy names (up to 64 characters each) or a policy logical expression (up to 64 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

extended-nh-encoding

Syntax	<pre>extended-nh-encoding [label-ipv4] [vpn-ipv4] no extended-nh-encoding</pre>
Context	<pre>config>router>bgp config>router>bgp>group</pre>

config>router>bgp>group>neighbor

Description This command configures BGP to advertise (at session OPEN) the capability to receive label IPv4, VPN IPv4 routes, or IPv6 next hops from the peers. These peers should not send such routes unless they receive notification of this capability. If the SR OS router receives a label IPv4 or VPN IPv4 route from a peer to which it did not advertise the necessary capability, the UPDATE message will be considered malformed and this will cause either session reset or **treat-as-withdraw** behavior depending on the error handling settings.

When the **no** form of the command is configured, the sending of an Extended NH Encoding BGP capability to the associated BGP peers is inherited from a higher configuration level or disabled (if configured at the BGP level).

Default no extended-nh-encoding

Parameters **label-ipv4** — Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=4, and next-hop AFI=2.
vpn-ipv4 — Instructs BGP to advertise an extended NH encoding capability for NLRI AFI=1, NLRI SAFI=128, and next-hop AFI=2.

family

Syntax **family** [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [ms-pw] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] [bgp-ls] [mcast-ipv6] [mcast-vpn-ipv6] [sr-policy-ipv4]
no family

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the set of BGP address families (AFI plus SAFI) to be supported by the applicable base router BGP sessions.

The **no** form of the command restores the default, which corresponds to unlabeled IPv4 unicast routes (AFI 1, SAFI 1) only.

Default family ipv4

Parameters **ipv4** — Adds support for the IPv4 unicast (unlabeled) address family.
label-ipv4 — Adds support for the IPv4 unicast (labeled) address family.
vpn-ipv4 — Adds support for the IPv4 VPN (SAFI 128) address family.
ipv6 — Adds support for the IPv6 unicast (unlabeled) address family.
label-ipv6 — Adds support for the IPv6 unicast (labeled) address family.
vpn-ipv6 — Adds support for the IPv6 VPN (SAFI 128) address family.
mcast-ipv4 — Adds support for the IPv4 multicast SAFI address family.

l2-vpn — Adds support for the L2 VPN address family.

mvpn-ipv4 — Adds support for the IPv4 multicast VPN address family.

mvpn-ipv6 — Adds support for the IPv6 multicast VPN address family.

mdt-safi — Adds support for the MDT SAFI address family.

ms-pw — Adds support for the multi-segment pseudowire address family.

flow-ipv4 — Adds support for the IPv4 flowspec address family.

flow-ipv6 — Adds support for the IPv6 flowspec address family.

route-target — Adds support for the address family for RT constrain routes.

mcast-vpn-ipv4 — Adds support for the IPv4 VPN multicast (SAFI 129) address family.

evpn — Adds support for the EVPN address family.

bgp-ls — Enables the advertisement of BGP-LS address family to the associated BGP neighbors.

mcast-ipv6 — Adds support for the IPv6 multicast SAFI address family.

mcast-vpn-ipv6 — Adds support for the IPv6 multicast routes from a VPRN over the provider network. This family is only applicable in the base BGP routing context.

sr-policy-ipv4 — Advertises the capability for AFI1/SAFI73, which corresponds to BGP routes that encode a segment routing policy to an IPv4 destination.

flowspec

Syntax	flowspec
Context	config>router>bgp
Description	This command enables the context to enable and disable flowspec validations.

graceful-restart

Syntax	[no] graceful-restart
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables BGP graceful restart helper procedures (the “receiving router” role defined in the standard) for address families included in the GR capabilities of both peers. SR OS can support GR helper functionality for IPv4, IPv6, VPN-IPv4, VPN-IPv6, Label-IPv4, Label-IPv6, L2-VPN, Route-Target (RTC), Flow-IPv4 (IPv4 flow-spec) and Flow-IPv6 (IPv6 flow-spec) routes.

If a neighbor covered by the GR helper mode restarts its control plane, forwarding can continue uninterrupted while the session is re-established and routes are re-learned.

The **no** form of the command disables graceful restart.

Default no graceful-restart

enable-notification

Syntax enable-notification
no enable-notification

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability then the session can be restarted gracefully (while preserving forwarding) if either peer needs to send a NOTIFICATION message due to some type of event or error.

Default no enable-notification

long-lived

Syntax [no] long-lived

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description This command enables the context to enter commands related to BGP Long-Lived Graceful-Restart (LLGR) procedures.

LLGR, known informally as BGP persistence, is an extension of BGP GR that allows a session to stay down for a longer period of time. During this time, learned routes are marked and re-advertised as stale but they can continue to be used as routes of last resort.

The LLGR handling of a session failure can be invoked immediately or it can be delayed until the end of the traditional GR restart window.

Default no long-lived

advertise-stale-to-all-neighbors

Syntax advertise-stale-to-all-neighbors [without-no-export | no without-no-export]

no advertise-stale-to-all-neighbors

Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived
Description	This command allows BGP routes marked as LLGR stale to be advertised to BGP peers that did not advertise the LLGR capability when the session was opened. The no version of this command causes advertisement behavior to follow the rule that stale routes cannot be advertised to a peer that does not understand or implement the LLGR capability. Stale routes are withdrawn towards such peers. When this command is configured with the without-no-export option, LLGR stale routes can be advertised to any peer (EBGP or IBGP) that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability, the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero. When this command is configured without the without-no-export option, LLGR stale routes are not advertised to any EBGP peer that did not signal the LLGR capability. Towards IBGP and confederation-EBGP peers that did not advertise the LLGR capability the LOCAL_PREFERENCE attribute in the advertised stale routes is automatically set to zero and a NO_EXPORT standard community is automatically added to the routes.
Default	no advertise-stale-to-all-neighbors
Parameters	without-no-export — Allows LLGR stale routes to be advertised to all peers, such that they can exit the local AS.

advertised-stale-time

Syntax	advertised-stale-time <i>seconds</i> no advertised-stale-time
Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>graceful-restart>long-lived>family config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived>family config>router>bgp>group>neighbor>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived>family
Description	This command sets the value of the long-lived stale time that is advertised by the router in its LLGR capability. When configured in the long-lived configuration context, advertised-stale-time applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the advertised-stale-time command in a family context. The no version of this command sets the advertised-stale-time value to 24 hours (86400 seconds).
Default	no advertised-stale-time

Parameters	<i>seconds</i> — Specifies the advertised long-lived stale time in seconds.
Values	0 to 16777215

family

Syntax	[no] family {ipv4 ipv6 label-ipv4 label-ipv6 vpn-ipv4 vpn-ipv6 l2-vpn route-target flow-ipv4 flow-ipv6}
Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived
Description	This command configures family-specific LLGR parameters for BGP peers.
Default	no family
Parameters	ipv4 — Specifies the IPv4 family. ipv6 — Specifies the IPv6 family. label-ipv4 — Specifies the label IPv4 family. label-ipv6 — Specifies the label IPv6 family. vpn-ipv4 — Specifies the VPN IPv4 family. vpn-ipv6 — Specifies the VPN IPv6 family. l2-vpn — Specifies the Layer 2 VPN family. route-target — Specifies the route target family. flow-ipv4 — Specifies the flow IPv4 family. flow-ipv6 — Specifies the flow IPv6 family.

restart-time

Syntax	restart-time seconds no restart-time
Context	config>router>bgp>graceful-restart config>router>bgp>group>graceful-restart config>router>bgp>group>neighbor>graceful-restart
Description	This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured, the default is 300.
Default	no restart time

Parameters	<i>seconds</i> — Specifies the restart-time that is advertised in the router's graceful-restart capability.
Values	0 to 4095 seconds
Default	config>router>bgp>graceful-restart: 120 seconds config>router>bgp>group>graceful-restart: 300 seconds config>router>bgp>group>neighbor>graceful-restart: 300 seconds

stale-routes-time

Syntax	stale-routes-time <i>time</i> no stale-routes-time
Context	config>router>bgp>graceful-restart config>router>bgp>group>graceful-restart config>router>bgp>group>neighbor>graceful-restart
Description	This command configures the maximum amount of time in seconds that stale routes should be maintained after a graceful restart is initiated. The no form of the command resets the stale routes time back to the default of 360 seconds.
Default	no stale-routes-time
Parameters	<i>time</i> — Specifies the amount of time that stale routes should be maintained after a graceful restart is initiated. Values 1 to 3600 seconds

validate-dest-prefix

Syntax	validate-dest-prefix no validate-dest-prefix
Context	config>router>bgp>flowspec
Description	This command enables or disables validation of received IPv4 and IPv6 flowspec routes that contain a destination-prefix subcomponent. A flowspec route with a destination-prefix subcomponent is considered invalid if both of the following are true: • it was originated outside the local AS of the receiving BGP router • the neighbor AS of the flowspec route does not match the neighbor AS of the best match BGP (unicast) route for the destination prefix or the neighbor AS of any longer match BGP (unicast) route for the destination prefix

An invalid route is retained in the BGP but it is not used for filtering traffic or propagated to other BGP routers.

The **no** form of the command disables the validation procedure based on destination-prefix.

Default no validate-dest-prefix

validate-redirect-ip

Syntax [no] **validate-redirect-ip**

Context config>router>bgp>flowspec

Description This command enables procedures to validate the **redirect-to-IPv4** action attached to flowspec IPv4 routes received by the BGP instance.

The SR OS flow-spec implementation supports the **redirect-to-IPv4** action encoded as an IPv4-address-specific BGP extended community.

When this command is configured, a flowspec IPv4 route is considered invalid and not installed as a filter rule if the flowspec IPv4 route is deemed to have originated in a different AS than the IP route that resolves the redirection IPv4 address. The originating AS of a flow-spec route is determined from its AS path.

The **no** form of the command disables the check described above.

Default no validate-redirect-ip

route-target-list

Syntax **route-target-list** *comm-id* [*comm-id*]
no route-target-list [*comm-id*]

Context config>router>bgp

Description This command specifies the route target(s) to be accepted from or advertised to peers. If the **route-target-list** is a non-null list, only routes with one or more of the given route targets are accepted from or advertised to peers.

The **route-target-list** is assigned at the global level and applies to all peers connected to the system.

This command is only applicable if the router is a route-reflector server.

The **no** form of the command with a specified route target community removes the specified community from the **route-target-list**. The **no** form of the command entered without a route target community removes all communities from the list.

Default no route-target-list

Parameters	<i>comm-id</i> — Specifies up to 15 route target communities.
Values	[target: {<i>ip-address:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i>} where: • <i>ip-address</i> — a.b.c.d • <i>comm-val</i> — 0 to 65535 • <i>2byte-asnumber</i> — 0 to 65535 • <i>ext-comm-val</i> — 0 to 4294967295 • <i>4byte-asnumber</i> — 0 to 4294967295

sr-policy-import

Syntax	[no] sr-policy-import
Context	config>router>bgp
Description	This command instructs BGP to import all statically-configured non-local segment routing policies from the segment routing DB into the BGP RIB so that they can be advertised, as originated routes, towards BGP peers supporting the sr-policy-ipv4 address family. The no form of the command instructs BGP to not import any statically defined segment routing policies into BGP.
Default	no sr-policy-import

third-party-nexthop

Syntax	third-party-nexthop no third-party-nexthop
Context	config>router>bgp config>router>bgp>group>neighbor
Description	Use this command to enable the router to send third-party next-hop to EBGp peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGp peer and advertised to another EBGp peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route. The no form of the command prevents BGP from performing any third party next-hop processing toward any single-hop EBGp peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.
Default	no third-party-nexthop

vpn-apply-export

Syntax	[no] vpn-apply-export
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command causes the base instance BGP export route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes. The no form of the command disables the application of the base instance BGP route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.
Default	no vpn-apply-export

vpn-apply-import

Syntax	[no] vpn-apply-import
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command causes the base instance BGP import route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes. The no form of the command disables the application of the base instance BGP import route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.
Default	no vpn-apply-import

leak-import

Syntax	leak-import <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i>] no leak-import
Context	config>router>bgp>rib-management>ipv4 config>router>bgp>rib-management>label-ipv4 config>router>bgp>rib-management>ipv6
Description	This command is used to specify route policies that control the importation of leak-eligible routes from the BGP RIB of another routing instance into the unlabeled-IPv4, unlabeled-IPv6, or labeled-IPv4 RIB of the base router. To leak a route from one routing instance to another, the origin and destination RIB types must be the same; for example, it is not possible to leak a route from an unlabeled-IPv4 RIB of a VPRN into the labeled-IPv4 RIB of the base router.

The **leak-import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine final action to accept or reject the route.

Only one of the 15 objects referenced by the **leak-import** command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters while the remaining 14 objects have a maximum length of 64 characters each.

When multiple **leak-import** commands are issued, the last command entered overrides the previous command.

When a **leak-import** policy is not specified, no BGP routes from other routing instances are leaked into the base router BGP RIB.

The **no** form of the command removes the policy association.

Default	no leak-import
Parameters	<i>plcy-or-long-expr</i> — Specifies up to 14 route policy names (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. <i>plcy-or-expr</i> — The route policy name (up to 64 characters long) or a policy logical expression (up to 64 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

route-table-import

Syntax	route-table-import <i>policy-name</i> no route-table-import
Context	config>router>bgp>rib-management>ipv4 config>router>bgp>rib-management>label-ipv4 config>router>bgp>rib-management>ipv6 config>router>bgp>rib-management>label-ipv6
Description	This command specifies the name of a route to control the importation of active routes from the IP route table into one of the BGP RIBs. If the route-table-import command is not configured, or if the command refers to an empty policy, all non-BGP routes from the IP route table are imported into the applicable RIB.

If the **route-table-import** command is configured, then routes dropped or rejected by the configured policy are not installed in the associated RIB. Rejected routes cannot be advertised to BGP peers associated with the RIB, but they can still be used to resolve BGP next-hops of routes in that RIB. If the active route for a prefix is rejected by the **route-table-import** policy, then the best BGP route for that prefix in the BGP RIB can be advertised to peers as though it is used.

Aggregate routes are always imported into each RIB, independent of the **route-table-import** policy.

Route modifications specified in the actions of a **route-table-import** policy are ignored and have no effect on the imported routes.

Default	no route-table-import
Parameters	<i>policy-name</i> — Specifies the name of a policy-statement (up to 64 characters).

helper-override-stale-time

Syntax	helper-override-stale-time <i>seconds</i> no helper-override-stale-time
Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>graceful-restart>long-lived>family config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived>family config>router>bgp>group>neighbor>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived>family
Description	This command overrides the LLGR stale-time advertised by a peer (in its LLGR capability) with a locally-configured value. When configured in the long-lived configuration context, helper-override-stale-time applies to all AFI/SAFI in the advertised LLGR capability except for any AFI/SAFI with a family-specific override. A family-specific override is configured with the helper-override-stale-time command in a family context. By default, the LLGR stale-time for an AFI/SAFI is the value signaled by the peer in the corresponding AFI/SAFI part of the LLGR capability.
Default	no helper-override-stale-time
Parameters	<i>seconds</i> — Specifies the locally imposed LLGR stale time in seconds. Values 0 to 16777215

forwarding-bits-set

Syntax	forwarding-bits-set {all non-fwd} no forwarding-bits-set

Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived
Description	This command determines the setting of the F bits in the GR and LLGR capabilities advertised by the router. When the F bit is set for an AFI/SAFI, it indicates that the advertising router was able to preserve forwarding state for the routes of that AFI/SAFI across the last restart. If a router restarts and does not set F=1, then when the session with a peer re-establishes the peer immediately deletes all LLGR stale routes it was preserving on behalf of the restarting router for the corresponding AFI/SAFI. This command allows the F bits for all advertised AFI/SAFI to be set to 1, or only the F bits for non-forwarding AFI/SAFI to be set to 1. Non-forwarding AFI/SAFI are the following configuration-related address families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.
Default	no forwarding-bits-set
Parameters	all — Specifies that the F bit for all AFI/SAFI should be set to 1. non-fwd — Specifies that the F bit for only non-forwarding AFI/SAFI should be set to 1. These AFI/SAFI correspond to the following families: L2-VPN, route-target, flow-IPv4, and flow-IPv6.

helper-override-restart-time

Syntax	helper-override-restart-time <i>seconds</i> no helper-override-restart-time
Context	config>router>bgp>graceful-restart>long-lived config>router>bgp>group>graceful-restart>long-lived config>router>bgp>group>neighbor>graceful-restart>long-lived
Description	This command overrides the restart-time advertised by a peer (in its GR capability) with a locally-configured value. This override applies only to AFI/SAFI that were included in the GR capability of the peer. The restart-time is always zero for AFI/SAFI not included in the GR capability. This command is useful if the local router wants to force LLGR phase to begin after a set time for all protected AFI/SAFI. By default, the restart time for all AFI/SAFI in the GR capability is the value signaled by the peer.
Default	no helper-override-restart-time
Parameters	seconds — The locally-imposed restart time for all AFI/SAFI included in the peer's GR capability. Values 0 to 4095

enforce-first-as

Syntax	enforce-first-as
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	When this command is configured so that it applies to an EBGp session, all routes (belonging to all address families) that are received from the EBGp peer are checked to ensure that the most recent autonomous system number (ASN) in the AS_PATH attribute of each route matches the configured peer-as of the session; if it does not match, then either the session is reset (if update-fault-tolerance is not enabled) or the session is left up but the route is treated as withdrawn (if update-fault-tolerance is enabled). Enabling or disabling this command on a session that is already up does not flap the session. When enforce-first-as is enabled, previously received routes are not checked for compliance with the rule. Enforcement applies only to routes received after the command is enabled and stops when the command is disabled.

error-handling

Syntax	error-handling
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command specifies whether updated BGP error handling procedures should be applied.

update-fault-tolerance

Syntax	[no] update-fault-tolerance
Context	config>router>bgp>error-handling config>router>bgp>group> error-handling config>router>bgp>group>neighbor>error-handling
Description	This command enables treat-as-withdraw and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.
Default	no update-fault-tolerance

group

Syntax	[no] group <i>name</i>
Context	config>router>bgp
Description	This command creates a context to configure a BGP peer group. The no form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shutdown before it can be deleted.
Default	no group
Parameters	<i>name</i> — Specifies the peer group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

aggregator-id-zero

Syntax	[no] aggregator-id-zero
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command sets the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths. When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute. When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level. The no form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no aggregator-id-zero

aigp

Syntax	[no] aigp
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables or disables Accumulated IGP (AIGP) path attribute support with one or more BGP peers. BGP path selection among routes with an associated AIGP metric is based on the end-to-end IGP metrics of the different BGP paths, even when these BGP paths span more than one AS and IGP instance. The effect of disabling AIGP (using the no form of the command or implicit) is to remove the AIGP attribute from advertised routes, if present, and to ignore the AIGP attribute in received routes.
Default	no aigp

disable-capability-negotiation

Syntax	[no] disable-capability-negotiation
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command disables capability negotiation. When the command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer. The no form of the command removes this command from the configuration and restores the normal behavior.
Default	no disable-capability-negotiation

dynamic-neighbor

Syntax	dynamic-neighbor
Context	config>router>bgp>group
Description	This command enables the context to configure dynamic BGP sessions for a peer group.

hold-time

Syntax	hold-time seconds [min seconds] no hold-time

Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the BGP hold time, expressed in seconds. The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used. Even though the implementation allows setting the keepalive time separately, the configured keepalive timer is overridden by the hold-time value under the following circumstances: 1. If the specified hold-time is less than the configured keepalive time, then the operational keepalive time is set to a third of the hold-time; the configured keepalive time is not changed. 2. If the hold-time is set to zero, then the operational value of the keepalive time is set to zero; the configured keepalive time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer. The no form of the command used at the global level reverts to the default value. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	hold-time 90
Parameters	seconds — Specifies the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently. Values 0, 3 to 65535 min seconds2 — Specifies the minimum hold-time that will be accepted for the session. If the peer proposes a hold-time lower than this value, the session attempt will be rejected.

ibgp-multipath

Syntax	[no] ibgp-multipath
Context	config>router>bgp
Description	This command enables IBGP multipath load balancing when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple nexthops. The no form of the command disables the IBGP multipath load balancing feature.

Default no ibgp-multipath

import

Syntax **import** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]]
no import

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies route policies that control the handling of inbound routes received from certain peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

The **import** command can reference up to 15 objects, where each object is either a policy logical expression or the name of a single policy. The objects are evaluated in the specified order to determine the modifications of each route and the final action to accept or reject the route.

Only one of the 15 objects referenced by the **import** command is allowed to be a policy logical expression consisting of policy names (enclosed in square brackets) and logical operators (AND, OR, NOT). The first of the 15 objects has a maximum length of 255 characters; the remaining 14 objects have a maximum length of 64 characters each.

When multiple **import** commands are issued, the last command entered overrides the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of the command removes the policy association.

Default no import

Parameters *plcy-or-long-expr* — Specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters long). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

plcy-or-expr — Specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 64 characters long). Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

initial-send-delay-zero

Syntax	[no] initial-send-delay-zero
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures BGP to send UPDATE messages announcing reachability information to a peer or set of peers immediately after the sessions become established with these peers. The no form of this command waits for min-route-advertisement time after each session is established before sending the first set of UPDATE messages.

keepalive

Syntax	keepalive seconds no keepalive
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The keepalive parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The keepalive value is generally one-third of the hold-time interval. Even though the implementation allows the keepalive value and the hold-time interval to be independently set, under the following circumstances, the configured keepalive value is overridden by the hold-time value: 1. If the specified keepalive value is greater than the configured hold-time, then the specified value is ignored and the keepalive is set to one third of the current hold-time value. 2. If the specified hold-time interval is less than the configured keepalive value, then the keepalive value is reset to one third of the specified hold-time interval. 3. If the hold-time interval is set to zero, then the configured value of the keepalive value is ignored. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer. The no form of the command used at the global level reverts to the default value The no form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	keepalive 30
Parameters	<i>seconds</i> — Specifies the keepalive timer, in seconds, expressed as a decimal integer.
Values	0 to 21845

label-preference

Syntax	label-preference <i>value</i> no label-preference
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the route preference for routes learned from labeled-unicast peers. This command can be configured at three levels: • Global level — applies to all peers • Group level — applies to all peers in the peer-group • Neighbor level — applies only to the specified peer The most specific value is used. The lower the preference, the higher the chance of the route being the active route. The no form of the command used at the global level reverts to the default <i>value</i> of 170. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no label-preference
Parameters	<i>value</i> — Specifies the route preference value.
Values	1 to 255

link-bandwidth

Syntax	link-bandwidth
Context	config>router>bgp>group

```
config>router>bgp>group>neighbor
```

Description	This command enables the configuration context for handling the link-bandwidth extended community attached to specific BGP routes.

When all used multipaths of an IP prefix correspond to BGP routes with a link-bandwidth extended community, the datapath is programmed to do weighted ECMP across the BGP next-hops in proportion to the bandwidth values.

accept-from-ebgp

Syntax **accept-from-ebgp** *family* [*family*]
no accept-from-ebgp

Context config>router>bgp>group>link-bandwidth
config>router>bgp>group>neighbor>link-bandwidth

Description	This command configures BGP to accept and use the link-bandwidth extended community attached to any route received from any EBGP peer in the scope of the command, as long as that route belongs to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGP peer and it should be discarded when received in any route from an EBGP peer. This command overrides the standard behavior.

Up to six families may be configured.

The **no** form of this command restores the default behavior of discarding the link-bandwidth extended community in any route received from an EBGP peer.

Default no accept-from-ebgp

Parameters *family* — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported.

Values	
ipv4	— Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes.
label-ipv4	— Adds a link-bandwidth extended community to labeled-unicast IPv4 routes.
vpn-ipv4	— Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes.
ipv6	— Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes.
label-ipv6	— Adds a link-bandwidth extended community to labeled-unicast IPv6 routes.
vpn-ipv6	— Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

add-to-received-ebgp

Syntax	add-to-received-ebgp <i>family</i> [<i>family</i>] no add-to-received-ebgp
Context	config>router>bgp>group>link-bandwidth config>router>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to automatically add a link-bandwidth extended community to every route received from a directly connected (single-hop) EBGp peer within the scope of the command, as long as that route belongs to one of the listed address families. The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGp peer. Up to six families may be configured. The no form of this command removes the link-bandwidth extended community added to received BGP routes.
Default	no add-to-received-ebgp
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported. Values - ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes. - label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. - vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes. - ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes. - label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes. - vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

aggregate-used-paths

Syntax	aggregate-used-paths <i>family</i> [<i>family</i>] no aggregate-used-paths
Context	config>router>bgp>group>link-bandwidth config>router>bgp>group>neighbor>link-bandwidth

Description	This command configures BGP to aggregate the bandwidth values from the link-bandwidth extended communities of the used multipaths towards an IP prefix when it is re-advertising a route with next-hop-self towards peers within the scope of the command, as long as the route belongs to one of the listed address families. Aggregation is not supported unless all of the used multipaths (up to the configured ECMP limit) correspond to received BGP routes with a link-bandwidth extended community. If add-path is also enabled toward the peer, then all of the add-paths advertised to the peer encode the aggregated bandwidth in a link-bandwidth extended community. Up to six families may be configured. The no form of the command disables aggregation in a next-hop-self scenario and the link-bandwidth extended community in the advertised route is a copy of the link-bandwidth extended community in the received route (which may have been added by import policy or by the effect of the add-to-received-ebgp command).
Default	no aggregate-used-paths
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported. Values - ipv4 — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes. - label-ipv4 — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. - vpn-ipv4 — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes. - ipv6 — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes. - label-ipv6 — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes. - vpn-ipv6 — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

send-to-ebgp

Syntax	send-to-ebgp <i>family</i> [<i>family</i>] no send-to-ebgp
Context	config>router>bgp>group>link-bandwidth config>router>bgp>group>neighbor>link-bandwidth
Description	This command configures BGP to allow link-bandwidth extended community to be sent in routes advertised to EBGp peers in the scope of the command, as long the routes belong to one of the listed address families.

The link-bandwidth extended community is encoded as a non-transitive type. This means that by default it should not be attached to any route advertised to an EBGp peer and it should be discarded when received in any route from an EBGp peer. This command overrides the standard behavior.

Up to six families may be configured.

The **no** form of the command restores the default behavior of stripping the link-bandwidth extended community from any route advertised to an EBGp peer.

Default	no send-to-ebgp
Parameters	<i>family</i> — Specifies the address families for which receiving the link-bandwidth extended community from EBGp peers should be supported.
Values	<i>ipv4</i> — Adds a link-bandwidth extended community to unlabeled unicast IPv4 routes. <i>label-ipv4</i> — Adds a link-bandwidth extended community to labeled-unicast IPv4 routes. <i>vpn-ipv4</i> — Adds a link-bandwidth extended community to IPv4 VPN (SAFI 128) routes. <i>ipv6</i> — Adds a link-bandwidth extended community to unlabeled unicast IPv6 routes. <i>label-ipv6</i> — Adds a link-bandwidth extended community to labeled-unicast IPv6 routes. <i>vpn-ipv6</i> — Adds a link-bandwidth extended community to IPv6 VPN (SAFI 128) routes.

link-state-import-enable

Syntax	[no] link-state-import-enable
Context	config>router>bgp
Description	This command enables the import of link-state information into the BGP-LS address family for advertisement to other BGP neighbors. The no form of this command disables the import of link state information into the BGP-LS address family.
Default	no link-state-import-enable

local-address

Syntax	local-address [<i>ip-int-name</i> <i>ip-address</i> <i>ipv6-address</i>] no local-address

Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the local IP address used by the group or neighbor when communicating with BGP peers. Outgoing connections use the local-address as the source of the TCP connection when initiating connections with a peer. When a local address is not specified, the router uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level. When set to a router interface, the local-address inherits the primary IPv4 or IPv6 address of the router interface depending on whether BGP is configured for IPv4 or IPv6. If the corresponding IPv4 or IPv6 address is not configured on the router interface, the BGP sessions that have this interface set as the local-address are kept down until an interface address is configured on the router interface. The no form of the command removes the configured local-address for BGP. The no form of the command used at the group level returns the configuration to the value defined at the global level. The no form of the command used at the neighbor level returns the configuration to the value defined at the group level.
Default	no local-address
Parameters	<i>ip-address</i> — Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. Values ipv4-address: • a.b.c.d (host bits must be 0) <i>ipv6-address</i> — Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IPv6 address on the router, either an interface or system IPv6 address. Values ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D <i>ip-int-name</i> — Specifies the IP interface name whose address the local address will inherit. The interface can be any network interface configured on the system.

local-as

Syntax	local-as <i>as-number</i> [private] [no-prepend-global-as] no local-as
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures a BGP local autonomous system (AS) number. In addition to the global AS number configured for BGP using the autonomous-system command, a local AS number can be configured to support various AS number migration scenarios. When the local-as command is applied to a BGP neighbor and the local-as is different from the peer-as, the session comes up as EBGP and by default the global-AS number and then (in that order) the local-as number are prepended to the AS_PATH attribute in outbound routes sent to the peer. In received routes from the EBGP peer, the local AS is prepended to the AS path by default, but this can be disabled with the private option. When the local-as command is applied to a BGP neighbor and the local-as is the same as the peer-as, the session comes up as IBGP, and by default, the global-AS number is prepended to the AS_PATH attribute in outbound routes sent to the peer. This configuration parameter can be set at three levels: global level (applies to all BGP peers), group level (applies to all BGP peers in group) or neighbor level (only applies to one specific BGP neighbor). Thus by specifying this at the neighbor level, it is possible to have a separate local-as for each BGP session. When the optional no-prepend-global-as command is configured, the global-as number is not added in outbound routes sent to an IBGP or EBGP peer. When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The private option can be added or removed dynamically by reissuing the command. Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no local-as

Parameters	<i>as-number</i> — Specifies the virtual autonomous system number expressed as a decimal integer.
Values	1 to 4294967295
	private — Specifies the local-as is hidden in paths learned from the peering.
	no-prepend-global-as — Specifies that the global-as is hidden in paths announced to the BGP peer.

local-preference

Syntax	local-preference <i>local-preference</i> no local-preference
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the local-preference integer set. The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	no local-preference
Parameters	<i>local-preference</i> — Specifies the local preference value to be used as the override value expressed as a decimal integer. Values 0 to 4294967295

loop-detect

Syntax	loop-detect { drop-peer discard-route ignore-loop off }

no loop-detect

Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures how the BGP peer session handles loop detection in the AS path. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.



Note: Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	loop-detect ignore-loop
Parameters	drop-peer — Sends a notification to the remote peer and drops the session. discard-route — Discards routes received from a peer with the same AS number as the router itself. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer. ignore-loop — Ignores routes with loops in the AS path but maintains peering. off — Disables loop detection.

med-out

Syntax	med-out { <i>number</i> igp-cost } no med-out
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no med-out
Parameters	<i>number</i> — Specifies the MED path attribute value, expressed as a decimal integer.
Values	0 to 4294967295
	igp-cost — Sets MED to the IGP cost of the given IP prefix.

min-route-advertisement

Syntax	min-route-advertisement <i>seconds</i> no min-route-advertisement
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the minimum interval, in seconds, between successive updates of a prefix towards a peer. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group), or neighbor level (only applies to specified peer). The most specific value is used. The no form of the command used at the global level reverts to default. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level. The rapid-update command can be used to override the peer-level min-route-advertisement time and applies the minimum setting (0 seconds) to routes belonging to address families specified by the rapid-update command; routes of other address families continue to be advertised according to the session-level MRAI setting.

The **rapid-update** and **rapid-withdrawal** commands may result in the routes being sent before the peer-level MRAl timer expires.

Default	min-route-advertisement 30
Parameters	<i>seconds</i> — Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.
Values	1 to 255

mp-bgp-keep

Syntax	[no] mp-bgp-keep
Context	config>router>bgp
Description	As a result of enabling this command, route refresh messages are no longer needed, or issued when VPN route policy changes are made; RIB-IN will retain all MP-BGP routes. The no form of the command is used to disable this feature.
Default	no mp-bgp-keep

multihop

Syntax	multihop ttl-value no multihop
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGp peer multiple hops away. The no form of the command is used to convey to the BGP instance that the EBGp peers are directly connected. The no form of the command used at the global level reverts to default. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	multihop 1 — EBGp peers are directly connected. multihop 64 — Ibgp

Parameters *ttl-value* — Specifies the TTL value, expressed as a decimal integer.

Values 1 to 255

multipath

Syntax **multipath** *max-paths* [**ebgp** *ebgp-max-paths*] [**ibgp** *ibgp-max-paths*] [**restrict** {**same-neighbor-as** | **exact-as-path**}]
no multipath

Context config>router>bgp

Description This command enables BGP multipath for all address families that support ECMP forwarding.

When multipath is enabled, traffic to the destination is load-shared across a set of paths (BGP routes) that the BGP decision process considers “equal” to the best path. The actual distribution of traffic over the multiple paths may be equal or unequal (that is, based on weights derived from the link-bandwidth extended community).

To qualify as a multipath, a non-best route must meet the following criteria (some criteria are controlled by this command).

- The route must be the same type of route as the best path (same AFI/SAFI and, in some cases, same next-hop resolution method)
- The route must be tied with the best path for all criteria of equal or greater importance than IGP cost to reach the BGP next-hop, except for criteria that are configured to be ignored
- The route must not have the same BGP next-hop as the best path or any other multipath
- The route must not cause the ECMP limit of the routing instance to be exceeded (configured using the **ecmp** command to a value in the range 1 to 64)
- The route must not cause the configured *max-paths* limit of the BGP instance to be exceeded

If the best path is an EBGp learned route and the **ebgp** option is used, the *ebgp-max-paths* limit overrides the *max-paths* limit. If the best path is an IBGP learned route and the **ibgp** option is used, the *ibgp-max-paths* limit overrides the *max-paths* limit.

All path limits are configurable up to a maximum of 64. Multipath is effectively disabled if a value is set to 1.

- The route must have the same neighbor AS in its AS path as the best path if the **restrict same-neighbor-as** option is configured. By default, any path with the same AS path length as the best path (regardless of neighbor AS) is eligible for multipath.
- The route must have the same AS path as the best path if the **restrict exact-as-path** option is configured. By default, any path with the same AS path length as the best path (regardless of the actual AS numbers) is eligible for multipath.

The **no** form of the command disables BGP multipath (equivalent to multipath 1).

Default no multipath

Parameters	<i>max-paths</i> — Specifies the maximum number of multipaths per prefix/NLRI if <i>ebgp-max-paths</i> or <i>ibgp-max-paths</i> does not apply. Values 1 to 64 <i>egp-max-paths</i> — Specifies the maximum number of multipaths per prefix/NLRI when the best path is an EBGp learned route. Values 1 to 64 <i>ibgp-max-paths</i> — Specifies the maximum number of multipaths per prefix/NLRI when the best path is an IBGP learned route. Values 1 to 64 restrict same-neighbor-as — Specifies that the non-best path must have the same neighbor AS in its AS path as the best path. restrict exact-as-path-as — Specifies that the non-best path must have the same AS path as the best path.

mvpn-vrf-import-subtype-new

Syntax	[no] mvpn-vrf-import-subtype-new
Context	config>router>bgp
Description	When enabled, the type/subtype in advertised routes is encoded as 0x010b. The no form of the command (the default) encodes the type/subtype as 0x010a (to preserve backwards compatibility).
Default	no mvpn-vrf-import-subtype-new

neighbor-trust

Syntax	neighbor-trust [vpn-ipv4] [vpn-ipv6] no neighbor-trust
Context	config>router>bgp
Description	This command enables a label security feature for prefixes of a VPN family at an inter-AS boundary. This label security feature allows the configuration of a router, acting in a PE and/or in an ASBR role, to accept packets of VPRN prefixes only from direct eBGP neighbors to which it advertised a VPRN label. The untrusted state identifies the participating interfaces. The router supports a maximum of 15 network interfaces that can participate in this feature.

At a high level, BGP tracks each direct eBGP neighbor over an untrusted interface and to which it sent a VPRN prefix label. For each of those VPRN prefixes, BGP programs a bit map in the ILM record that indicates, on per-untrusted interface basis, whether the matching received packets must be forwarded or dropped.

The **no** form of the command disables the inter-AS security feature for the VPN family.

Parameters **vpn-ipv4** — Enables the inter-AS label security for VPN IPv4 family.
 vpn-ipv6 — Enables the inter-AS label security for VPN IPv6 family.

next-hop-resolution

Syntax **next-hop-resolution**

Context config>router>bgp

Description This command enables the context to configure next-hop resolution parameters.

labeled-routes

Syntax **labeled-routes**

Context config>router>bgp>next-hop-res

Description This command enables the context to configure labeled route options for next-hop resolution.

allow-static

Syntax **allow-static**
 no allow-static

Context config>router>bgp>next-hop-res>labeled-routes

Description This command allows the BGP next-hop of label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6 routes received from any EBGP or IBGP peer to be resolved using static routes, except for static default routes (0/0 and ::/0).

A static route is less preferred than a local or interface route for resolving the BGP next-hop of labeled route, but more preferred than other IGP routes or tunnels.



Note: A label-IPv4 or label-IPv6 route can be resolved by a static blackhole route, even when the **allow-static** command is not configured, but only if the static blackhole route is the longest prefix match (LPM) static route for the BGP next-hop address.

Default no allow-static

rr-use-route-table

Syntax **rr-use-route-table**
no rr-use-route-table

Context config>router>bgp>next-hop-res>labeled-routes

Description This command enables BGP to perform a lookup of IGP routes in the route table to resolve the BGP next-hop of label-IPv4 and label-IPv6 routes. This is useful for a Route Reflector (RR) that does not participate in tunnel signaling protocols such as LDP and RSVP and therefore, does not have tunnels to resolve the BGP next-hops of label-unicast routes.

Configure the **disable-route-table-install** command before you configure the **rr-use-route-table** command because forwarding would otherwise be incorrect for cases where label routes are resolved this way.

Default no rr-use-route-table

transport-tunnel

Syntax **transport-tunnel**

Context config>router>bgp>next-hop-res>labeled-routes

Description This command enables the context to configure options for the next-hop resolution of BGP labeled routes (VPN-IP and labeled-unicast) using tunnels in TTM. The context allows the selection of different tunnel resolution options for different types of BGP labeled routes: label-unicast IPv4, label-unicast IPv6, and VPN-IP routes (both VPN-IPv4 and VPN-IPv6).

By default (if this context and the resolution options are not configured), these routes resolve only to LDP tunnels.

If the **resolution** option is explicitly set to **disabled**, the default binding to LDP tunnel resumes. If **resolution** is set to **any**, then any supported tunnel type is allowed and the selection is based on the lowest numerical TTM preference value.

family

Syntax **family {label-ipv4 | label-ipv6 | vpn}**

Context config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel

Description This command configures the address family context for configuring next-hop resolution of BGP label routes.

- Parameters**
- label-ipv4** — Enters the context for configuring next-hop-resolution options for labeled-unicast IPv4 routes.
 - label-ipv6** — Enters the context for configuring next-hop-resolution options for labeled-unicast IPv6 routes.
 - vpn** — Enters the context for configuring next-hop-resolution options for VPN-IPv4 and VPN-IPv6 routes when they are not imported into any VPRN service.

enforce-strict-tunnel-tagging

- Syntax** `[no] enforce-strict-tunnel-tagging`
- Context** `config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family`
`config>router>bgp>next-hop-resolution>shortcut-tunn>family`
- Description** This command forces the system to only consider LSPs marked with an **admin-tag** for next-hop resolution. Untagged LSPs are not be considered.
- The **no** form of this command reverts to the default behavior. While tagged RSVP and SR-TE LSPs will be considered first, the system can fall back to using tagged LSPs that are not explicitly excluded by a route admin tag policy and untagged LSPs of other types and not exclude them.
- Default** `no enforce-strict-tunnel-tagging`

resolution

- Syntax** `resolution {any | filter | disabled}`
- Context** `config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family`
- Description** This command configures the resolution mode in the resolution of BGP label routes using tunnels to BGP peers.
- Parameters**
- any** — Enables the binding to any supported tunnel type in the BGP label route context following TTM preference.
 - filter** — Enables the binding to the subset of tunnel types configured under **resolution-filter**.
 - disabled** — Disables the resolution of BGP label routes using tunnels to BGP peers.

resolution-filter

- Syntax** `resolution-filter`
- Context** `config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family`

Description This command enters the configuration mode to set resolution filter types.

bgp

Syntax [no] **bgp**

Context config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter
config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter

Description This command selects BGP tunneling for next-hop resolution and specifies the IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32.

IPv4 prefixes

ldp

Syntax [no] **ldp**

Context config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter
config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter

Description This command selects LDP tunneling for next-hop resolution and specifies the LDP tunnels in the tunnel table corresponding to /32 IPv4 FECs and /128 IPv6 FECs.

rib-api

Syntax [no] **rib-api**

Context config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter
config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter

Description This command enables tunnels programmed using the RibApi gRPC service for use in resolving the next hops of label-IPv4 or label-IPv6 routes.

rsvp

Syntax [no] **rsvp**

Context config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter
config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter

Description	This command selects RSVP tunneling for next-hop resolution and specifies RSVP tunnels in a tunnel table to IPv4 destinations. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

sr-isis

Syntax	[no] sr-isis
Context	config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment-routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or, in case of a tie, the lowest numbered IS-IS instance.

sr-ospf

Syntax	[no] sr-ospf
Context	config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command selects the Segment Routing (SR) tunnel type programmed by an OSPF instance in TTM for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPF instance or, in case of a tie, the lowest numbered OSPF instance.

sr-ospf3

Syntax	[no] sr-ospf3
Context	config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command selects the IPv6 segment routing tunnel type programmed by an OSPFv3 instance in the TTMv6 for next-hop resolution of BGP routes and labeled routes. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPFv3 instance or, in case of a tie, the lowest-numbered OSPFv3 instance.

sr-policy

Syntax	[no] sr-policy
Context	config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command enables the use of segment routing policies to resolve the BGP next hop of certain BGP routes (depending on the context). The segment routing policies that are considered are statically configured in the local router or learned by BGP routes (AFI 1/SAFI 73). For a BGP route to be resolved by an SR policy, the highest numbered color extended community attached to BGP route must match the color of the SR policy. Next hop resolution of VPN IP routes by SR policies is not supported.

sr-te

Syntax	[no] sr-te
Context	config>router>bgp>next-hop-resolution>shortcut-tunn>family>resolution-filter config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command selects the Segment Routing (SR) tunnel type programmed by a traffic engineered (TE) instance in TTM for next-hop resolution. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel ID.

udp

Syntax	[no] udp
Context	config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family>resolution-filter
Description	This command selects UDP tunnel in TTM for next-hop resolution.

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>router>bgp>next-hop-resolution

Description	This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next-hops to MPLS tunnels. If a BGP next-hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved; if the route N is accepted by the policy then it becomes the resolving route for R. The default next-hop resolution policy (when the no policy command is configured) is to use the longest matching active route in RTM that is not a BGP route (unless use-bgp-routes is configured), an aggregate route or a subscriber management route.
Default	no policy
Parameters	<i>policy-name</i> — Specifies the route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the config>router>policy-options context.

shortcut-tunnel

Syntax	shortcut-tunnel
Context	config>router>bgp>next-hop-resolution
Description	This command creates the context to configure the tunnel types that can be used to resolve unlabeled IPv4 and IPv6 BGP routes. The following tunnel types are supported for resolving IPv4 routes and IPv6 routes with IPv4-mapped IPv6 next-hop addresses: bgp, ldp, rsvp, sr-isis, sr-ospf, sr-policy and sr-te. In this context: • bgp — refers to IPv4 tunnels created by receiving BGP label-unicast IPv4 routes for /32 IPv4 prefixes • ldp — refers to /32 and shorter length LDP FEC prefixes imported into the tunnel table. For IPv4 NLRI, BGP selects the LDP FEC that is the longest-prefix-match (LPM) of the BGP next-hop address. For IPv6 NLRI, BGP selects the /32 FEC that is an exact match of the BGP next-hop address. • rsvp — refers to RSVP tunnels in the tunnel table to IPv4 destinations. This option allows BGP to use the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback interface of the remote BGP router. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel id. • sr-isis — refers to segment routing tunnels (shortest path) to IPv4 destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance.

- **sr-ospf** — refers to segment routing tunnels (shortest path) to IPv4 destinations reachable by the OSPF protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference OSPF instance or (in case of a tie) the lowest numbered OSPF instance.
- **sr-policy** — refers to segment routing policies with an IPv4 endpoint that are statically configured in the local router or learned through BGP routes (AFI 1/SAFI 73). For BGP to resolve the next hop of an unlabeled IPv4 or IPv6 route using a segment routing policy the highest numbered color extended community attached to the IPv4 or IPv6 route must match the color of the segment routing policy.
- **sr-te** — refers to traffic engineered (TE) segment routing tunnels. This option allows BGP to use the best metric SR-TE tunnel to the address of the BGP next-hop. In the case of multiple SR-TE tunnels with the same lowest metric, BGP selects the tunnel with the lowest tunnel id.
- **udp** — refers to MPLSoUDPoIPv4 tunnels set up by action of the BGP import policies

The following tunnel types are supported for resolving IPv6 routes with IPv6 next-hops that are not IPv4-mapped IPv6 addresses: ldp, sr-isis, and sr-policy. In this context:

- **ldp** — refers to /128 LDP FEC prefixes in the tunnel table. BGP selects the /128 FEC that is an exact match of the BGP next-hop address.
- **sr-isis** — refers to segment routing tunnels (shortest path) to IPv6 destinations reachable by the IS-IS protocol. This option allows BGP to use the segment routing tunnel in the tunnel table submitted by the lowest preference IS-IS instance or (in case of a tie) the lowest numbered IS-IS instance.
- **sr-policy** — refers to segment routing policies with a null IPv4 endpoint (0.0.0.0) that are statically configured in the local router or learned through BGP routes (AFI 1/SAFI 73). For BGP to resolve the next hop of an IPv6 route using a segment routing policy the highest numbered color extended community attached to the IPv6 route must match the color of the segment routing policy and its color bits must be set to '01' or '10'.

family

Syntax	family {ipv4 ipv6}
Context	config>router>bgp>next-hop-res>shortcut-tunnel
Description	This command creates the context to configure next-hop resolution of unlabeled IPv4 or unlabeled IPv6 routes by certain tunnel types in the tunnel table.
Parameters	ipv4 — Configuration applies to unlabeled IPv4 BGP routes. ipv6 — Configuration applies to unlabeled IPv6 BGP routes.

resolution

Syntax	resolution {any filter disabled}

Context	config>router>bgp>next-hop-res>shortcut-tunn>family
Description	This command configures the resolution mode in the resolution of BGP prefixes using tunnels to BGP peers.
Parameters	any — Enables the binding to any supported tunnel type in BGP shortcut context following TTM preference. filter — Enables the binding to the subset of tunnel types configured under resolution-filter. disabled — Disables the resolution of BGP prefixes using tunnels to BGP peers.

resolution-filter

Syntax	resolution-filter [bgp] [ldp] [rsvp] [sr-isis] [sr-ospf] [sr-policy] [sr-te]
Context	config>router>bgp>next-hop-res>shortcut-tunn>family
Description	This command configures the subset of tunnel types that can be used in the resolution of BGP unlabeled routes.
Parameters	bgp — Selects the BGP label route tunnel type. ldp — Selects the LDP tunnel type. rsvp — Selects the RSVP-TE tunnel type. sr-isis — Selects the SR tunnel type programmed by an IS-IS instance in TTM. sr-ospf — Selects the SR tunnel type programmed by an OSPF instance in TTM. sr-policy — Selects the SR tunnel type programmed by an SR policy instance in TTM. sr-te — Selects the SR tunnel type programmed by a TE instance in TTM.

use-bgp-routes

Syntax	[no] use-bgp-routes
Context	config>router>bgp>next-hop-res
Description	This command specifies whether to use BGP routes to recursively resolve the BGP next-hop of unlabeled IPv4 and unlabeled IPv6 routes. Up to four levels of recursion are supported. By default, BGP routes are not considered by the next-hop resolution process.
Default	no use-bgp-routes

weighted-ecmp

Syntax	[no] weighted-ecmp
Context	config>router>bgp>next-hop-res
Description	This command enables weighted ECMP for next-hop tunnel selection for 6PE. When weighted ECMP is enabled, the RSVP-TE tunnel used to forward 6PE packets to the ECMP next hop is chosen according to the outcome of the hash on the packet at the normalized load-balancing weight of the tunnel. The no version of this command disables weighted ECMP for next-hop tunnel selection for 6PE.
Default	no weighted-ecmp

optimal-route-reflection

Syntax	optimal-route-reflection
Context	config>router>bgp
Description	This command creates the optimal route reflection context.

spf-wait

Syntax	spf-wait <i>max-wait</i> [<i>initial-wait initial-wait</i>] [<i>second-wait second-wait</i>] no spf-wait
Context	config>router>bgp>optimal-route-reflection
Description	This command controls the interval between consecutive SPF calculations performed by the TE DB in support of BGP optimal route reflection. The time parameters of this command implement an exponential back-off algorithm. The no form of this command causes a return to default values.
Default	no spf-wait
Parameters	<i>max-wait</i> — Specifies the maximum interval in seconds between two consecutive SPF calculations. Values 1 to 600 Default 60

initial-wait *initial-wait* — Specifies the initial SPF calculation delay in seconds after a topology change.

Values 1 to 300

Default 5

second-wait *second-wait* — Specifies the delay in seconds between the first and second SPF calculation and the exponential delay increment.

Values 1 to 300

Default 15

location

Syntax	location <i>location-id</i> [primary-ip-address <i>ipv4-address</i>] [secondary-ip-address <i>ipv4-address</i>] [tertiary-ip-address <i>ipv4-address</i>]
Context	config>router>bgp>optimal-route-reflection
Description	This command configures the location ID for the for the route reflector. A BGP neighbor can be associated with a location if it is a route-reflector client.
Parameters	<i>location-id</i> — Specifies an optimal-route-reflection location. Values 1 to 16 <i>ipv4-address</i> — Specifies the primary, secondary, or tertiary IP address. Values primary ipv4-address, secondary ipv4-address, tertiary ipv4-address,

primary-ip-address

Syntax	primary-ip-address <i>ipv4-address</i> no primary-ip-address
Context	config>router>bgp>orr>location
Description	This command specifies the primary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location. If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address. The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location. The no form of this command removes the primary IP address information.

Default	no primary-ip-address
Parameters	<i>ipv4-address</i> — Specifies the primary IPv4 address of a location expressed in dotted decimal notation.
Values	a.b.c.d

primary-ipv6-address

Syntax	primary-ipv6-address <i>ipv6-address</i> no primary-ipv6-address
Context	config>router>bgp>orr>location
Description	This command specifies the primary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location. If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address. The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location. The no form of this command removes the primary IPv6 address information.
Default	no primary-ipv6-address
Parameters	<i>ipv6-address</i> — Specifies the primary IPv6 address of a location expressed in dotted decimal notation.
Values	ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D

secondary-ip-address

Syntax	secondary-ip-address <i>ipv4-address</i> no secondary-ip-address
Context	config>router>bgp>orr>location
Description	This command specifies the secondary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location.

If the TE DB is unable to find a node in its topology database that matches the primary address, then the TE DB tries to find a node with the matching secondary address. If this attempt also fails, the TE DB then tries to find a node with the matching tertiary address.

The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location.

The **no** form of this command removes the secondary IP address information.

Default	no secondary-ip-address
Parameters	<i>ipv4-address</i> — Specifies the secondary IPv4 address of a location, expressed in dotted decimal notation.
Values	a.b.c.d

secondary-ipv6-address

Syntax	secondary-ipv6-address <i>ipv6-address</i> no secondary-ipv6-address
Context	config>router>bgp>orr>location
Description	This command specifies the secondary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location. If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address. The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location. The no form of this command removes the secondary IPv6 address information.
Default	no secondary-ipv6-address
Parameters	<i>ipv6-address</i> — Specifies the secondary IPv6 address of a location.
Values	ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D

tertiary-ip-address

Syntax	tertiary-ip-address <i>ipv4-address</i> no tertiary-ip-address
Context	config>router>bgp>orr>location
Description	This command specifies the tertiary IP address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location. If the TE DB is unable to find a node in its topology database that matches the primary address, then the TE DB tries to find a node with the matching secondary address. If this attempt also fails, the TE DB then tries to find a node with the matching tertiary address. The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location. The no form of this command removes the tertiary IP address information.
Default	no tertiary-ip-address
Parameters	<i>ipv4-address</i> — Specifies the tertiary IPv4 address of a location, expressed in dotted decimal notation. Values a.b.c.d

tertiary-ipv6-address

Syntax	tertiary-ipv6-address <i>ipv6-address</i> no tertiary-ipv6-address
Context	config>router>bgp>orr>location
Description	This command specifies the tertiary IPv6 address of a reference location used for BGP optimal route reflection. Up to three IPv4 addresses and three IPv6 addresses can be specified per location. If the TE DB is unable find a node in its topology database that matches a primary address of the location, then it tries to find a node matching a secondary address. If this attempt also fails, the TE DB tries to find a node matching a tertiary address. The IP addresses specified for a location should be topologically “close” to a set of clients that should all receive the same optimal path for that location. The no form of this command removes the tertiary IPv6 address information.
Default	no tertiary-ipv6-address

Parameters *ipv6-address* — Specifies the tertiary IPv6 address of a location.

- Values** *ipv6-address*:
- x:x:x:x:x:x:x (eight 16-bit pieces)
 - x:x:x:x:x:d.d.d.d
 - x: [0 to FFFF]H
 - d: [0 to 255]D

outbound-route-filtering

Syntax **[no] outbound-route-filtering**

Context config>router>bgp
 config>router>bgp>group
 config>router>bgp>group>neighbor

Description This command opens the configuration tree for sending or accepting BGP filter lists from peers (outbound route filtering).

Default no outbound-route-filtering

peer-tracking-policy

Syntax **peer-tracking-policy *policy-name***
 no peer-tracking-policy

Context config>router>bgp

Description This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The default peer-tracking policy (when the **no peer-tracking-policy** command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.



Note: When **peer-tracking** is configured, the peer-tracking policy should only permit one of **direct-interface** or **direct** routes to be advertised to a BGP peer. Advertising both routes will cause the best route to oscillate.

Default no peer-tracking-policy

Parameters *policy-name* — Specifies the route policy name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

extended-community

Syntax **[no] extended-community**

Context config>router>bgp>outbound-route-filtering
config>router>bgp>group>outbound-route-filtering
config>router>bgp>group>neighbor>outbound-route-filtering

Description The extended-community command opens the configuration tree for sending or accepting extended-community based BGP filters.

For the **no** version of the command to work, all sub-commands (**send-orf**, **accept-orf**) must be removed first.

Default no extended-community

accept-orf

Syntax **[no] accept-orf**

Context config>router>bgp>outbound-route-filtering>extended-community
config>router>bgp>group>outbound-route-filtering>extended-community
config>router>bgp>group>neighbor>outbound-route-filtering>extended-community

Description This command instructs the router to negotiate the receive capability in the BGP ORF negotiation with a peer, and to accept filters that the peer wishes to send.

The **no** form of the command causes the router to remove the accept capability in the BGP ORF negotiation with a peer, and to clear any existing ORF filters that are currently in place.

Default no accept-orf

send-orf

Syntax **send-orf** [*comm-id*]
no send-orf [*comm-id*]

Context config>router>bgp>outbound-route-filtering>extended-community
config>router>bgp>group>outbound-route-filtering>extended-community
config>router>bgp>group>neighbor>outbound-route-filtering>extended-community

Description	This command instructs the router to negotiate the send capability in the BGP outbound route filtering (ORF) negotiation with a peer. This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD. The no form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer. The no form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer. If the <i>comm-id</i> parameters are not exclusively route target communities then the router will extract appropriate route targets and use those. If, for some reason, the <i>comm-id</i> parameters specified contain no route targets, then the router will not send an ORF.
Default	no send-orf
Parameters	<i>comm-id</i> — Specifies up to 32 community policies, which must consist exclusively of route target extended communities. If it is not specified, then the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs and locally configured route targets. Values [target: {<i>ip-address:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i>} where: • <i>ip-address</i> — a.b.c.d • <i>comm-val</i> — 0 to 65535 • <i>2byte-asnumber</i> — 0 to 65535 • <i>ext-comm-val</i> — 0 to 4294967295 • <i>4byte-asnumber</i> — 0 to 4294967295

neighbor

Syntax	[no] neighbor <i>ip-address</i>
Context	config>router>bgp>group
Description	This command creates a BGP peer/neighbor instance within the context of the BGP group. This command can be issued repeatedly to create multiple peers and their associated configuration. The no form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively shutdown before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Default	no neighbor
Parameters	<i>ip-address</i> — Specifies the IP address of the BGP peer router in dotted decimal notation.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum, mandatory for link local addresses

advertise-ldp-prefix

Syntax	[no] advertise-ldp-prefix
Context	config>router>bgp>group>neighbor
Description	This command, when configured for a session that supports the IPv4 labeled-unicast address family, allows (subject to BGP export policies) active /32 LDP FEC prefixes to be advertised to the BGP peer with an RFC 3107 label, even though there may be BGP paths for the same prefix.
Default	no advertise-ldp-prefix

next-hop-self

Syntax	[no] next-hop-self
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures BGP to advertise routes to members of a group or to a specific neighbor using a local address of the BGP instance as the BGP next-hop address. Note that next-hop-self is set without exception, regardless of the route source (EBGP or IBGP) or its family. When used with VPN-IPv4 and VPN-IPv6 routes the enable-rr-vpn-forwarding command should also be configured. The no form of the command uses protocol standard behavior to decide whether or not to set next-hop-self in advertised routes.
Default	no next-hop-self

next-hop-unchanged

Syntax	next-hop-unchanged [label-ipv4] [label-ipv6] no next-hop-unchanged
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables unchanged BGP next-hops when sending BGP routes to peers in this group. The no form of the command disables unchanged BGP next-hops.
Default	no next-hop-unchanged

override-tunnel-elc

Syntax	[no] override-tunnel-elc
Context	config>router>bgp
Description	This command enables or disables entropy label capability (ELC) on BGP tunnels. When this command is enabled, the system assumes that all far ends for BGP tunnels are entropy-label-capable, regardless of any received capability signaling. This ensures that the entropy label will be inserted on BGP tunnels in the absence of capability signaling support by the far end. This is a system-wide configuration, since efficient entropy label operation requires that all LSRs in a network support entropy labels. This command should be used with care, particularly in inter-AS use cases, since entropy label capability may differ between domains.
Default	no override-tunnel-elc

passive

Syntax	[no] passive
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	Enables/disables passive mode for the BGP group or neighbor. When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer. The no form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no passive

peer-as

Syntax **peer-as** *as-number*

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Parameters *as-number* — Specifies the autonomous system number expressed as a decimal integer.

Values 1 to 4294967295

path-mtu-discovery

Syntax [**no**] **path-mtu-discovery**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables path MTU discovery for the associated TCP connections. In doing so, the MTU for the associated TCP session will be initially set to the egress interface MTU. The DF bit will also be set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it will send back an ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.

The **no** form of the command disables path MTU discovery.

Default no path-mtu-discovery

preference

Syntax	[no] preference <i>preference</i>
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the route preference for routes learned from the configured peers. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The lower the preference the higher the chance of the route being the active route. The router assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF. The no form of the command used at the global level reverts to default value. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.
Default	preference 170
Parameters	<i>preference</i> — Specifies the route preference expressed as a decimal integer. Values 1 to 255

purge-timer

Syntax	purge-timer <i>minutes</i> no purge-timer
Context	config>router>bgp
Description	When the system sends a VPN-IP Route-Refresh to a peer it sets all the VPN-IP routes received from that peer (in the RIB-IN) to stale and starts the purge-timer. If the routes are not updated (refreshed) before the purge-timer has expired then the routes are removed. The BGP purge timer configures the time before stale routes are purged. The no form of the command reverts to the default.
Default	purge-timer 10

Parameters	<i>minutes</i> — Specifies the maximum time before stale routes are purged.
Values	1 to 60

rapid-update

Syntax	rapid-update [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [evpn] [label-ipv4] [label-ipv6] [vpn-ipv4] [vpn-ipv6] [mcast-vpn-ipv4] [mcast-vpn-ipv6] no rapid-update
Context	config>router>bgp
Description	This command enables and disables BGP rapid update for specified address families. If rapid update is enabled for a set of address families, and a route belonging to a family in that set is received by the router and chosen for propagation to certain BGP peers, the remaining time on the MRAI timer of these peers is ignored and the route is transmitted immediately, along with all other pending routes for these peers (including routes of address families not specified in the rapid-update command). The rapid-update command overrides the peer-level min-route-advertisement time and applies the minimum setting (0 seconds) to routes belonging to specified address families; routes of other address families continue to be advertised according to the session-level MRAI setting. The no form of the command disables rapid update for all address families.
Default	no rapid-update
Parameters	l2-vpn — Specifies the BGP rapid update for the 12-byte Virtual Switch Instance identifier (VSI-ID) value consisting of the 8-byte route distinguisher (RD) followed by a 4-byte value. mvpn-ipv4 — Specifies BGP rapid update for the mvpn-ipv4 address family. The mvpn-ipv4 address is a variable size value consisting of the 1-byte route type, 1-byte length and variable size that is route type specific. Route type defines encoding for the route type specific field. Length indicates the length in octets of the route type specific field. mdt-safi — Specifies BGP rapid update for the mdt-safi address family. The address is a 16-byte value consisting of 12-byte route distinguisher (RD) followed by a 4-byte group address. mvpn-ipv6 — Specifies BGP rapid update for the mvpn-ipv6 address family. evpn — Specifies BGP rapid update for the evpn address family by including or removing EVPN routes from the set of routes that can trigger rapid update. label-ipv4 — Includes or removes label-ipv4 routes from the set of routes that can trigger rapid update. label-ipv6 — Includes or removes label-ipv6 routes from the set of routes that can trigger rapid update.

vpn-ipv4 — Includes or removes vpn-ipv4 routes from the set of routes that can trigger rapid update.

vpn-ipv6 — Includes or removes vpn-ipv6 routes from the set of routes that can trigger rapid update.

mcast-vpn-ipv4 — Includes or removes mcast-vpn-ipv4 routes from the set of routes that can trigger rapid update.

mcast-vpn-ipv6 — Includes or removes mcast-vpn-ipv6 routes from the set of routes that can trigger rapid update.

rapid-withdrawal

Syntax	[no] rapid-withdrawal
Context	config>router>bgp
Description	This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates. The no form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.
Default	no rapid-withdrawal

prefix-limit

Syntax	prefix-limit <i>family limit</i> [threshold percentage] [idle-timeout {<i>minutes</i> forever} log-only] [post-import] no prefix-limit <i>family</i>
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures the maximum number of BGP routes that can be received from a peer before some administrative action is taken. The administrative action can be the generation of a log event or taking down the session. If a session is taken down, then it can be brought back up automatically after an idle-timeout period, or else it can be configured to stay down ('forever') until the operator performs a reset. No prefix limits for any address family are configured by default. The prefix-limit command allows each address family to have its own limit; a set of address family limits can be applied to one neighbor or to all neighbors in a group. The no form of the command removes the prefix-limit.

Parameters	log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down. <i>percentage</i> — Specifies the threshold value (as a percentage) that triggers a warning message to be sent. Values 1 to 100 <i>family</i> — Specifies the address family to which the limit applies. Values ipv4, label-ipv4, vpn-ipv4, ipv6, label-ipv6, vpn-ipv6, mcast-ipv4, l2-vpn, mvpn-ipv4, mdt-safi, ms-pw, flow-ipv4, route-target, mcast-vpn-ipv4, mvpn-ipv6, flow-ipv6, evpn, mcast-ipv6, bgp-ls <i>limit</i> — Specifies the number of routes that can be learned from a peer expressed as a decimal integer. Values 1 to 4294967295 <i>minutes</i> — Specifies duration in minutes before automatically re-establishing a session. Values 1 to 1024 forever — Specifies that the session is reestablished only after clear router bgp command is executed. post-import — Specifies that the limit should be applied only to the number of routes that are accepted by import policies.

remove-private

Syntax	remove-private [limited] [skip-peer-as] no remove-private
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers. When the remove-private parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration. The set of AS numbers that are defined by IANA as private are in the range of 64512 through 65535, and 4200000000-4294967295, inclusive. The no form of the command used at the global level reverts to default value. The no form of the command used at the group level reverts to the value defined at the global level. The no form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no remove-private
Parameters	limited — This optional keyword removes private ASNs up to the first public ASN encountered. It then stops removing private ASNs. skip-peer-as — This optional keyword causes this command to not remove a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

rib-management

Syntax	rib-management
Context	config>router>bgp
Description	This command enables the context to configure RIB management parameters.

router-id

Syntax	router-id <i>ip-address</i> no router-id
Context	config>router>bgp
Description	This command specifies the router ID to be used with this BGP instance. Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID. It is possible to configure an SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address. When no router-id is configured for BGP, the system interface IP address is used.
Default	no router-id
Parameters	<i>ip-address</i> — Specifies the router ID, expressed as any non-zero value in the range 0.0.0.1 to 255.255.255.255 (or when converted to decimal it can have any value in the range 1-4294967295). It is recommended to use the system IPv4 address.

split-horizon

Syntax	[no] split-horizon
Context	config>router>bgp config>router>bgp>group

config>router>bgp>group>neighbor

Description This command enables the use of split-horizon. Split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGP and IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.

Default no split-horizon

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer.

The **no** form of the command disables TTL security.

Default no ttl-security

Parameters *min-ttl-value* — Specifies the minimum TTL value for an incoming packet.

Values 1 to 255

Default 1

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description **This command designates** the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the router derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default	no type
Parameters	internal — Configures the peer as internal. external — Configures the peer as external.

5.13.2.3 BGP BMP Commands

bmp

Syntax	bmp
Context	config
Description	This command enables the context to configure BGP Monitoring Protocol (BMP) parameters.

station

Syntax	station <i>station-name</i> [create] no station <i>station-name</i>
Context	config>bmp
Description	The command configures the BMP monitoring station name. The no form of the command removes the station name from the configuration.
Parameters	<i>station-name</i> — Specifies the station name of the BMP monitoring station up to 32 characters. create — Keyword used to create the station name. The create keyword requirement can be enabled or disabled in the environment>create context.

connection

Syntax	connection
Context	config>bmp>station
Description	This command enables the context to configure connection parameters for the BMP monitoring station.

connect-retry

Syntax	connect-retry <i>seconds</i> no connect-retry
Context	config>bmp>station>connection
Description	This command configures the BMP connect retry timer value. When this timer expires, BMP tries to reconnect to the configured monitoring station. This timer is applicable when the connection to the monitoring station is not yet established. The no form of the command reverts to the default value.
Default	connect-retry 120
Parameters	<i>seconds</i> — Specifies the BMP connect retry timer in seconds. Values 1 to 65535

local-address

Syntax	local-address <i>ip-address</i> <i>ipv6-address</i> no local-address
Context	config>bmp>station>connection
Description	This command configures the local IP address used by the local router when communicating with the BMP monitoring station. This configuration is optional. Outgoing connections use the local-address as the source of the TCP connection when initiating connections with a monitoring station. The BMP session may flap when this parameter is changed. Shut down the BMP session before changing the values. The no form of the command removes the configured local-address for the BMP session. The default is to use the system IP address.
Default	local-address ip-address (system IP address)
Parameters	<i>ip-address</i> — Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. Values ipv4-address: • a.b.c.d (host bits must be 0)

ipv6-address — Specifies the local address expressed in dotted decimal notation. Allowed value is a valid routable IPv6 address on the router, either an interface or system IPv6 address.

Values *ipv6-address*:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

router

Syntax **router service-name** *service-name*
router *router-instance*
no router

Context config>bmp>station>connection

Description This command configures the router instance to be used to connect to the associate BMP monitoring station.

The **no** form of the command removes the parameters from the configuration.

Parameters *service-name* — Specifies the name associated with the VPRN service through which the BMP monitoring station connection should traverse.

router-instance — Specifies the routing instance where the lead pool resides.

Values *router-name* | *vprn-service-id*
router-name: "Base" Default - Base
vprn-svc-id: 1 to 2147483647
service-name: The service name up to 64 characters in length.

station-address

Syntax **station-address** *ip-address* | *ipv6-address* **port** *port*
no station-address

Context config>bmp>station>connection

Description This command configures the IP address and TCP port number of the remote BMP monitoring station. This is a mandatory parameter and must be configured before the associated station can transitioned out of the shut down state.

The **no** form of the command removes the configured station IP address and port number for the BMP session. The **no station-address** command cannot be accepted unless the BMP or station instance is shut down.

Parameters	<i>ip-address</i> — Specifies the station address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. Values ipv4-address: • a.b.c.d (host bits must be 0) <i>ipv6-address</i> — Specifies the station address expressed in dotted decimal notation. Allowed value is a valid routable IPv6 address on the router, either an interface or system IPv6 address. Values ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D <i>port</i> — Specifies the TCP (destination) port number to be used when establishing the connection to the associated BMP station. Values 1 to 65535

tcp-keepalive

Syntax	tcp-keepalive
Context	config>bmp>station>connection
Description	This command enables the context to configure TCP keepalive parameters for the station.

keep-count

Syntax	keep-count <i>count</i> no keep-count
Context	config>bmp>station>connection>tcp-keepalive
Description	This command configures the number of missed keepalives before the TCP connection is declared down. The no form of the command reverts to the default.
Default	keep-count 4
Parameters	<i>count</i> — Specifies the number of missed keepalives before the TCP connection is declared down. Values 3 to 100

keep-idle

Syntax	keep-idle <i>idle</i> no keep-idle
Context	config>bmp>station>connection>tcp-keepalive
Description	This command configures the time until the first TCP keepalive probe is sent. The no form of the command reverts to the default.
Default	keep-idle 600
Parameters	<i>idle</i> — Specifies the time, in seconds, until the first TCP keepalive probe is sent. Values 1 to 100000

keep-interval

Syntax	keep-interval <i>interval</i> no keep-interval
Context	config>bmp>station>connection>tcp-keepalive
Description	This command configures the time between two TCP keepalives probes. The no form of the command reverts to the default.
Default	keep-interval 15
Parameters	<i>interval</i> — Specifies the time, in seconds, between two TCP keepalives probes. Values 1 to 100000

family

Syntax	family [ipv4] [ipv6] [label-ipv4] [label-ipv6] [vpn-ipv4] [vpn-ipv6] no family
Context	config>bmp>station
Description	This command configures the address families that are reported to a BMP monitoring station. The no form of the command reverts to the default.
Default	family ipv4
Parameters	ipv4 — Adds support for the IPv4 unicast (unlabeled) address family. ipv6 — Adds support for the IPv6 unicast (unlabeled) address family.

label-ipv4 — Adds support for the IPv4 unicast (labeled) address family.

label-ipv6 — Adds support for the IPv6 unicast (labeled) address family.

vpn-ipv4 — Adds support for the IPv4 VPN (SAFI 128) address family.

vpn-ipv6 — Adds support for the IPv6 VPN (SAFI 128) address family.

initiation-message

Syntax	initiation-message [<i>initiation-message</i>] no initiation-message
Context	config>bmp>station
Description	This command configures a free-form initiation message for a type 0 TLV to be sent to the BMP monitoring station. The message is transmitted when a BMP monitoring station establishes a connection to the device. Information can be provided to the BMP station system administrator (for example, a contact phone number). The initiation message includes a type 1 TLV containing the SNMP sysDescr value specified in RFC 1213, <i>Management Information Base for Network Management of TCP/IP-based internets: MIB-II</i>, and a type 2 TLV containing the SNMP sysName value also from RFC 1213. The string in the initiation-message is UTF-8 encoded. The no form of the command removes initiation message from the configuration and causes a free-form message to be included in the type 0 information TLV and the corresponding tlv-length is made 0.
Parameters	<i>initiation-message</i> — Specifies an initiation message up to 256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

report-local-routes

Syntax	[no] report-local-routes
Context	config>bmp>station
Description	This command enables local route reporting to the BMP monitoring station. The no form of the command disables the local route reporting.

stats-report-interval

Syntax	stats-report-interval [<i>seconds</i>] no stats-report-interval

Context	config>bmp>station
Description	This command configures the frequency of sending statistics reporting messages to the BMP monitoring station The no form of the command removes the interval from the configuration.
Parameters	<i>seconds</i> — Specifies the frequency of sending statistics reporting messages, in seconds, to the BMP monitoring station. Values 15 to 65535

monitor

Syntax	[no] monitor
Context	config>router>bgp>monitor config>router>bgp>group config>router>bgp>group>neighbor config>service>vrpn>bgp>group config>service>vrpn>bgp>group>neighbor
Description	This command enables the context to configure monitor parameters. The no form of the command disables BGP BMP monitoring.

route-monitoring

Syntax	route-monitoring pre-policy [post-policy] route-monitoring post-policy no route-monitoring
Context	config>router>bgp>monitor config>router>bgp>group>monitor config>router>bgp>group>neighbor>monitor config>service>vrpn>bgp>monitor config>service>vrpn>bgp>group>monitor config>service>vrpn>bgp>group>monitor
Description	This command specifies if BMP sends pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none. The no form of the command disables sending of route-monitoring messages.
Parameters	pre-policy — Enables sending pre-policy route monitoring messages using the pre-policy path attribute values, if available. post-policy — Enables sending post-policy route monitoring messages using the post-policy path attribute values, if available.

station

Syntax	station all station <i>name</i> [<i>name</i>] no station
Context	config>router>bgp>group>monitor config>router>bgp>group>neighbor>monitor config>router>bgp>monitor config>service>vprn>bgp>group>monitor config>service>vprn>bgp>group>neighbor>monitor
Description	This command configures the set of BMP monitoring stations for which BMP messages are to be sent, at the global BGP instance level, per group or for a particular neighbor. Whatever value is configured for the station parameter at the most specific BGP hierarchy level is used. • If a station list or the no station command is configured at a neighbor context, then that value is used. • If no station command is configured at the neighbor context, the group value is used. • If a station list or the no station command is configured at a group context, then that value is used. • If no station command is configured at the group context, the global value is used. • If a station list or the no station command is configured at the global context, then that value is used. • If no station command is configured at the global context, then a no station is assumed. The no form of the command disables sending BMP messages to BMP monitoring stations.
Parameters	<i>name</i> — Specifies up to eight station names up to 32 characters. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. all — Specifies all configured stations.

5.13.2.4 Other BGP-Related Commands

autonomous-system

Syntax	autonomous-system <i>autonomous-system-number</i> no autonomous-system

Context	config>router
Description	This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself. If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (shutdown/no shutdown) the BGP instance or rebooting the system with the new configuration.
Default	no autonomous-system
Parameters	<i>autonomous-system-number</i> — Specifies the autonomous system number, expressed as a decimal integer. Values 1 to 4294967295

confederation

Syntax	confederation <i>confed-as-num</i> [members <i>as-number</i> [<i>as-number</i>]] no confederation no confederation <i>confed-as-num</i> members <i>as-number</i> [<i>as-number</i>]
Context	config>router
Description	This command creates confederation autonomous systems within an AS. This technique is used to reduce the number of IBGP sessions required within an AS. Route reflection is the other technique that is commonly deployed to reduce the number of IBGP sessions. The no form of the command deletes the specified member AS from the confederation. When members are not specified in the no statement, the entire list is removed and confederations is disabled. When the last member of the list is removed, confederations is disabled.
Default	no confederation
Parameters	<i>confed-as-num</i> — Specifies the confederation AS number expressed as a decimal integer. Values 1 to 65535 <i>as-num</i> — Specifies the AS number(s) of members that are part of the confederation expressed as a decimal integer. Configure up to 15 members per <i>confed-as-num</i>.

router-id

Syntax	router-id <i>router-id</i> no router-id
Context	config>router
Description	This command configures the router ID for the router instance. The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID. When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs. To force the new router ID to be used, issue the shutdown and no shutdown commands for each protocol that uses the router ID, or restart the entire router. The system uses the system interface address (which is also the loopback address). If a system interface address is not configured, use the last 32 bits of the chassis MAC address. The no form of the command to reverts to the default value.
Default	no router-id
Parameters	<i>router-id</i> — Specifies the 32 bit router ID, expressed in dotted decimal notation or as a decimal value.

5.14 Show, Clear, and Debug Command Reference

5.14.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

5.14.1.1 Show Commands

```
show
  — router [router-instance]
    — bgp
      — auth-keychain keychain-name
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] [ipv4]
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] ipv6
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] label-ipv4
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] label-ipv6
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] mcast-ipv4
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] mcast-ipv6
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] mvpn-ipv4
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] vpn-ipv4
      — damping [ip-prefix[ip-prefix-length]] [damp-type] [detail] vpn-ipv6
      — group [name] [detail]
      — inter-as-label [next-hop ip-address | backup-next-hop ip-address]
        [received-label received-label]
      — neighbor [ip-address [detail]]
      — neighbor [as-number [detail]]
      — neighbor ip-address [family [type mvpn-type] [evpn-type]] filter1 [brief]
        [community comm-id]
      — neighbor ip-address [family] filter2
      — neighbor as-number [family] filter2
      — neighbor ip-address orf [filter3]
      — neighbor ip-address graceful-restart
      — neighbor [dynamic]
      — next-hop [family [service-id service-id]] [ip-address [detail]]
      — optimal-route-reflection [bgp-nh-info [location location-id] [bgp-nh ip-
        address]]
      — paths
      — policy-test plcy-or-long-expr [plcy-or-expr [plcy-or-expr.] [family family] [rd rd]
        prefix ip-prefix[ip-prefix-length] [longer] [neighbor ip-address] [display-
        rejects] [detail | brief]
      — prefix-label
      — routes [family] [brief]
      — routes [family] prefix [detail | longer | hunt [brief]]
```

- **routes** *family* [*prefix*] [**detail** | **longer** | **hunt** [**brief**]]
- **routes** [*family* [**type** *mvpn-type*]] **community** *comm-id*
- **routes** [*family* [**type** *mvpn-type*]] **aspath-regex** *reg-ex*
- **routes** *vpn-ipv4* *prefix* [**rd** *rd*] [**detail** | **longer** | **hunt** [**brief**]]
- **routes** *vpn-ipv6* *prefix* [**rd** *rd*] [**detail** | **longer** | **hunt** [**brief**]]
- **routes** *mvpn-ipv4* **type** *mvpn-type* {**rd** *rd* | **originator-ip** *ip-address* | **source-ip** *ip-address* | **group-ip** *ip-address* | **source-as** *as-number*} [**hunt** | **detail**]
- **routes** [*family* [**l2vpn-type**]] [**brief**]
- **routes** [*family* [**l2vpn-type**]] **community** *comm-id*
- **routes** [*family* [**l2vpn-type**]] **aspath-regex** *reg-ex*
- **routes** *evpn* {[**inclusive-mcast**] | [**ip-prefix**] | [**mac**]}
- **routes** *evpn* *inclusive-mcast* [**hunt** | **detail**] [**rd** *rd*] [**originator-ip** *ip-address*] [**next-hop** *ip-address*] [**community** *comm-id*] [**tag** *vni-id*]
- **routes** *evpn* *ip-prefix* [**hunt** | **detail**] [**rd** *rd*] [**prefix** *ip-prefix/mask*] [**community** *comm-id*] [**tag** *vni-id*] [**next-hop** *ip-address*]
- **routes** *evpn* *mac* [**hunt** | **detail**] [**rd** *rd*] [**next-hop** *ip-address*] [**mac-address** *mac-address*] [**community** *comm-id*] [**tag** *vni-id*]
- **routes** *l2-vpn* *l2vpn-type* {[**rd** *rd*] | [**siteid** *site-id*] | [**veid** *veid*] [**offset** *vpls-base-offset*]}
- **routes** *bgp-ls* [**hunt**] [**node** | **link** | **ipv4-prefix** [*ipv4-prefix/mask-len*]]
- **routes** *mdt-safi* [**rd** *rd*] [**grp-address** *mcast-grp-address*] [**brief**]
- **routes** *ms-pw* [**rd** *rd*] [**aii-type2** *aii-type2*] [**brief**]
- **routes** *flow-ipv4*
- **routes** *route-target* [**source-as** *as-number*] [**brief**] [**aspath-regex** *reg-exp*] [**community** *comm-id*]
- **routes** *route-target* [**rtc-prefix** *rtc-prefix*] [**hunt**] [**brief**] [**aspath-regex** *reg-exp*] [**community** *comm-id*]
- **routes** *route-target* *rtc-prefix* *rtc-prefix* [**aspath-regex** *reg-exp*] [**community** *comm-id*]
- **routes** *route-target* [**rtc-prefix** *rtc-prefix*] [**detail** | **longer**] [**aspath-regex** *reg-exp*] [**community** *comm-id*]
- **summary** [**all**]
- **summary** [*family* *family*] [**neighbor** *ip-address*]
- **fib** *slot-number* [**family**] [*ip-prefix/prefix-length*] [**longer**] [**secondary**] [**qos**] [**accounting-class**] [**all**]
- **fib** *slot-number* **extensive** [*ip-prefix/prefix-length*] [**family**] [**all**]
- **fib** *slot-number* [**family**] **summary**
- **fib** *slot-number* **nh-table-usage**
- **fib** **all** **summary**
- **route-table** [**family**] [*ip-prefix/prefix-length*] [**longer** | **exact** | **protocol** *protocol-name*] [**all**] [**next-hop-type** *type*][**qos**][**alternative**] [**accounting-class**]
- **route-table** [**family**] **summary**
- **route-table** **tunnel-endpoints** [*ip-prefix/prefix-length*] [**longer** | **exact**] [**detail**]
- **route-table** [**family**] [*ip-prefix/prefix-length*] [**longer** | **exact** | **protocol** *protocol-name*] **extensive** [**all**]
- **bmp**
 - **station** [*station-name*]

5.14.1.2 Clear Commands

clear

- router
 - bgp
 - **damping** [{*ip-prefix/ip-prefix-length*] [**neighbor** *ip-address*]} | **group** *name*
 - **flap-statistics** [{*ip-prefix/mask* [**neighbor** *ip-address*] | [**group** *group-name*] | **regex** *reg-exp* | **policy** *policy-name*}]
 - **neighbor** {*ip-address* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound** | **hard**]
 - **neighbor** {*ip-address* | **as** *as-number* | **external** | **all**} **statistics**
 - **neighbor** *ip-address* **end-of-rib**
 - **protocol**
 - bmp
 - **station** *station-name*
 - **station** **all**

5.14.1.3 Debug Commands

- debug
 - router
 - bgp
 - **events** [**neighbor** *ip-address* | **group** *name*]
 - **no events**
 - **graceful-restart** [**neighbor** *ip-address* | **group** *name*]
 - **no graceful-restart**
 - **keepalive** [**neighbor** *ip-address* | **group** *name*]
 - **no keepalive**
 - **notification** [**neighbor** *ip-address* | **group** *name*]
 - **no notification**
 - **open** [**neighbor** *ip-address* | **group** *name*]
 - **no open**
 - [**no**] **outbound-route-filtering**
 - **packets** [**neighbor** *ip-address* | **group** *name*]
 - **no packets**
 - **route-refresh** [**neighbor** *ip-address* | **group** *name*]
 - **no route-refresh**
 - **rtm** [**neighbor** *ip-address* | **group** *name*]
 - **no rtm**
 - **socket** [**neighbor** *ip-address* | **group** *name*]
 - **no socket**
 - **timers** [**neighbor** *ip-address* | **group** *name*]\
 - **no timers**
 - **update** [**neighbor** *ip-address* | **group** *name*]
 - **no update**
 - bmp
 - **events** [**station** *station-name*]
 - **no events**
 - **packets** [**station** *station-name*]
 - **no packets**

5.14.2 Command Descriptions

5.14.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

router

Syntax	router [<i>router-instance</i>]				
Context	show				
Description	Displays router instance information.				
Parameters	<i>router-instance</i> — Specify either the router-name or service-id. <table><tr><td>Values</td><td>router-name: Base, management service-id: 1 to 2147483647</td></tr><tr><td>Default</td><td>Base</td></tr></table>	Values	router-name: Base, management service-id: 1 to 2147483647	Default	Base
Values	router-name: Base, management service-id: 1 to 2147483647				
Default	Base				

bgp

Syntax	bgp
Context	show>router
Description	Enables the context to display BGP related information.

auth-keychain

Syntax	auth-keychain [<i>keychain</i>]
Context	show>router>bgp show>router>bgp>group show>router>bgp>group>neighbor
Description	This command displays BGP sessions using a particular authentication keychain.
Parameters	<i>keychain</i> — Specifies an existing keychain name.
Output	The following output is an example of authentication keychain information.

Sample Output

```
*A:ALA-48# show router 2 bgp auth-keychain
=====
Sessions using key chains
=====
Peer address          Group          Keychain name
-----
10.20.1.3             1             eta_keychain1
30.1.0.2              1             eta_keychain1
=====
*A:ALA-48#
*A:ALA-48>config>router>bgp# show router bgp group "To_AS_10000"
=====
BGP Group : To_AS_10000
-----
Group          : To_AS_10000
-----
Group Type      : No Type          State          : Up
Peer AS         : 10000             Local AS       : 200
Local Address   : n/a              Loop Detect    : Ignore
Import Policy   : None Specified / Inherited
Export Policy    : ospf3
Hold Time       : 90              Keep Alive     : 30
Cluster Id      : 0.0.0.100      Client Reflect : Enabled
NLRI            : Unicast         Preference     : 170
TTL Security    : Disabled       Min TTL Value  : n/a
Graceful Restart : Enabled       Stale Routes Time: 360
Auth key chain  : testname

List of Peers
- 10.0.0.8 :
    To_Router B - EBGp Peer
Total Peers   : 1              Established    : 0
-----
Peer Groups : 1
=====
*A:ALA-48>config>router>bgp#

*A:ALA-48>config>router>bgp# show router bgp neighbor 10.0.0.8
=====
BGP Neighbor
-----
Peer   : 10.0.0.8
Group  : To_AS_10000
-----
Peer AS      : 10000          Peer Port     : 0
Peer Address : 10.0.0.8       Local Port    : 0
Local AS     : 200
Local Address : 0.0.0.0
Peer Type    : External
State        : Active        Last State     : Idle
Last Event   : stop
Last Error   : Cease
Local Family : IPv4
Remote Family : Unused
Hold Time    : 90             Keep Alive     : 30
Active Hold Time : 0         Active Keep Alive : 0
```

```

Cluster Id           : 0.0.0.100
Preference           : 99
Recd. Paths          : 0
IPv4 Recd. Prefixes  : 0
IPv4 Suppressed Pfxs : 0
VPN-IPv4 Recd. Pfxs  : 0
Mc IPv4 Recd. Pfxs.  : 0
Mc IPv4 Suppr. Pfxs  : 0
IPv6 Recd. Prefixes  : 0
Input Queue          : 0
i/p Messages         : 0
i/p Octets           : 0
i/p Updates          : 0
TTL Security         : Disabled
Graceful Restart     : Enabled
Advertise Inactive   : Disabled
Advertise Label       : None
Auth key chain        : testname
Local Capability      : RouteRefresh MP-BGP
Remote Capability     :
Import Policy         : None Specified / Inherited
Export Policy         : ospf3
Num of Flaps          : 0
IPv4 Active Prefixes : 0
VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Active Pfxs : 0
Mc IPv4 Active Pfxs. : 0
IPv6 Suppressed Pfxs : 0
IPv6 Active Prefixes : 0
Output Queue         : 0
o/p Messages         : 0
o/p Octets           : 0
o/p Updates          : 0
Min TTL Value        : n/a
Stale Routes Time    : 360
Peer Tracking        : Disabled

```

```

-----
Neighbors : 1
=====

```

```

*A:ALA-48>config>router>bgp#

```

```

*A:ALA-48>config>router>bgp# show router bgp auth-keychain testname

```

```

=====
Sessions using key chain: keychain
=====

```

Peer address	Group	Keychain name
10.0.0.8	To_AS_10000	testname

```

=====
*A:ALA-48>config>router>bgp#

```

damping

Syntax	damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] [ipv4] damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] ipv6 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] label-ipv4 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] label-ipv6 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] mcast-ipv4 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] mcast-ipv6 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] mvpn-ipv4 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] vpn-ipv4 damping [<i>ip-prefix</i>[/<i>ip-prefix-length</i>]] [<i>damp-type</i>] [detail] vpn-ipv6
Context	show>router>bgp
Description	This command displays BGP routes which have been dampened due to route flapping. This command can be entered with or without a route parameter. When the keyword detail is included, more detailed information displays. When only the command is entered (without any parameters included except detail), then all dampened routes are listed. When a parameter is specified, then the matching route or routes are listed. When a decayed, history, or suppressed keyword is specified, only those types of dampened routes are listed.
Parameters	<i>ip-prefix</i> — Displays damping information for the specified IP prefix and length. Values ipv4-prefix: • a.b.c.d (host bits must be 0) ipv4-prefix-length: [0 to 32] ipv6-prefix: • x:x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D ipv6-prefix-length: [0 to 128] <i>damp-type</i> — Specifies the type of damping to display. Values decayed — Displays damping entries that are decayed but are not suppressed. history — Displays damping entries that are withdrawn but have history. suppressed — Displays damping entries suppressed because of route damping. detail — Displays detailed information.

ipv4 — Displays IPv4 information.

ipv6 — Displays IPv6 information.

label-ipv4 — Displays for labeled-unicast IPv4 routes.

label-ipv6 — Displays for labeled-unicast IPv6 routes.

vpn-ipv4 — Displays VPN IPv4 information.

vpn-ipv6 — Displays VPM IPv6 information.

Output The following output is an example of BGP damping information, and [Table 56](#) describes the output fields.

Sample Output

```
*A:ALA-12# show router 3 bgp damping
=====
BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
Flag  Network          From          Reuse          AS-Path
-----
ud*i  10.149.7.0/24       10.0.28.1     00h00m00s     60203 65001 19855 3356
                        1239  22406
si    10.155.6.0/23      10.0.28.1     00h43m41s     60203 65001 19855 3356
                        2914  7459
si    10.155.8.0/22      10.0.28.1     00h38m31s     60203 65001 19855 3356
                        2914  7459
si    10.155.12.0/22     10.0.28.1     00h35m41s     60203 65001 19855 3356
                        2914  7459
si    10.155.22.0/23     10.0.28.1     00h35m41s     60203 65001 19855 3356
                        2914  7459
si    10.155.24.0/22     10.0.28.1     00h35m41s     60203 65001 19855 3356
                        2914  7459
si    10.155.28.0/22     10.0.28.1     00h34m31s     60203 65001 19855 3356
                        2914  7459
si    10.155.40.0/21     10.0.28.1     00h28m24s     60203 65001 19855 3356
                        7911  7459
si    10.155.48.0/20     10.0.28.1     00h28m24s     60203 65001 19855 3356
                        7911  7459
ud*i  10.8.140.0/24       10.0.28.1     00h00m00s     60203 65001 19855 3356
                        4637  17447
ud*i  10.8.141.0/24       10.0.28.1     00h00m00s     60203 65001 19855 3356
                        4637  17447
ud*i  10.9.0.0/18        10.0.28.1     00h00m00s     60203 65001 19855 3356
                        3561  9658  6163
. . .
ud*i  10.213.184.0/23   10.0.28.1     00h00m00s     60203 65001 19855 3356
                        6774  6774  9154
-----
*A:ALA-12#
```


*A:ALA-12# show router 3 bgp damping detail

```
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
-----
Network : 10.149.7.0/24
-----
Network      : 10.149.7.0/24      Peer      : 10.0.28.1
NextHop      : 10.0.28.1        Reuse time : 00h00m00s
Peer AS     : 60203            Peer Router-Id : 10.32.27.203
Local Pref  : none
Age         : 00h22m09s        Last update : 02d00h58m
FOM Present : 738              FOM Last upd. : 2039
Number of Flaps : 2            Flags       : ud*i
Path        : 60203 65001 19855 3356 1239 22406
Applied Policy : default-damping-profile
-----
Network : 10.142.48.0/20
-----
Network      : 10.142.48.0/20    Peer      : 10.0.28.1
NextHop      : 10.0.28.1        Reuse time : 00h00m00s
Peer AS     : 60203            Peer Router-Id : 10.32.27.203
Local Pref  : none
Age         : 00h00m38s        Last update : 02d01h20m
FOM Present : 2011             FOM Last upd. : 2023
Number of Flaps : 2            Flags       : ud*i
Path        : 60203 65001 19855 3356 3561 5551 1889
Applied Policy : default-damping-profile
-----
Network : 10.200.128.0/19
-----
Network      : 10.200.128.0/19  Peer      : 10.0.28.1
NextHop      : 10.0.28.1        Reuse time : 00h00m00s
Peer AS     : 60203            Peer Router-Id : 10.32.27.203
Local Pref  : none
Age         : 00h00m38s        Last update : 02d01h20m
FOM Present : 2011             FOM Last upd. : 2023
Number of Flaps : 2            Flags       : ud*i
Path        : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.203.192.0/18
-----
Network      : 10.203.192.0/18  Peer      : 10.0.28.1
NextHop      : 10.0.28.1        Reuse time : 00h00m00s
Peer AS     : 60203            Peer Router-Id : 10.32.27.203
Local Pref  : none
Age         : 00h00m07s        Last update : 02d01h20m
FOM Present : 1018             FOM Last upd. : 1024
Number of Flaps : 1            Flags       : ud*i
Path        : 60203 65001 19855 1299 702 1889
```

```
Applied Policy      : default-damping-profile
-----
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping 10.203.192.0/18 detail
=====
      BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
      Legend -
      Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
      Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
      BGP Damped Routes 10.203.192.0/18
=====
      Network : 10.203.192.0/18
-----
      Network      : 10.203.192.0/18      Peer      : 10.0.28.1
      NextHop      : 10.0.28.1           Reuse time : 00h00m00s
      Peer AS      : 60203                Peer Router-Id : 10.32.27.203
      Local Pref   : none
      Age          : 00h00m42s            Last update : 02d01h20m
      FOM Present  : 2003                 FOM Last upd. : 2025
      Number of Flaps : 2                  Flags      : ud*i
      Path         : 60203 65001 19855 3356 702 1889
      Applied Policy : default-damping-profile
-----
      Paths : 1
=====
*A:ALA-12#
*A:ALA-12# show router 3 bgp damping suppressed detail
=====
      BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
      Legend -
      Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
      Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
      BGP Damped Routes (Suppressed)
=====
      Network : 10.142.48.0/20
-----
      Network      : 10.142.48.0/20      Peer      : 10.0.28.1
      NextHop      : 10.0.28.1           Reuse time : 00h29m22s
      Peer AS      : 60203                Peer Router-Id : 10.32.27.203
      Local Pref   : none
      Age          : 00h01m28s            Last update : 02d01h20m
      FOM Present  : 2936                 FOM Last upd. : 3001
      Number of Flaps : 3                  Flags      : si
      Path         : 60203 65001 19855 3356 702 1889
      Applied Policy : default-damping-profile
-----
      Network : 10.200.128.0/19
-----
      Network      : 10.200.128.0/19     Peer      : 10.0.28.1
      NextHop      : 10.0.28.1           Reuse time : 00h29m22s
      Peer AS      : 60203                Peer Router-Id : 10.32.27.203
      Local Pref   : none
      Age          : 00h01m28s            Last update : 02d01h20m
```

```

FOM Present      : 2936                      FOM Last upd.    : 3001
Number of Flaps  : 3                        Flags          : si
Path             : 60203 65001 19855 3356 702 1889
Applied Policy   : default-damping-profile
-----
Network : 10.203.240.0/20
-----
Network      : 10.203.240.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 10.206.0.0/17
-----
Network      : 10.206.0.0/17      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203              Peer Router-Id : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

```

Table 56 **Show Damping Field Descriptions**

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
Network	Route IP prefix and mask length for the route.
Flag(s)	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, then the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
Network	The IP prefix and mask length for the route.
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.

Table 56 Show Damping Field Descriptions (Continued)

Label	Description (Continued)
AS Path	The BGP AS path for the route.
Peer	The router ID of the advertising router.
NextHop	BGP nexthop for the route.
Peer AS	The autonomous system number of the advertising router.
Peer Router-Id	The router ID of the advertising router.
Local Pref	BGP local preference path attribute for the route.
Age	The time elapsed since the service was enabled.
Last update	The time when BGP was updated last in second/minute/hour (SS:MM:HH) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of flaps in the neighbor connection.
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

group

Syntax `group [name] [detail]`

Context `show>router>bgp`

Description This command displays group information for a BGP peer group. This command can be entered with or without parameters.

When this command is entered without a group name, information about all peer groups displays.

When the command is issued with a specific group name, information only pertaining to that specific peer group displays.

The 'State' field displays the BGP group's operational state. Valid states are:

Up — BGP global process is configured and running.

Down — BGP global process is administratively shutdown and not running.

Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters *name* — Displays information for the BGP group specified.
detail — Displays detailed information.

Output Standard and Detailed Group Output

The following table describes the standard and detailed command output fields for a BGP group.

Table 57 BGP Group Output Fields

Label	Description
Group	Displays the BGP group name.
Group Type	No Type — Peer type not configured. External — Peer type configured as external BGP peers. Internal — Peer type configured as internal BGP peers.
State	Disabled — The BGP peer group has been operationally disabled. Down — The BGP peer group is operationally inactive. Up — The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
Connect Retry	The configured or inherited connect retry timer value.
Authentication	None — No authentication is configured. MD5 — MD5 authentication is configured.
Bfd	Yes — BFD is enabled. No — BFD is disabled.
Local Pref	The configured or inherited local preference value.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.

Table 57 BGP Group Output Fields (Continued)

Label	Description
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group. 1 to 4294967295 — The maximum number of routes BGP can learn from a peer.
Passive	Disabled — BGP attempts to establish a BGP connection with neighbor in the specified peer group. Enabled — BGP will not actively attempt to establish a BGP connection with neighbor in the specified peer group.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to neighbors in the peer group. Enabled — BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group. Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group. Enabled BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
Damping	Disabled — The peer group is configured not to dampen route flaps. Enabled — The peer group is configured to dampen route flaps.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.

Table 57 BGP Group Output Fields (Continued)

Label	Description
Cluster Id	The configured route reflector cluster ID.
	None No cluster ID has been configured
Client Reflect	Disabled — The BGP route reflector will not reflect routes to this neighbor. Enabled — The BGP route reflector is configured to reflect routes to this neighbor.
NLRI	The type of NLRI information that the specified peer group can accept. Unicast — IPv4 unicast routing information can be carried.
Preference	The configured route preference value for the peer group.
List of Peers	A list of BGP peers configured under the peer group.
List of Static Peers	A list of statically configured BGP peers configured under the peer group.
List of Dynamic Peers	A list of dynamic BGP peers configured under the peer group.
Total Peers	The total number of peers configured under the peer group.
Established	The total number of peers that are in an established state.

Sample Output

```

A:ALA-12# show router bgp group
=====
BGP Groups
-----
Group           : To_AS_40000
-----
Description      : Not Available
Group Type       : No Type
Peer AS          : 40000
Local Address    : n/a
Export Policy    : direct2bgp
Hold Time        : 90
Cluster Id       : None
NLRI             : Unicast
State            : Up
Local AS         : 65206
Loop Detect      : Ignore
Keep Alive       : 30
Client Reflect   : Enabled
Preference       : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_SR1
- 10.0.0.15      : To_H-215

```

```

Total Peers      : 5                      Established      : 2
=====
A:ALA-12#

```

Sample Detailed Output

```

A:ALA-12# show router bgp group detail
=====
BGP Groups (detail)
-----
Group           : To_AS_40000
-----
Description      : Not Available
Group Type       : No Type
Peer AS          : 40000
Local Address    : n/a
Connect Retry    : 20
Local Pref       : 100
Multihop         : 0 (Default)
Min Route Advt.  : 30
Prefix Limit     : No Limit
Next Hop Self    : Disabled
Remove Private   : Disabled
Export Policy    : direct2bgp
Hold Time        : 90
Cluster Id       : None
NLRI             : Unicast

State            : Up
Local AS        : 65206
Loop Detect     : Ignore
Authentication  : None
MED Out         : 0

Min AS Originate : 15
Passive          : Disabled
Aggregator ID 0  : Disabled
Damping          : Disabled

Keep Alive      : 30
Client Reflect  : Enabled
Preference      : 170

List of Peers
- 10.0.0.1      : To_Jukebox
- 10.0.0.12     : Not Available
- 10.0.0.13     : Not Available
- 10.0.0.14     : To_SR1
- 10.0.0.15     : To_H-215

Total Peers      : 5                      Established      : 2
=====
A:ALA-12#

```

```

A:SetupCLI>show>router>bgp# group
=====
BGP Group
-----
Group           : bgp_group_1 34567890123456789012
-----
Description      : Testing the length of the group value for the DESCRIPTION
                  parameter of BGP
Group Type       : No Type
Peer AS          : n/a
Local Address    : n/a
Import Policy    : test i1
                  : test i2
                  : test i3
                  : test i4
                  : test i5 890123456789012345678901
Export Policy    : test e1
                  : test e2

State            : Up
Local AS        : 100
Loop Detect     : Ignore

```



```

: test e3
: test e4
: test e5 890123456789012345678901
Hold Time      : 120          Keep Alive      : 30
Cluster Id     : None         Client Reflect  : Disabled
NLRI           : Unicast      Preference     : 101
TTL Security   : Disabled     Min TTL Value : n/a
Graceful Restart : Disabled    Stale Routes Time: n/a
Auth key chain  : n/a         Bfd Enabled   : Yes

List of Peers
- 3.3.3.3 :
  Testing the length of the neighbor value for the DESCRIPTION parameter of
  BGP
Total Peers      : 1          Established      : 0
-----
Peer Groups : 1
=====
A:SetupCLI>show>router>bgp#
```

inter-as-label

Syntax	inter-as-label [next-hop <i>ip-address</i> backup-next-hop <i>ip-address</i>] [received-label <i>received-label</i>]
Context	show>router>bgp
Description	This command displays BGP inter-AS label information.
Parameters	<i>ip-address</i> — Displays information for the specified IP address. Values ipv4-address: • a.b.c.d (host bits must be 0) <i>received-label</i> — Displays information for the received label. Values 0 to 429496729

neighbor

Syntax	neighbor [<i>ip-address</i> [detail]] neighbor [<i>as-number</i> [detail]] neighbor <i>ip-address</i> [<i>family</i> [type <i>mvpn-type</i>] [<i>evpn-type</i>]] <i>filter1</i> [brief] [community <i>comm-id</i>] neighbor <i>ip-address</i> [<i>family</i>] <i>filter2</i> neighbor <i>as-number</i> [<i>family</i>] <i>filter2</i> neighbor <i>ip-address</i> orf [<i>filter3</i>] neighbor <i>ip-address</i> graceful-restart neighbor [dynamic]

Context show>router>bgp

Description This command displays BGP neighbor information and can be entered with or without any parameters. When this command is issued without any parameters, information is displayed for all BGP peers. When the command is issued with a specific IP address or ASN, information is displayed only for the specified peer or peers with the same AS.

When either **received-routes** or **advertised-routes** is specified, then the routes received from or sent to the specified peer is listed (see the second output example, below).

When either **history** or **suppressed** is specified, then the routes learned from those peers that either have a history or are suppressed (respectively) are listed.

The “State” field displays the BGP peer’s protocol state. In addition to the standard protocol states, this field can also display the “Disabled” operational state, which indicates the peer is operationally disabled and must be restarted by the operator.



Note: The information generated by this command is not available by SNMP.

Parameters *as-number* — Displays information for the specified AS number.

Values 1 to 4294967295

brief — Displays information in a brief format. This parameter is only supported with **received-routes** and **advertised-routes**.

community — Displays all routes with the specified BGP community.

comm-id — Specifies community IDs, up to 72 characters.

Values [*as-num:comm-val* | *ext-comm* | *well-known-comm*]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *ext-comm* — the extended community, defined as one of the following:

- {*target* | *origin*}:*ip-address:comm-val*
- {*target* | *origin*}:*asnum:ext-comm-val*
- {*target* | *origin*}:*ext-asnum:comm-val*
- **bandwidth**:*asnum:val-in-mbps*
- **ext:4300**:*ovstate*
- **ext**:*value1:value2*
- **flowspec-set**:*ext-asnum:group-id*

where:

- *target* — route target
- *origin* — route origin

- *ip-address* — a.b.c.d
- *ext-comm-val* — 0 to 4294967295
- *ext-asnum* — 0 to 4294967295
- **bandwidth** — bandwidth
- *val-in-mbps* — 0 to 16777215
- **ext** — extended
- **ext:4300** — origin verification
- *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
- *value1* — 0000 to FFFF
- *value2* — 0 to FFFFFFFFFF
- **flowspec-set** — flow-spec set
- *group-id* — 0 to 16383
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **llgr-stale** | **no-llgr** | **blackhole**

detail — Displays detailed information.

dynamic — Displays information for dynamic BGP neighbors.

evpn-type — Displays information for the specified EVPN route type.

- Values**
- auto-disc** — Displays the auto discovery routes for this peer.
 - mac** — Displays the MAC routes for this peer.
 - inclusive-mcast** — Displays the inclusive multicast routes for this peer.
 - eth-seg** — Displays the Ethernet segment routes for this peer.
 - ip-prefix** — Displays the IPv4 routes for this peer.
 - ipv6-prefix** — Displays the IPv6 routes for this peer.

family — Specifies the type of routing information to display.

- Values**
- evpn** — Displays the BGP EVPN routes for this peer.
 - flow-ipv4** — Displays the BGP IPv4 flow-spec routes for this peer.
 - flow-ipv6** — Displays the BGP IPv6 flow-spec routes for this peer.
 - ipv4** — Displays the BGP IPv4 routes for this peer.
 - ipv6** — Displays the BGP IPv6 routes for this peer.
 - l2-vpn** — Displays the BGP L2 VPN routes for this peer.
 - label-ipv4** — Displays the BGP IPv4 labeled unicast routes for this peer.
 - label-ipv6** — Displays the BGP IPv6 labeled unicast routes for this peer.
 - mcast-ipv4** — Displays the BGP multicast IPv4 routes for this peer.
 - mcast-ipv6** — Displays the BGP multicast IPv6 routes for this peer.

mcast-vpn-ipv4 — Displays the BGP multicast VPN IPv4 routes for this peer.

mcast-vpn-ipv6 — Displays the BGP multicast VPN Ipv6 routes for this peer.

mdt-safi — Displays the BGP multicast VPN MDT-SAFI routes for this peer.

ms-pw — Displays the BGP Multi-segment PW routes for this peer.

mvpn-ipv4 — Displays the BGP NG-MVPN IPv4 routes for this peer.

mvpn-ipv6 — Displays the BGP NG-MVPN IPv6 routes for this peer.

route-target — Displays the BGP route-target constrain routes for this peer.

vpn-ipv4 — Displays the BGP VPN IPv4 routes for this peer.

vpn-ipv6 — Displays the BGP VPN IPv6 routes for this peer.

filter1 — Display information for the specified IP address.

Values **received-routes** — Displays the number of routes received from this peer.
advertised-routes — Displays the number of routes advertised to this peer.

filter2 — Display information for the specified AS number.

Values **history** — Display statistics for dampened routes.
suppressed — Display the number of paths from this peer that have been suppressed by damping.

filter3 — Displays path information for the specified IP address.

Values **send** — Displays the number of paths sent to this peer.
receive — Displays the number of paths received from this peer.

graceful-restart — Displays neighbors configured for graceful restart.

ip-address — Display information for the specified IP address.

Values ipv4-address:
 • a.b.c.d (host bits must be 0)
 ipv6-address:
 • x:x:x:x:x:x:x [-interface]
 • x:x:x:x:x:d.d.d.d [-interface]
 • x: [0 to FFFF]H
 • d: [0 to 255]D
 • interface: 32 characters maximum, mandatory for link local addresses

mvpn-type — Displays information for the specified MVPN route type.

Values

- intra-ad** — Displays the intra-ad routes for this peer.
- inter-ad** — Displays the inter-ad routes for this peer.
- spmsi-ad** — Displays the spmsi-ad routes for this peer.
- leaf-ad** — Displays the leaf-ad routes for this peer.
- source-ad** — Displays the source-ad routes for this peer.
- shared-join** — Displays the shared-join routes for this peer.
- source-join** — Displays the source-join routes for this peer.

orf — Displays outbound route filtering for the BGP instance. Outbound Route Filtering (ORF) is used to inform a neighbor of targets (using **target-list**) that it is willing to receive. This mechanism helps lessen the update exchanges between neighbors and saves CPU cycles to process routes that could have been received from the neighbor only to be dropped or ignored.

Output Standard and Detailed Neighbor

The following table describes the standard and detailed command output fields for a BGP neighbor.

Table 58 BGP Neighbor Fields

Label	Description
Peer	The IP address of the configured BGP peer.
Group	The BGP peer group to which this peer is assigned.
Peer AS	The configured or inherited peer AS for the peer group.
Peer Address	The configured address for the BGP peer.
Peer Port	The TCP port number used on the far-end system.
Local AS	The configured or inherited local AS for the peer group.
Local Address	The configured or inherited local address for originating peering for the peer group.
Local Port	The TCP port number used on the local system.
Peer Type	External — Peer type configured as external BGP peers. Internal — Peer type configured as internal BGP peers.
Bfd	Yes — BFD is enabled. No — BFD is disabled.
Dynamic Peer	Yes — The session is dynamic (unconfigured). No — The session is statically configured.

Table 58 BGP Neighbor Fields (Continued)

Label	Description
State	Idle — The BGP peer is not accepting connections. (<i>Shutdown</i>) is displayed in addition, if the peer is administratively disabled. Active — BGP is listening for and accepting TCP connections from this peer. Connect — BGP is attempting to establish a TCP connections from this peer. Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer. Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION. Established — BGP has successfully established a peering and is exchanging routing information.
Last State	Idle — The BGP peer is not accepting connections. Active — BGP is listening for and accepting TCP connections from this peer. Connect — BGP is attempting to establish a TCP connections from this peer. Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer. Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
Last Event	start — BGP has initialized the BGP neighbor. stop — BGP has disabled the BGP neighbor. open — BGP transport connection opened. close — BGP transport connection closed. openFail — BGP transport connection failed to open. error — BGP transport connection error. connectRetry — Connect retry timer expired. holdTime — Hold time timer expired. keepAlive — Keepalive timer expired. recvOpen — Receive an OPEN message. revKeepalive — Receive a KEEPALIVE message. recvUpdate — Receive an UPDATE message. recvNotify — Receive a NOTIFICATION message. None — No events have occurred.
Last Error	Displays the last BGP error and subcode to occur on the BGP neighbor.

Table 58 BGP Neighbor Fields (Continued)

Label	Description
Connect Retry	The configured or inherited connect retry timer value.
Local Pref.	The configured or inherited local preference value.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Damping	Disabled — BGP neighbor is configured not to dampen route flaps. Enabled — BGP neighbor is configured to dampen route flaps.
Loop Detect	Ignore — The BGP neighbor is configured to ignore routes with an AS loop. Drop — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected. Off — AS loop detection is disabled for the neighbor.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Authentication	None — No authentication is configured. MD5 — MD5 authentication is configured.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor. Enabled — BGP will send only its own IP address as the BGP nexthop in route updates to the neighbor.
AggregatorID Zero	Disabled — The BGP Neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates. Enabled — The BGP Neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor. Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.

Table 58 BGP Neighbor Fields (Continued)

Label	Description
Passive	Disabled — BGP will actively attempt to establish a BGP connection with the specified neighbor. Enabled — BGP will not actively attempt to establish a BGP connection with the specified neighbor.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group. 1 to 4294967295 — The maximum number of routes BGP can learn from a peer.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state.
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state.
Cluster Id	The configured route reflector cluster ID. None — No cluster ID has been configured.
Client Reflect	Disabled — The BGP route reflector is configured not to reflect routes to this neighbor. Enabled — The BGP route reflector is configured to reflect routes to this neighbor.
Preference	The configured route preference value for the peer group.
Num of Flaps	The number of route flaps in the neighbor connection.
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.
i/p Messages	Total number of packets received from the BGP neighbor.
o/p Messages	Total number of packets sent to the BGP neighbor.

Table 58 BGP Neighbor Fields (Continued)

Label	Description
i/p Octets	Total number of octets received from the BGP neighbor.
o/p Octets	Total number of octets sent to the BGP neighbor.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.

Sample Output

```
A:ALA-48# show router bgp neighbor
=====
BGP Neighbor
-----
Peer : 10.0.0.5          Group : headquarters1
-----
Peer AS      : 300Peer Port      : 0
Peer Address : 10.0.0.5
Local AS     : 200Local Port    : 0
Local Address : 10.0.0.104
Peer Type    : External          Dynamic Peer    : Yes
State       : Active             Last State     : Idle
Last Event   : stop
Last Error   : Cease
Local Family : IPv4              Remote Family   : Unused
Hold Time    : 90                Keep Alive     : 30
Active Hold Time : 0            Active Keep Alive: 0
Cluster Id   : 0.0.0.100
Preference   : 170              Num of Flaps    : 0
Recd. Prefixes : 0              Active Prefixes : 0
Recd. Paths    : 0              Suppressed Paths : 0
Input Queue    : 0              Output Queue    : 0
i/p Messages   : 0              o/p Messages    : 0
i/p Octets     : 0              o/p Octets      : 0
i/p Updates    : 0              o/p Updates     : 0
TTL Security   : Enabled        Min TTL Value   : 255
Graceful Restart : Disabled     Stale Routes Time: n/a
Local Capability : RouteRefresh MP-BGP
Remote Capability:
Import Policy   : None Specified / Inherited
Export Policy   : None Specified / Inherited
-----
Peer : 10.0.0.91          Group : Santa Clara
-----
Peer AS      : 100Peer Port      : 0
Peer Address : 10.0.0.91
Local AS     : 200Local Port    : 0
Local Address : 10.0.0.103
Peer Type    : External
State       : Connect           Last State     : Active
Last Event   : openFail
Last Error   : Cease
Local Family : IPv4              Remote Family   : Unused
```

```

Hold Time      : 90
Active Hold Time : 0
Cluster Id     : 0.0.0.100
Preference     : 170
Recd. Prefixes : 0
Recd. Paths    : 0
Input Queue    : 0
i/p Messages   : 0
i/p Octets     : 0
i/p Updates    : 0
TTL Security   : Disabled
Graceful Restart : Disabled
Local Capability : RouteRefresh MP-BGP
Remote Capability:
Import Policy   : None Specified / Inherited
Export Policy   : None Specified / Inherited
...

```

```
-----
A:ALA-48#
```

```
A:ALA-48# show router 2 bgp neighbor 10.20.1.3
```

```
=====
BGP Neighbor
=====
```

```
Peer   : 10.20.1.3
```

```
Group  : 1
```

```
-----
Peer AS      : 100
Peer Address  : 10.20.1.3
Local AS     : 100
Local Address : 10.20.1.2
Peer Type    : Internal
State        : Established
Last Event   : recvKeepAlive
Last Error   : Cease
Local Family : IPv4
Remote Family : IPv4
Hold Time    : 3
Active Hold Time : 3
Cluster Id   : None
Preference   : 170
Recd. Paths  : 1
IPv4 Recd. Prefixes : 11
IPv4 Suppressed Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0
IPv6 Recd. Prefixes : 0
Input Queue  : 0
i/p Messages : 471
i/p Octets   : 3241
i/p Updates  : 4
TTL Security : Disabled
Advertise Inactive : Disabled
Advertise Label : None
Auth key chain : eta_keychain1
Local Capability : RouteRefresh MP-BGP
Remote Capability : RouteRefresh MP-BGP
Import Policy  : None Specified / Inherited

Peer Port      : 49725
Local Port     : 179
Last State     : Established
Keep Alive     : 1
Active Keep Alive : 1
Num of Flaps   : 0
IPv4 Active Prefixes : 10
VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Active Pfxs : 0
Mc IPv4 Active Pfxs. : 0
IPv6 Suppressed Pfxs : 0
IPv6 Active Prefixes : 0
Output Queue  : 0
o/p Messages  : 473
o/p Octets    : 3241
o/p Updates   : 4
Min TTL Value  : n/a
Peer Tracking  : Disabled

```

```

Export Policy      : static2bgp
-----
Neighbors : 1
=====
A:ALA-48#

A:ALA-12# show router bgp neighbor 10.0.0.11 orf
=====
BGP Neighbor 10.0.0.11 ORF
=====
Send List (Automatic)
-----
target:65535:10
target:65535:20
=====
A:ALA-12

A:ALA-22 show router bgp neighbor 10.0.0.1 orf
=====
BGP Neighbor 10.0.0.1 ORF
=====
Receive List
-----
target:65535:10
target:65535:20
=====
A:ALA-22

```

Sample Detailed Output

```

A:ALA-12# show router bgp neighbor detail
=====
BGP Neighbor (detail)
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205          Peer Port      : 0
Peer Address  : 10.0.0.15     Local Port     : 0
Local AS     : 65206          Local Address  : 10.0.0.16
Peer Type    : External
State        : Active         Last State     : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Connect Retry : 20            Local Pref.    : 100
Min Route Advt. : 30          Min AS Orig.   : 15

Damping      : Disabled       Loop Detect     : Ignore
MED Out      : No MED Out     Authentication : None
Next Hop Self : Disabled      AggregatorID Zero: Disabled
Remove Private : Disabled     Passive        : Disabled
Prefix Limit  : No Limit
Hold Time     : 90            Keep Alive     : 30
Active Hold Time : 0          Active Keep Alive: 0
Cluster Id    : None          Client Reflect  : Enabled
Preference    : 170           Num of Flaps   : 0

```

```

Recd. Prefixes : 0
Recd. Paths : 0
Input Queue : 0
i/p Messages : 0
i/p Octets : 0
i/p Updates : 0
Export Policy : direct2bgp
Active Prefixes : 0
Suppressed Paths : 0
Output Queue : 0
o/p Messages : 0
o/p Octets : 0
o/p Updates : 0

```

```

=====
A:ALA-12#

```

```

*A:SetupCLI>show>router>bgp# neighbor

```

```

=====
BGP Neighbor

```

```

=====
Peer : 3.3.3.3
Group : bgp_group_1 34567890123456789012
-----

```

```

Peer AS : 20 Peer Port : 0
Peer Address : 3.3.3.3 Local Port : 0
Local AS : 100
Local Address : 0.0.0.0
Peer Type : Internal
State : Active Last State : Idle
Last Event : stop
Last Error : Cease
Local Family : IPv4
Remote Family : Unused
Hold Time : 10 Keep Alive : 30
Active Hold Time : 0 Active Keep Alive : 0
Cluster Id : 2.2.3.4
Preference : 101 Num of Flaps : 0
Recd. Paths : 0
IPv4 Recd. Prefixes : 0 IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0 VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0 VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0 Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0 IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0 IPv6 Active Prefixes : 0
Input Queue : 0 Output Queue : 0
i/p Messages : 0 o/p Messages : 0
i/p Octets : 0 o/p Octets : 0
i/p Updates : 0 o/p Updates : 0
TTL Security : Disabled Min TTL Value : n/a
Graceful Restart : Enabled Stale Routes Time : 360
Advertise Inactive : Disabled Peer Tracking : Enabled
Advertise Label : None Bfd Enabled : Yes
Auth key chain : n/a
Local Capability : RouteRefresh MP-BGP
Remote Capability :
Import Policy : test i1
: test i2
: test i3
: test i4
: test i5 890123456789012345678901
Export Policy : test e1
: test e2
: test e3
: test e4

```

```

: test e5 890123456789012345678901
-----
Neighbors : 1
=====

*A:~RR>config>router>bgp>group# show router bgp neighbor 2.2.2.2 detail

=====
BGP Neighbor
=====
-----
Peer          : 2.2.2.2
Description   : (Not Specified)
Group         : cisco
-----
Peer AS       : 65002           Peer Port      : 53257
Peer Address  : 2.2.2.2
Local AS      : 65002           Local Port     : 179
Local Address : 1.1.1.1
Peer Type     : Internal
State         : Established     Last State      : Established
Last Event    : recvKeepAlive
Last Error    : Cease (Connection Collision Resolution)
Local Family  : IPv4
Remote Family : IPv4
Connect Retry : 120             Local Pref.     : 100
Min Route Advt. : 30
Multihop      : 0 (Default)    AS Override     : Disabled
Damping       : Disabled       Loop Detect      : Ignore
MED Out       : No MED Out     Authentication   : None
Next Hop Self : Disabled       AggregatorID Zero : Disabled
Remove Private : Disabled      Passive         : Disabled
Peer Identifier : 2.2.2.2       Fsm Est. Trans  : 1
Fsm Est. Time : 01h17m16s     InUpd Elap. Time : 01h19m03s
Hold Time     : 90             Keep Alive      : 30
Min Hold Time : 0
Active Hold Time : 90          Active Keep Alive : 30
Cluster Id    : None           Client Reflect   : Disabled
Preference    : 170           Num of Update Flaps : 0
Recd. Paths   : 1
IPv4 Recd. Prefixes : 1       IPv4 Active Prefixes : 1
IPv4 Suppressed Pfxs : 0       VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0       VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0       Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0       IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0       IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0       VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
Mc IPv6 Recd. Pfxs. : 0       Mc IPv6 Active Pfxs. : 0
Mc IPv6 Suppr. Pfxs : 0       L2-VPN Suppr. Pfxs : 0
L2-VPN Recd. Pfxs : 0       L2-VPN Active Pfxs : 0
MVPN-IPv4 Suppr. Pfxs : 0     MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0     MDT-SAFI Suppr. Pfxs : 0
MDT-SAFI Recd. Pfxs : 0       MDT-SAFI Active Pfxs : 0
Flow-IPv4 Suppr. Pfxs : 0     Flow-IPv4 Recd. Pfxs : 0
Flow-IPv4 Active Pfxs : 0     Rte-Tgt Suppr. Pfxs : 0
Rte-Tgt Recd. Pfxs : 0       Rte-Tgt Active Pfxs : 0
Backup IPv4 Pfxs : 0         Backup IPv6 Pfxs : 0
Mc Vpn Ipv4 Recd. Pf*: 0     Mc Vpn Ipv4 Active P*: 0

```

```

Mc Vpn Ipv4 Suppr. P*: 0
Backup Vpn IPv4 Pfxs : 0
Input Queue           : 0
i/p Messages          : 158
i/p Octets            : 3090
i/p Updates           : 1
MVPN-IPv6 Suppr. Pfxs: 0
MVPN-IPv6 Active Pfxs: 0
Flow-IPv6 Suppr. Pfxs: 0
Flow-IPv6 Active Pfxs: 0
Evpn Suppr. Pfxs      : 0
Evpn Active Pfxs      : 0
MS-PW Suppr. Pfxs     : 0
MS-PW Active Pfxs     : 0
TTL Security          : Disabled
Graceful Restart      : Disabled
Restart Time          : n/a
Advertise Inactive    : Disabled
Advertise Label        : None
Auth key chain         : n/a
Disable Cap Nego       : Disabled
Flowspec Validate     : Disabled
Aigp Metric           : Disabled
Damp Peer Oscillatio*: Disabled
GR Notification        : Disabled
Rem Idle Hold Time    : 00h00m00s
Next-Hop Unchanged    : None
L2 VPN Cisco Interop  : Disabled
Local Capability       : RtRefresh MPBGP 4byte ASN
Remote Capability      : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi*: Disabled
Remote AddPath Capab*: Send - None
                      : Receive - None
Import Policy          : link-bw
Export Policy          : None Specified / Inherited
Origin Validation      : N/A
EBGP Link Bandwidth   : n/a
IPv4 Rej. Pfxs        : 0
VPN-IPv4 Rej. Pfxs    : 0
Mc IPv4 Rej. Pfxs     : 0
MVPN-IPv4 Rej. Pfxs   : 0
Flow-IPv4 Rej. Pfxs   : 0
L2-VPN Rej. Pfxs      : 0
Rte-Tgt Rej. Pfxs     : 0
Mc Vpn Ipv4 Rej. Pfxs: 0
Backup Vpn IPv6 Pfxs : 0
Output Queue          : 0
o/p Messages          : 157
o/p Octets            : 3009
o/p Updates           : 0
MVPN-IPv6 Recd. Pfxs : 0
Flow-IPv6 Recd. Pfxs : 0
Evpn Recd. Pfxs       : 0
MS-PW Recd. Pfxs      : 0
Min TTL Value         : n/a
Stale Routes Time     : n/a
Peer Tracking         : Disabled
Bfd Enabled           : Disabled
Default Route Tgt     : Disabled
Split Horizon         : Disabled
Update Errors         : 0
Fault Tolerance       : Disabled

```

```

=====
Prefix Limits Per Address Family
=====

```

Family	Limit	Idle Timeout	Threshold	Log Only	Post Import
ipv4	1000	forever	90	Disabled	Enabled
vpnIpv4	1000	forever	90	Disabled	Enabled

```

-----
Neighbors : 1
=====

```

```

* indicates that the corresponding row element may have been truncated.

```

Advertised and Received Routes Output

The following table describes the command output for both the standard and detailed information for a neighbor.

Table 59 **Standard and Detailed BGP Neighbor Output Fields**

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then it is the same value as the AS.
Flag	u — used s — suppressed h — history d — decayed * — valid i — igp e — egp ? — incomplete > — best
Network	Route IP prefix and mask length for the route.
Next Hop	BGP nexthop for the route.
LocalPref	BGP local preference path attribute for the route.
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route.
AS Path	The BGP AS path for the route.

Sample Output

```
A:Dut-F# show router bgp neighbor 1.101.201.201 received-routes
=====
BGP Router ID:10.20.1.101      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p -
purge Origin codes : i - IGP, e - EGP, ? - incomplete
=====
BGP Routes
=====
Flag  Network                               LocalPref  MED
      Nexthop (Router)                     Path-Id    Label
      As-Path
```

```

-----
u*>i 10.20.1.201/32          n/a      None
      1.101.201.201        238      -
      200
*i   10.20.1.202/32          n/a      2000
      1.101.201.201        3752     -
      200
u*>i 10.20.1.231/32          n/a      1000
      1.101.201.201        3751     -
      200
-----
Routes : 3
=====
A:Dut-F#

A:Dut-F# show router bgp neighbor 1.101.201.201 advertised-routes
=====
BGP Router ID:10.20.1.101      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p -
purge Origin codes : i - IGP, e - EGP, ? - incomplete
=====
BGP Routes
=====
Flag  Network                LocalPref  MED
      Nexthop (Router)
      As-Path
-----
i     10.20.1.101/32          n/a      None
      1.101.201.101        88      -
      100
i     10.20.1.102/32          n/a      2000
      1.101.201.101        1118     -
      100
i     10.20.1.131/32          n/a      2000
      1.101.201.101        790      -
      100
-----
Routes : 3
=====
A:Dut-F#

A:core_east# show router bgp neighbor 10.193.0.10 graceful-restart
=====
BGP Neighbor 10.193.0.10 Graceful Restart
=====
Graceful Restart locally configured for peer: Enabled
Peer's Graceful Restart feature           : Enabled
NLRI(s) that peer supports restart for    : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that peer saved forwarding for    : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that restart is negotiated for    : None
NLRI(s) of received end-of-rib markers    : IPv4-Unicast
NLRI(s) of all end-of-rib markers sent    : IPv4-Unicast
Restart time locally configured for peer  : 120 seconds
Restart time requested by the peer        : 390 seconds
Time stale routes from peer are kept for  : 360 seconds
Graceful restart status on the peer       : Not currently being helped

```



```
Number of Restarts          : 328
Last Restart at            : 08/20/2006 12:22:06
=====
A:core_east#
```

next-hop

- Syntax** **next-hop** [*family* [**service-id** *service-id*]] [*ip-address* [**detail**]]
- Context** show>router>bgp
- Description** Displays BGP next-hop information.
- Parameters** *family* — Specify the type of routing information to be distributed by the BGP instance.
- Values** ipv4, mcast-ipv4, ipv6, mcast-ipv6, label-ipv4, label-ipv6, vpn-ipv4, vpn-ipv6, evpn, mcast-vpn-ipv4, mcast-vpn-ipv6
- ip-address* — Displays the next hop information for the specified IP address.
- Values** ipv4-address:
- a.b.c.d (host bits must be 0)
- ipv6-address:
- x:x:x:x:x:x:x (eight 16-bit pieces)
 - x:x:x:x:x:d.d.d.d
 - x: [0 to FFFF]H
 - d: [0 to 255]D
- detail** — Display the longer, more detailed version of the output.
- Output** Show Next-Hop Output

The following table describes the command output fields for a BGP next hop.

Table 60 BGP Next Hop Output Fields

Label	Description
BGP ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Next Hop	The next-hop address.
Resolving Prefix	Displays the prefix of the best next hop.
Owner	Displays the routing protocol used to derive the best next hop.
Preference	Displays the BGP preference attribute for the routes.

Table 60 BGP Next Hop Output Fields (Continued)

Label	Description
Reference Count	Displays the number of routes using the resolving prefix.
Resolved Next Hop	The IP address of the next hop.

Sample Output

```
*A:Dut-C# show router bgp next-hop
=====
BGP Router ID:10.20.1.3      AS:5000      Local AS:5000
=====

BGP Next Hop
=====
Next Hop                                Pref Owner
  Resolving Prefix                      Metric
  Resolved Next Hop                     Ref. Count
-----
10.20.1.1                                7    RSVP
   10.20.1.1/32                          1000
   10.10.2.1                              2
10.20.1.2                                7    RSVP
   10.20.1.2/32                          1000
   10.10.3.2                              2
10.20.1.4                                7    RSVP
   10.20.1.4/32                          1000
   10.10.11.4                             2
-----
Next Hops : 3

A:ALA-49>show>router>bgp# next-hop 192.168.2.194
=====
BGP Router ID : 10.10.10.104      AS : 200      Local AS : 200
=====
BGP Next Hop
=====
Next Hop      Resolving      Owner  Preference Reference  Resolved
               Prefix                               Count      Next Hop
-----
A:ALA-49>show>router>bgp# next-hop 10.10.10.104
```

optimal-route-reflection

Syntax	optimal-route-reflection [bgp-nh-info [location <i>location-id</i>] [bgp-nh <i>ip-address</i>]]
Context	show>router>bgp
Description	This command displays the calculated IGP costs that are used by BGP in order to perform optimal route reflection (ORR).

Parameters **bgp-nh-info** — Displays the calculated IGP cost from a specific location ID to all prefixes that contain the BGP next-hop addresses.

location-id — Specifies the ORR location identifier.

Values 1 to 16

ip-address — Specifies the IPv4 or IPv6 BGP next-hop address.

Values ipv4-address: a.b.c.d
 ipv6-address:
 • x:x:x:x:x:x:x (eight 16-bit pieces)
 • x:x:x:x:x:d.d.d.d
 • x: [0 to FFFF]H
 • d: [0 to 255]D

Output The following table describes the command output fields for an optimal route reflection.

Table 61 ORR Fields

Label	Description
Next Hop Loc	The address of the next-hop location.
Dest-Prefix	Displays the longest prefix match for the BGP next-hop address from the point of view of a specific location. Displays “unreachable” if no route exists to the BGP next-hop from the location.
DB-Source	Displays the source of the reachability information in the TE database.
Type	Displays the type of route or reachability, either intra-area or inter-area.
Proto	Displays the source protocol that is providing reachability to the BGP next-hop address.
Metric	Displays the IGP cost to reach the BGP next-hop from the location reference point.
Preference	Displays the routing preference of the route that is providing reachability to the BGP next-hop address. This is derived from standard protocol preference values.

Sample Output

```
*A:Dut-C# show router bgp optimal-route-reflection bgp-nh-info
=====
ORR BGP-NH Table (Router: Base)
=====
Location 1:
    Primary      : 1.1.1.1
```

```
Secondary      : -
Tertiary       : -
Primary-ipv6   : 2001::1
Secondary-ipv6 : -
Tertiary-ipv6  : -
Age            : 00h00m00s
Spf wait       : 60
Initial wait   : 5
Second wait    : 15
-----
Next Hop
  Loc      Dest-Prefix
-----
10.0.0.10
  1        Unreachable
192.0.2.1
  1        Unreachable
-----
No. of BGP-NHs: 2
=====
```

paths

- Syntax** paths
- Context** show>router>bgp
- Description** This command displays a summary of BGP path attributes.
- Output** The following table describes the command output fields for a BGP path.

Table 62 BGP Paths Output Fields

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Path	The AS path attribute.
Origin	EGP — The NLRI is learned by an EGP protocol. IGP — The NLRI is interior to the originating AS. INCOMPLETE — NLRI was learned another way.
Next Hop	The advertised BGP next-hop.
MED	The Multi-Exit Discriminator value.

Table 62 BGP Paths Output Fields (Continued)

Label	Description
Local Preference	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
Refs	The number of routes using a specified set of path attributes.
ASes	The number of autonomous system numbers in the AS path attribute.
Segments	The number of segments in the AS path attribute.
Flags	EBGP-learned — Path attributes learned by an EBGP peering. IBGP-Learned — Path attributes learned by an IBGP peering.
Aggregator	The route aggregator ID.
Community	The BGP community attribute list.
Originator ID	The originator ID path attribute value.
Cluster List	The route reflector cluster list.

Sample Output

```

=====
BGP Router ID : 10.0.0.14          AS : 65206   Local AS : 65206
=====
BGP Paths
=====
Path: 60203 65001 19855 3356 15412
-----
Origin          : IGP                Next Hop        : 10.0.28.1
MED             : 60203              Local Preference : none
Refs            : 4                  ASes           : 5
Segments        : 1
Flags           : EBGP-learned
Aggregator      : 15412 62.216.140.1
-----
Path: 60203 65001 19855 3356 1      1236 1236 1236 1236
-----
Origin          : IGP                Next Hop        : 10.0.28.1
MED             : 60203              Local Preference : none
Refs            : 2                  ASes           : 9
Segments        : 1
Flags           : EBGP-learned

```

policy-test

Syntax **policy-test** *plcy-or-long-expr* [*plcy-or-expr* [*plcy-or-expr*]] [**family** *family*] [**rd** *rd*] **prefix** *ip-*

*prefix**ip-prefix-length* [**longer**] [**neighbor** *ip-address*] [**display-rejects**] [**detail** | **brief**]

- Context** show>router>bgp>routes
- Description** This command allows an operator to evaluate an existing policy against the RIB to identify what prefixes are matched/not matched by the policy prior to attaching it to a routing neighbor or instance.
- Parameters** *plcy-or-long-expr* — Specifies up to 14 policy names or long expression.
- Values** *policy-name, long-expr*
- plcy-or-expr* — Specifies the policy name or expression.
- Values** *policy-name, ong-expr*
- policy-name* — Specifies the policy name. Must be the name of an existing configured and committed policy, up to 64 characters.
- long-expr* — Specifies a long expression up to 255 characters.
- expr* — Specifies an expression up to 64 characters.
- family* — ipv4, ipv6, vpn-ipv4, vpn-ipv6, label-ipv4, label-ipv6
- Default** ipv4
- rd* — {*ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}
- ip-prefix/ip-prefix-length* — Specifies an IPv4 or IPv6 prefix/mask to be evaluated. The keyword **longer** may be specified to evaluate longer prefix matches. (optional)
- neighbor* — Specifies a BGP neighbor.
- display-rejects** — Displays routes that were rejected by the policy. If not specified, only a count of rejected routes are shown.
- detail** — When the policy modifies route attributes, it displays the modifications made to the routes. This command requires an exact prefix to be specified.
- Output** The following output is an example of policy test information.

Sample Output

```
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 0.0.0.0/0
longer neighbor 220.0.0.2
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network                               LocalPref  MED
      Nexthop                             Path-Id    VPNLabel
      As-Path
-----
```

```

Accepted by Policy
u*>? 4.0.0.6/32
220.0.0.2
14
None
None
None
-

-----
Total Routes : 17 Routes rejected : 16
=====
A:sim-1# show router bgp policy-test bgpprefix6 prefix 0.0.0.0/0 longer neighbor
220.0.0.2
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag Network
-----
Accepted by Policy
u*>? 4.0.0.6/32
-----
Total Routes : 17 Routes rejected : 16
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 0.0.0.0/0
longer neighbor 220.0.0.2 display-rejects brief
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag Network
-----
Rejected by Default action
u*>? 2.2.2.2/32
Rejected by Default action
u*>? 4.0.0.1/32
Rejected by Default action
u*>? 4.0.0.2/32
Rejected by Default action
u*>? 4.0.0.3/32
Rejected by Default action
u*>? 4.0.0.4/32
Rejected by Default action
u*>? 4.0.0.5/32
Accepted by Policy
u*>? 4.0.0.6/32
Rejected by Default action
u*>? 6.0.0.1/32
Rejected by Default action
u*>? 7.0.0.1/32
Rejected by Default action
u*>i 10.0.4.0/24
Rejected by Default action

```

```

*i    10.12.0.0/24
Rejected by Default action
*i    10.14.0.0/24
Rejected by Default action
u*>i  10.24.0.0/24
Rejected by Default action
*i    12.12.12.12/32
Rejected by Default action
*i    220.0.0.2/32
Rejected by Default action
*i    220.0.0.3/32
Rejected by Default action
u*>i  221.0.0.2/32
-----
Total Routes : 17 Routes rejected : 16
=====
A:sim-1# show router bgp policy-test bgpprefix6 prefix 0.0.0.0/0 longer neighbor
220.0.0.2 display-rejects
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network                      LocalPref  MED
      Nexthop                      Path-Id    VPNLabel
      As-Path
-----
Rejected by Default action
u*>?  2.2.2.2/32                      None       None
      220.0.0.2                      None       -
      14
Rejected by Default action
u*>?  4.0.0.1/32                      None       None
      220.0.0.2                      None       -
      14
Rejected by Default action
u*>?  4.0.0.2/32                      None       None
      220.0.0.2                      None       -
      14
Rejected by Default action
u*>?  4.0.0.3/32                      None       None
      220.0.0.2                      None       -
      14
Rejected by Default action
u*>?  4.0.0.4/32                      None       None
      220.0.0.2                      None       -
      14
Rejected by Default action
u*>?  4.0.0.5/32                      None       None
      220.0.0.2                      None       -
      14
Accepted by Policy
u*>?  4.0.0.6/32                      None       None
      220.0.0.2                      None       -
      14

```



```

Rejected by Default action
u*>? 6.0.0.1/32          None      None
      220.0.0.2          None      -
      14
Rejected by Default action
u*>? 7.0.0.1/32          None      None
      220.0.0.2          None      -
      14
Rejected by Default action
u*>i 10.0.4.0/24          None      None
      220.0.0.2          None      -
      14
Rejected by Default action
*i 10.12.0.0/24          None      20
      220.0.0.2          None      -
      14
Rejected by Default action
*i 10.14.0.0/24          None      None
      220.0.0.2          None      -
      14
Rejected by Default action
u*>i 10.24.0.0/24         None      None
      220.0.0.2          None      -
      14
Rejected by Default action
*i 12.12.12.12/32        None      20
      220.0.0.2          None      -
      14
Rejected by Default action
*i 220.0.0.2/32          None      None
      220.0.0.2          None      -
      14
Rejected by Default action
*i 220.0.0.3/32          None      10
      220.0.0.2          None      -
      14
Rejected by Default action
u*>i 221.0.0.2/32         None      None
      220.0.0.2          None      -
      14

```

Total Routes : 17 Routes rejected : 16

=====

```

A:sim-1# show router bgp policy-test bgpprefix6 prefix 4.0.0.1/32 detail neighbor
220.0.0.2 display-rejects

```

=====

```

BGP Router ID:11.11.11.10      AS:11      Local AS:11

```

=====

```

Legend -

```

```

Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

```

=====

```

BGP IPv4 Routes

```

=====

```

Rejected by Default action
Network      : 4.0.0.1/32
Nextthop     : 220.0.0.2
Path Id      : None

```

```

From          : 220.0.0.2
Res. Nexthop  : 10.14.0.4
Local Pref.   : None
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : None
Connector     : None
Community     : target:65530:20
Cluster       : No Cluster Members
Originator Id : None
Fwd Class     : None
Flags         : Used Valid Best Incomplete
Route Source  : External
AS-Path       : 14
Interface Name : to-sim-6
Aggregator     : None
MED            : None
Peer Router Id : 14.14.14.10
Priority       : None

```

```

-----
Total Routes : 1 Routes rejected : 1
=====

```

```

A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.6/32
neighbor 220.0.0.2

```

```

=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====

```

```

Legend -

```

```

Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

```

```

=====
BGP IPv4 Routes
=====

```

```

Accepted by Policy
-----

```

```

Original Attributes

```

```

Network       : 4.0.0.6/32
Nexthop       : 220.0.0.2
Path Id       : None
From          : 220.0.0.2
Res. Nexthop  : 10.14.0.4
Local Pref.   : n/a
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : None
Connector     : None
Community     : target:65530:20
Cluster       : No Cluster Members
Originator Id : None
Fwd Class     : None
Flags         : Used Valid Best Incomplete
Route Source  : External
AS-Path       : 14
Interface Name : to-sim-6
Aggregator     : None
MED            : None
Peer Router Id : 14.14.14.10
Priority       : None

```

```

Modified Attributes

```

```

Network       : 4.0.0.6/32
Nexthop       : 220.0.0.2
Path Id       : None
From          : 220.0.0.2
Res. Nexthop  : 10.14.0.4

```

```

Local Pref.      : None
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
AIGP Metric     : None
Connector       : None
Community       : 2:11 2:12 2:13 target:65530:20
Cluster         : No Cluster Members
Originator Id   : None
Peer Router Id  : 14.14.14.10
Fwd Class       : None
Priority        : None
Flags           : Used Valid Best Incomplete
Route Source    : External
AS-Path         : 14

```

```

-----
Routes : 1
=====

```

```

A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.6/32
longer neighbor 220.0.0.2

```

```

=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network                LocalPref  MED
      Nexthop                Path-Id    VPNLabel
      As-Path
-----
Accepted by Policy
u*>?  4.0.0.6/32              None       None
      220.0.0.2              None       -
      14
-----

```

```

Routes : 1
=====

```

```

A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.6/32
longer neighbor 220.0.0.2 detail

```

```

=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network
-----
Accepted by Policy
u*>?  4.0.0.6/32
-----
Routes : 1
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.0/24
longer neighbor 220.0.0.2 brief

```

```
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag Network
-----
Accepted by Policy
u*>? 4.0.0.6/32
-----
Total Routes : 6 Routes rejected : 5
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.0/24
longer neighbor 220.0.0.2 display-rejects detail
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag Network
-----
Rejected by Default action
u*>? 4.0.0.1/32
Rejected by Default action
u*>? 4.0.0.2/32
Rejected by Default action
u*>? 4.0.0.3/32
Rejected by Default action
u*>? 4.0.0.4/32
Rejected by Default action
u*>? 4.0.0.5/32
Accepted by Policy
u*>? 4.0.0.6/32
-----
Total Routes : 6 Routes rejected : 5
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.0/24
longer neighbor 220.0.0.2 display-rejects
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag Network                                LocalPref MED
Nexthop                                     Path-Id    VPNTLabel
As-Path
```

```

-----
Rejected by Default action
u*>?  4.0.0.1/32                None      None
      220.0.0.2                None      -
      14
Rejected by Default action
u*>?  4.0.0.2/32                None      None
      220.0.0.2                None      -
      14
Rejected by Default action
u*>?  4.0.0.3/32                None      None
      220.0.0.2                None      -
      14
Rejected by Default action
u*>?  4.0.0.4/32                None      None
      220.0.0.2                None      -
      14
Rejected by Default action
u*>?  4.0.0.5/32                None      None
      220.0.0.2                None      -
      14
Accepted by Policy
u*>?  4.0.0.6/32                None      None
      220.0.0.2                None      -
      14
-----
Total Routes : 6 Routes rejected : 5
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.0/24
longer neighbor 220.0.0.2 display-rejects brief
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP IPv4 Routes
=====
Flag  Network
-----
Rejected by Default action
u*>?  4.0.0.1/32
Rejected by Default action
u*>?  4.0.0.2/32
Rejected by Default action
u*>?  4.0.0.3/32
Rejected by Default action
u*>?  4.0.0.4/32
Rejected by Default action
u*>?  4.0.0.5/32
Accepted by Policy
u*>?  4.0.0.6/32
-----
Total Routes : 6 Routes rejected : 5
=====
A:sim-1# show router bgp policy-test bgpprefix6 family ipv4 prefix 4.0.0.0/24

```

```

longer neighbor 220.0.0.2
=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network                      LocalPref  MED
      Nexthop                      Path-Id    VPNLabel
      As-Path
-----
Accepted by Policy
u*>?  4.0.0.6/32                      None       None
      220.0.0.2                      None       -
      14
-----
Total Routes : 6 Routes rejected : 5

```

```

A:sim-1# show router bgp policy-test bgpprefix44rj family vpn-ipv4 prefix 0.0.0.0/0
longer neighbor display-rejects

```

```

=====
BGP Router ID:11.11.11.10      AS:11      Local AS:11
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv4 Routes
=====
Flag  Network                      LocalPref  MED
      Nexthop                      Path-Id    VPNLabel
      As-Path
-----
Accepted by Policy
u*>i  1:30:192.14.15.0/24              None       None
      220.0.0.2                      None       131069
      14
Accepted by Policy
u*>i  65530:20:8.0.0.1/32              None       None
      220.0.0.2                      None       131070
      14
Accepted by Policy
u*>i  65530:20:10.0.3.0/24             None       None
      220.0.0.2                      None       131070
      14 101
Accepted by Policy
u*>i  65530:20:10.13.0.0/24            None       None
      220.0.0.2                      None       131070
      14 101
Accepted by Policy
u*>i  65530:20:10.23.0.0/24            None       None
      220.0.0.2                      None       131070
      14 101

```

Accepted by Policy		
u*>i 65530:20:13.13.13.13/32	None	None
220.0.0.2	None	131070
14 101		
Accepted by Policy		
u*>i 65530:20:20.20.20.5/32	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:20.20.20.6/32	None	None
220.0.0.2	None	131070
14		
Rejected by Policy Entry = 10		
u*>i 65530:20:44.44.44.0/24	None	None
220.0.0.2	None	131070
14 101		
Accepted by Policy		
u*>i 65530:20:192.14.15.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.16.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.17.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.18.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.19.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.20.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.21.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.22.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.23.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:192.14.25.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
u*>i 65530:20:196.34.0.0/24	None	None
220.0.0.2	None	131070

```

      14
Accepted by Policy
u*>i  220.0.0.2:50:192.50.50.0/24          None      None
      220.0.0.2                          None      131067
      14
Accepted by Policy
u*>i  220.0.0.2:50:220.0.0.2/32           None      None
      220.0.0.2                          None      131067
      14
-----
Total Routes : 22 Routes rejected : 1
=====

```

prefix-label

Syntax	prefix-label
Context	show>router>bgp
Description	This command displays BGP Prefix label information.

routes

Syntax

```

routes [family] [brief]
routes [family] prefix [detail | longer | hunt [brief]]
routes family prefix [detail | longer | hunt [brief]]
routes [family [type mvpn-type]] community comm-id
routes [family [type mvpn-type]] aspath-regex reg-ex
routes vpn-ipv4 prefix [rd rd] [detail | longer | hunt [brief]]
routes vpn-ipv6 prefix [rd rd] [detail | longer | hunt [brief]]
routes mvpn-ipv4 type mvpn-type {rd rd | originator-ip ip-address | source-ip ip-address
  | group-ip ip-address | source-as as-number} [hunt | detail]
routes [family [l2vpn-type]] [brief]
routes [family [l2vpn-type]] community comm-id
routes [family [l2vpn-type]] aspath-regex reg-ex
routes l2-vpn l2vpn-type {[rd rd] | [site-id site-id] | [veid veid] [offset vpls-base-offset ]}
routes mdt-safi [rd rd] [grp-address mcast-grp-address] [brief]
routes ms-pw [rd rd] [aii-type2 aii-type2] [brief]
routes flow-ipv4
routes evpn inclusive-mcast [hunt | detail] [rd rd] [originator-ip ip-address] [next-hop
  ip-address] [community comm-id] [tag vni-id]
routes evpn ip-prefix [hunt | detail] [rd rd] [prefix ip-prefix/mask] [community comm-id]
  [tag vni-id] [next-hop ip-address]
routes evpn mac [hunt | detail] [rd rd] [next-hop ip-address] [mac-address mac-address]
  [community comm-id] [tag vni-id]

```



```

routes bgp-ls bgp-ls [hunt] [node | link | ipv4-prefix [ipv4-prefix/mask-len]]
route-target [source-as as-number] [brief] [aspath-regex reg-exp] [community comm-id]
route-target [rtc-prefix rtc-prefix] [hunt] [brief] [aspath-regex reg-exp] [community comm-id]
route-target rtc-prefix rtc-prefix [aspath-regex reg-exp] [community comm-id]
route-target [rtc-prefix rtc-prefix] [detail | longer] [aspath-regex reg-exp] [community comm-id]

```

Context	show>router>bgp
Description	This command displays BGP route information. When this command is issued without any parameters, then the entire BGP routing table displays. When this command is issued with an IP prefix/mask or IP address, then the best match for the parameter displays.
Parameters	<i>family</i> — Specify the type of routing information to be distributed by the BGP instance. Values evpn — Displays the BGP information related to Ethernet VPN. ipv4 — Displays only those BGP peers that have the IPv4 family enable and not those capable of exchanging IP-VPN routes. vpn-ipv4 — Displays the BGP peers that are IP-VPN capable. ipv6 — Displays the BGP peers that are IPv6 capable. mcast-ipv4 — Displays the BGP peers that are mcast-ipv4 capable. <i>mvpn-type</i> — Specifies the mvpn-type Values intra-ad, inter-ad, spmsi-ad, leaf-ad, source-ad, shared-join, and source-join. received — Specifies to show the BGP routes received from the neighbor. <i>prefix</i> — Specifies the prefix to display. Values - ip-address: a.b.c.d - ipv4-prefix: a.b.c.d - ipv4-prefix-length: [0 to 32] - ipv6-prefix [/pref* ipv6-prefix]: - x:x:x:x:x:x:x (eight 16-bit pieces) - x:x:x:x:x:d.d.d.d - x: [0 to FFFF]H - d: [0 to 255]D - prefix-length: [0 to 128] <i>filter</i> — Specifies route criteria. Values hunt—Displays entries for the specified route in the RIB-In, RIB-Out, and RTM. longer—Displays the specified route and subsets of the route. detail—Display the longer, more detailed version of the output.

aspath-regex “*reg-exp*” — Displays all routes with an AS path matching the specified regular expression *reg-exp*.

community — Displays all routes with the specified BGP community.

comm-id — Specifies community IDs, up to 72 characters.

Values [as-num:comm-val | ext-comm | well-known-comm]

where:

- *as-num* — 0 to 65535
- *comm-val* — 0 to 65535
- *ext-comm* — the extended community, defined as one of the following:

- {*target* | *origin*}:*ip-address*:*comm-val*
- {*target* | *origin*}:*asnum*:*ext-comm-val*
- {*target* | *origin*}:*ext-asnum*:*comm-val*
- **bandwidth**:*asnum*:*val-in-mbps*
- **ext:4300**:*ovstate*
- **ext** *lt:value1:value2*
- **flowspec-set**:*ext-asnum*:*group-id*

where:

- *target* — route target
- *origin* — route origin
- *ip-address* — a.b.c.d
- *ext-comm-val* — 0 to 4294967295
- *ext-asnum* — 0 to 4294967295
- **bandwidth** — bandwidth
- *val-in-mbps* — 0 to 16777215
- **ext** — extended
- **ext:4300** — origin verification
- *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
- *value1* — 0000 to FFFF
- *value2* — 0 to FFFFFFFFFFFFFF
- **flowspec-set** — flow-spec set
- *group-id* — 0 to 16383
- *well-known-comm* — **null** | **no-export** | **no-export-subconfed** | **no-advertise** | **blackhole**

brief — Provides a summarized display of the set of peers to which a BGP route is advertised.

rd — Allows more precise definition of the RD vs. prefix for VPN-IPv6 routes.

Values ip-addr:comm-val

2byte-asnumber:ext-comm-val

4byte-asnumber:comm-val}

veid — Specifies a two byte identifier that represents the local bridging instance in a VPLS and is advertised through the BGP NLRI. This value must be lower than or equal to the max-ve-id.

Values 0 to 4294967295

vpls-base-offset — Specifies a two byte identifier advertised through the NLRI that is used to indicate which VE-ID should use the advertised NLRI at the receiving PE according to the following rule: if the offset ≤ local VE-ID ≤ offset+VBS-1 (VBS = virtual block size = 8 in our implementation) then the NLRI is processed. Otherwise it is ignored. The NLRI with this offset is generated as soon as the first VE-ID value between (offset, offset + VBS-1) is advertised in the network.

Values 0 to 4294967295

site-id — Specifies a two byte identifier usually employed for the BGP multi-homing solution. It identifies the BGP multi-homing site associated with one or a set of objects (SAPs, pseudowires or combination). The site-id must be identical between the two PEs carrying the connection to the access device multi-homed to the PEs.

Values 0 to 4294967295

l2vpn-type — Specifies a 12-byte Virtual Switch Instance identifier (VSI-ID) type.

Values bgp-ad, bgp-vpls, multi-homing

ms-pw [rd rd] [aii-type2 aii-type2] [brief] — Displays routes for ms-pw family.

route-target — Displays a summary of route-target constrain routes for this peer.

rtc-prefix — *source-as:{ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val}/prefix-length*

Values *ip-addr:* a.b.c.d
comm-val: 0 to 65535
2byte-asnumber: 0 to 65535
ext-comm-val: 0 to 4294967295
4byte-asnumber: 0 to 4294967295
prefix-length: 0 to 96

routes evpn inclusive-mcast [hunt | detail] [rd rd] [originator-ip ip-address] [next-hop ip-address] [community comm-id] [tag vni-id] — Displays inclusive multicast routes for evpn family.

routes evpn ip-prefix [hunt | detail] [rd rd] [prefix ip-prefix/mask] [community comm-id] [tag vni-id] [next-hop ip-address] — Displays inclusive ip-prefix routes for evpn family.

routes evpn mac [hunt | detail] [rd rd] [next-hop ip-address] [mac-address mac-address] [community comm-id] [tag vni-id] — Displays MAC routes for the EVPN family.

tag *vni-id* — Displays all routes with the specified ethernet-tag. For VXLAN tunnels, the ethernet-tag encodes the VNI (VXLAN Network Identifier).

Values 1 to 16777215

Output BGP Route

The following table describes the command output fields for BGP routes.

Table 63 BGP Routes Output Fields

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Route Dist.	Displays the route distinguisher identifier attached to routes that distinguishes the VPN it belongs.
VPN Label	Displays the label generated by the PEs label manager.
Network	The IP prefix and mask length.
Nexthop	The BGP nexthop.
From	The advertising BGP neighbor's IP address.
Res. Nexthop	The resolved nexthop.
Local Pref.	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
Flag	u — used s — suppressed h — history d — decayed * — valid i — igp e — egp ? — incomplete > — best S — sticky
Aggregator AS	The aggregator AS value. none — Aggregator AS attributes are not present.

Table 63 BGP Routes Output Fields (Continued)

Label	Description
Aggregator	The aggregator attribute value. none — Aggregator attributes are not present.
Atomic Aggr.	Atomic —Not Atomic — The atomic aggregator flag is not set. The atomic aggregator flag is set.
MED	The MED metric value. none — MED metrics are present.
Community	The BGP community attribute list.
Cluster	The route reflector cluster list.
Originator Id	The originator ID path attribute value. none — The originator ID attribute is not present.
Peer Router Id	The router ID of the advertising router.
AS-Path	The BGP AS path attribute.
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed.

Table 63 BGP Routes Output Fields (Continued)

Label	Description
TieBreakReason	Displays the step in the BGP decision process where a BGP route lost the tie-break with the next better BGP route for the same prefix. LocalPref — This route is not the best because the next better route has a higher LOCAL_PREF. AIGP — This route is not the best because the next better route has a lower derived AIGP metric value. ASPathLen — This route is not the best because the next better route has a shorter AS PATH length. Origin — This route is not the best because the next better route has a lower Origin value. MED — This route is not the best because the next better route has a lower MED, and MED comparison of the routes was allowed. IBGP — This IBGP route is not the best because the next better route is an EBGP route. NHCost — This route is not the best because the next better route has a lower metric value to reach the BGP NEXT HOP. BGPID — This route is not the best because the next better route has a lower Originator ID or BGP Identifier. ClusterLen — This route is not the best because the next better route has a shorter Cluster list length. PeerIP — This route is not the best because the next better route has a lower neighbor IP address.

Sample Output

```

*A:7750SR7-PE# show router bgp routes 10.10.10.5
=====
BGP Router ID:10.20.1.3      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p -
purge Origin codes : i - IGP, e - EGP, ? - incomplete
=====
BGP IPv4 Routes
=====
Flag  Network                               LocalPref  MED
      Nexthop (Router)                     Path-Id    Label
      As-Path
-----
u*>?  10.10.10.0/24                          None       None
      10.20.1.4                             None       -
      200 300
-----

```

Routes : 1

```
*A:Dut-C# show router bgp routes l2-vpn detail
=====
BGP Router ID:10.20.1.3      AS:1000      Local AS:1000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP L2VPN Routes
=====
Route Type      : AutoDiscovery
Route Dist.     : 10.20.1.1:1
Prefix          : 10.20.1.1
Nextthop        : 10.20.1.1
From            : 10.20.1.1
Res. Nextthop   : n/a
Local Pref.     : 100
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
AIGP Metric     : Not Atomic
Community       : target:4455:4455 target:1.20.30.40:6543
                  l2-vpn/vrf-imp:100.1.200.1:65535
Cluster         : No Cluster Members
Originator Id   : None
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : No As-Path
-----
PMSI Tunnel Attribute :
Tunnel-type      : RSVP-TE P2MP LSP      Flags           : Leaf not required
MPLS Label       : 0
P2MP-ID          : 1001                  Tunnel-ID        : 61440
Extended-Tunne*  : 10.20.1.1
```

```
*A:Dut-C# show router bgp routes l2-vpn detail
=====
BGP Router ID:10.20.1.3      AS:1000      Local AS:1000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP L2VPN Routes
=====
Route Type      : AutoDiscovery
Route Dist.     : 10.20.1.1:1
Prefix          : 10.20.1.1
Nextthop        : 10.20.1.1
From            : 10.20.1.1
Res. Nextthop   : n/a
Local Pref.     : 100
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
AIGP Metric     : Not Atomic
Community       : target:4455:4455 target:1.20.30.40:6543
                  l2-vpn/vrf-imp:100.1.200.1:65535
Cluster         : No Cluster Members
Originator Id   : None
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : No As-Path
-----
PMSI Tunnel Attribute :
Tunnel-type      : RSVP-TE P2MP LSP      Flags           : Leaf not required
MPLS Label       : 0
P2MP-ID          : 1001                  Tunnel-ID        : 61440
Extended-Tunne*  : 10.20.1.1
```

```

12-vpn/vrf-imp:100.1.200.1:65535

Cluster      : No Cluster Members
Originator Id : None                      Peer Router Id : 10.20.1.1
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : No As-Path
-----
PMSI Tunnel Attribute :
Tunnel-type   : RSVP-TE P2MP LSP          Flags           : Leaf not required
MPLS Label    : 0
P2MP-ID       : 1001                      Tunnel-ID        : 61440
Extended-Tunne*: 10.20.1.1

*A:Dut-C>show>router>bgp# routes 12-vpn 10.20.1.1 rd 10.20.1.1:1 hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP L2VPN-AD Routes
=====
Route Type      : AutoDiscovery
Route Dist.     : 10.20.1.1:1
Prefix          : 10.20.1.1
Nextthop        : 10.20.1.1
From            : 10.20.1.2
Res. Nextthop   : n/a
Local Pref.     : 100
Aggregator AS   : None                      Interface Name : NotAvailable
Atomic Aggr.    : Not Atomic                Aggregator     : None
MED             : 0
AIGP Metric     : None
Connector       : None
Community       : target:1.20.30.40:6543
                  12-vpn/vrf-imp:100.1.200.1:65535
Cluster         : 1.1.1.1
Originator Id   : 10.20.1.1                  Peer Router Id : 10.20.1.2
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : No As-Path
-----
RIB Out Entries
-----
Routes : 1
=====
*A:Dut-C>show>router>bgp#

*A:Dut-C>show>router>bgp# routes vpn-ipv4 6.6.6.6/32 rd 10.20.1.4:1 hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid

```



```

Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv4 Routes
=====
-----
RIB In Entries
-----
Network      : 6.6.6.6/32
Nexthop      : 10.20.1.4
Route Dist.  : 10.20.1.4:1          VPN Label      : 131070
Path Id      : None
From         : 10.20.1.4
Res. Nexthop : n/a
Local Pref.  : 100                  Interface Name : int_to_D
Aggregator AS : None                Aggregator    : None
Atomic Aggr. : Not Atomic           MED           : None
AIGP Metric  : None
Connector    : None
Community    : target:100:100
Cluster      : No Cluster Members
Originator Id : None                Peer Router Id : 10.20.1.4
Fwd Class    : None                Priority       : None
Flags        : Used Valid Best Incomplete
Route Source : Internal
AS-Path      : 106
VPRN Imported : 1
-----
RIB Out Entries
-----
-----
Routes : 1
=====
*A:Dut-C>show>router>bgp#

*A:Dut-C>show>router>bgp# routes vpn-
ipv4 6.6.6.6 32 hunt << SAME AS ABOVE BUT RD NOT SPECIFIED.I.E. ANY RD (RD is option
al).
=====
BGP Router ID:10.20.1.3          AS:None          Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv4 Routes
=====
-----
RIB In Entries
-----
Network      : 6.6.6.6/32
Nexthop      : 10.20.1.4
Route Dist.  : 10.20.1.4:1          VPN Label      : 131070
Path Id      : None
From         : 10.20.1.4
Res. Nexthop : n/a
Local Pref.  : 100                  Interface Name : int_to_D
Aggregator AS : None                Aggregator    : None
Atomic Aggr. : Not Atomic           MED           : None
AIGP Metric  : None

```

```

Connector      : None
Community      : target:100:100
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id : 10.20.1.4
Fwd Class      : None                Priority       : None
Flags          : Used Valid Best Incomplete
Route Source   : Internal
AS-Path        : 106
VPRN Imported  : 1
-----
RIB Out Entries
-----
Routes : 1
=====
*A:Dut-C>show>router>bgp#

*A:Dut-C>show>router>bgp# routes 3FFE::606:609/128 vpn-ipv6 hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
RIB In Entries
-----
Network        : 3FFE::606:609/128
Nexthop        : ::FFFF:A14:104
Route Dist.    : 10.20.1.4:1          VPN Label      : 131070
Path Id        : None
From           : 10.20.1.4
Res. Nexthop   : n/a
Local Pref.    : 100
Aggregator AS  : None                 Interface Name : int_to_D
Atomic Aggr.   : Not Atomic           Aggregator     : None
AIGP Metric    : None                 MED            : None
Connector      : None
Community      : target:100:100
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id : 10.20.1.4
Fwd Class      : None                Priority       : None
Flags          : Used Valid Best Incomplete
Route Source   : Internal
AS-Path        : 106
VPRN Imported  : 1
-----
RIB Out Entries
-----
Routes : 1
=====
*A:Dut-C>show>router>bgp#

```

```
*A:Dut-C>show>router>bgp# routes vpn-ipv6 3FFE::606:607/128 rd 10.20.1.4:1 hunt
=====
BGP Router ID:10.20.1.3          AS:None          Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
-----
RIB In Entries
-----
Network      : 3FFE::606:607/128
Nexthop      : ::FFFF:A14:104
Route Dist.  : 10.20.1.4:1          VPN Label      : 131070
Path Id      : None
From         : 10.20.1.4
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                Interface Name : int_to_D
Atomic Aggr. : Not Atomic           Aggregator     : None
AIGP Metric  : None                MED            : None
Connector    : None
Community    : target:100:100
Cluster      : No Cluster Members
Originator Id : None                Peer Router Id  : 10.20.1.4
Fwd Class    : None                Priority        : None
Flags        : Used Valid Best Incomplete
Route Source  : Internal
AS-Path       : 106
VPRN Imported : 1
-----
RIB Out Entries
-----
Routes : 1
=====
*A:Dut-C>show>router>bgp# routes vpn-ipv6 3FFE::606:607/128 rd 10.20.1.4:2 hunt
=====
BGP Router ID:10.20.1.3          AS:None          Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
No Matching Entries Found
=====
*A:Dut-C>show>router>bgp#

*A:Dut-C# show router bgp routes hunt 1.1.1.1/32
=====
BGP Router ID:10.20.1.3          AS:5000          Local AS:5000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
```

```

Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP IPv4 Routes
=====
-----
RIB In Entries
-----
Network      : 1.1.1.1/32
Nexthop      : 10.20.1.1
From         : 10.20.1.1
Res. Nexthop : 10.20.1.1 (RSVP LSP: 1)
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
Community    : No Community Members
Cluster      : No Cluster Members
Originator Id : None
Flags        : Used Valid Best Incomplete
AS-Path      : No As-Path
Interface Name : ip-10.10.2.3
Aggregator    : None
MED           : None
Peer Router Id : 10.20.1.1
-----
RIB Out Entries
-----
-----
Routes : 1
=====

```

```

A:ALA-12>config>router>bgp# show router bgp routes family ipv4
=====
BGP Router ID : 10.10.10.103      AS : 200      Local AS : 200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Flag  Network      Nexthop      LocalPref  MED
     VPN Label      As-Path
-----
No Matching Entries Found
=====
A:ALA-12>config>router>bgp#

```

```

A:ALA-12>config>router>bgp# show router bgp routes 13.1.0.0/24 de
=====
BGP Router ID : 10.128.0.161 AS : 65535 Local AS : 65535
=====
Legend - Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Original Attributes
Network      : 13.1.0.0/24      Nexthop      : 10.20.1.20
Route Dist.  : 10070:100        VPN Label    : 152784
From         : 10.20.1.20      Res. Nexthop : 10.130.0.2
Local Pref.  : 100

```

```

Aggregator AS      :      none      Aggregator      :      none
Atomic Aggr.      :      Not Atomic MED      :      none
Community          :      target:10070:1
Cluster            :      No Cluster Members
Originator Id      :      None      Peer Router Id   : 10.20.1.20
Flags              :      Used Valid Best IGP
AS-Path            :      10070 {14730}
Modified Attributes

```

```

Network :13.1.0.0/24 Nexthop :10.20.1.20
Route Dist.: 10001:100 VPN Label :152560
From :10.20.1.20 Res. Nexthop :10.130.0.2
Local Pref.:100
Aggregator AS: none Aggregator:none
Atomic Aggr.:Not Atomic MED :none
Community :target:10001:1
Cluster :No Cluster Members
Originator Id:None Peer Router Id:10.20.1.20
Flags :Used Valid Best IGP
AS-Path :No As-Path

```

...

A:ALA-12>config>router>bgp#

A:SR-12# show router bgp routes 100.0.0.0/30 hunt

```

=====
BGP Router ID : 10.20.1.1   AS : 100Local AS : 100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
RIB In Entries
-----
Network      : 100.0.0.0/30
Nexthop      : 10.20.1.2
Route Dist.  : 10.20.1.2:1VPN Label: 131070
From         : 10.20.1.2
Res. Nexthop : 10.10.1.2
Local Pref.  : 100Interface Name: to-sr7
Aggregator AS : noneAggregator: none
Atomic Aggr. : Not AtomicMED: none
Community    : target:10.20.1.2:1
Cluster      : No Cluster Members
Originator Id : NonePeer Router Id: 10.20.1.2
Flags        : Used Valid Best IGP
AS-Path      : No As-Path
VPRN Imported : 1 2 10 12
-----
RIB Out Entries
-----
Routes : 1
=====
A:SR-12#

```

```

*A:praragon-sim1# /show router bgp routes mvpn-ipv4
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      LocalPref  MED  VPNLabel
      RD
      Nextthop
      As-Path
      SourceAS
      SourceIP
      GroupIP
-----
u*>i  Intra-Ad        10.20.1.4         100        0
      1:1             -                 -
      10.20.1.4       -                 -
      No As-Path      -
u*>i  Source-Ad        -                 100        0
      1:1             -                 -
      10.20.1.4       130.100.1.2
      No As-Path      227.0.0.0
u*>i  Source-Join      -                 100        0
      1:1             200               -
      10.20.1.4       150.100.1.2
      No As-Path      226.0.0.0
-----
Routes : 3
=====
*A:praragon-sim1#

*A:praragon-sim1# show router bgp routes mvpn-ipv4 brief
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      SourceIP
      RD
      SourceAS
      GroupIP
-----
u*>i  Intra-Ad        10.20.1.4         -
      1:1             -                 -
u*>i  Source-Ad        -                 130.100.1.2
      1:1             -                 227.0.0.0
u* >i Source-Join      -                 150.100.1.2
      1:1             200               226.0.0.0
-----
Routes : 3
=====
*A:praragon-sim1#

*A:praragon-sim1# show router bgp routes mvpn-ipv4 type source-join source-as

```

```

200 source-ip 150.100.1.2 group-ip 226.0.0.0 detail
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best

=====
BGP MVPN-IPv4 Routes
=====
Route Type      : Source-Join
Route Dist.     : 1:1
Source AS       : 200
Source IP       : 150.100.1.2
Group IP        : 226.0.0.0
Nextthop       : 10.20.1.4
From            : 10.20.1.4
Res. Nextthop   : 0.0.0.0
Local Pref.     : 100
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
Community       : target:10.20.1.3:2
Cluster         : No Cluster Members
Originator Id   : None
Flags           : Used Valid Best IGP
AS-Path         : No As-Path
Interface Name  : NotAvailable
Aggregator      : None
MED             : 0
Peer Router Id  : 10.20.1.4

-----
Routes : 1
=====
*A:praragon-sim1#

*A:Dut-C# show router bgp routes mvpn-ipv4 type spmsi-ad detail
=====
BGP Router ID:10.20.1.3      AS:46000      Local AS:46000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
              l - leaked
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP MVPN-IPv4 Routes
=====
Original Attributes

Route Type      : Spmsi-Ad
Route Dist.     : 10.1.200.41:1
Originator IP   : 10.20.1.4
Source IP       : 10.1.101.2
Group IP        : 225.100.0.0

<snip>

Last Modified   : 00h18m52s
VPRN Imported   : 1

-----
PMSI Tunnel Attribute :
Tunnel-type      : None
Flags            : Leaf required

```

```

MPLS Label      : 0
-----

=====
*A:Dut-C#

*A:Dut-C# show router bgp routes mvpn-ipv6 type spmsi-ad detail
=====
BGP Router ID:10.20.1.3      AS:46000      Local AS:46000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP MVPN-IPv6 Routes
=====
Original Attributes

Route Type      : Spmsi-Ad
Route Dist.     : 10.1.200.41:1
Originator IP   : 10.20.1.4
Source IP       : 2001:10:1:101::2
Group IP        : ff0e:225:100::

<snip>

VPRN Imported   : 1
-----

PMSI Tunnel Attribute :
Tunnel-type        : None                      Flags          : Leaf required
MPLS Label         : 0
-----

=====
*A:Dut-C#

*A:Dut-C# show router bgp routes ms-pw
=====
BGP Router ID:10.20.1.3      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP MSPW Routes
=====
Flag  Network      RD
      Nexthop      AII-Type2/Preflen
      As-Path
-----
?    3:10.20.1.3    100:3
      10.20.1.5     3:10.20.1.3:0/64
      200 100
?    3:10.20.1.3    100:4
      10.20.1.5     3:10.20.1.3:0/64
      200 100
u*>? 6:10.20.1.6     100:6

```



```

10.20.1.5          6:10.20.1.6:0/64
200 300 400
-----
Routes : 3
=====

*A:DUT# show router bgp routes ipv4 detail
=====
BGP Router ID:1.1.1.1          AS:100          Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
Original Attributes

Network      : 11.1.1.1/32
Nexthop      : 192.168.1.1
Path Id      : None
From         : 192.168.1.1
Res. Nexthop : 192.168.1.1
Local Pref.  : n/a
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : 100
Community    : None
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
Flags        : Used Valid Best Incomplete
Route Source : External
AS-Path      : 200 400 500
Interface Name : net
Aggregator    : None
MED          : 5000
Peer Router Id : 2.2.2.2
Priority      : None

Modified Attributes

Network      : 11.1.1.1/32
Nexthop      : 192.168.1.1
Path Id      : None
From         : 192.168.1.1
Res. Nexthop : 192.168.1.1
Local Pref.  : None
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : 110
Community    : None
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
Flags        : Used Valid Best Incomplete
Route Source : External
AS-Path      : 200 400 500
Interface Name : net
Aggregator    : None
MED          : 5000
Peer Router Id : 2.2.2.2
Priority      : None
-----
Routes : 1
=====

```

```

*A:DUT# show router bgp routes 11.1.1.1/32 hunt
=====
BGP Router ID:1.1.1.1          AS:100          Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
RIB In Entries
-----
Network      : 11.1.1.1/32
Nextthop     : 192.168.1.1
Path Id      : None
From         : 192.168.1.1
Res. Nextthop : 192.168.1.1
Local Pref.  : None
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : 110
Community    : None
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
Flags        : Used Valid Best Incomplete
Route Source  : External
AS-Path       : 200 400 500
-----
RIB Out Entries
-----
Network      : 11.1.1.1/32
Nextthop     : 1.1.1.1
Path Id      : None
To           : 3.3.3.3
Res. Nextthop : n/a
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : 150
Community    : None
Cluster      : No Cluster Members
Originator Id : None
Origin       : Incomplete
AS-Path       : 200 400 500
-----
Routes : 2
=====

*A:PCE Server 226# show router bgp routes bgp-ls hunt link
=====
BGP Router ID:38.120.48.226    AS:65000      Local AS:65000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
              l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete
=====

```

```

BGP-LS Link NLRIs
=====
-----
RIB In Entries
-----
Network:
Type           : LINK-NLRI
Protocol       : ISIS Level-2          Identifier      : 0xa
Local Node descriptor:
  Autonomous System : 0.0.253.232
  Link State Id     : 10
  IGP Router Id     : 0x38120048184
Remote Node descriptor:
  Autonomous System : 0.0.253.232
  Link State Id     : 10
  IGP Router Id     : 0x38120048223
Link descriptor:
  IPV4 Interface Addr: 10.0.14.184
  IPV4 Neighbor Addr : 10.0.14.223
Nexthop        : 38.120.48.199
From           : 38.120.48.199
Res. Nexthop   : 0.0.0.0
Local Pref.    : 100
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
AIGP Metric    : None
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None
Flags          : Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0
Add Paths Send : Default
Last Modified  : 02h27m50s
Interface Name : NotAvailable
Aggregator     : None
MED            : None
Peer Router Id : 38.120.48.199
Dest Class     : 0
-----
Link State Attribute TLVs :
Administrative group (color) : 0x0
Maximum link bandwidth : 100000 Kbps
Max. reservable link bandwidth : 100000 Kbps
Unreserved bandwidth0 : 100000 Kbps
Unreserved bandwidth1 : 100000 Kbps
Unreserved bandwidth2 : 100000 Kbps
Unreserved bandwidth3 : 100000 Kbps
Unreserved bandwidth4 : 100000 Kbps
Unreserved bandwidth5 : 100000 Kbps
Unreserved bandwidth6 : 100000 Kbps
Unreserved bandwidth7 : 100000 Kbps
TE Default Metric : 100
IGP Metric : 100
Adjacency Segment Identifier (Adj-SID) : flags 0x30 weight 0 sid 262136
-----
Network:
Type           : LINK-NLRI
Protocol       : ISIS Level-2          Identifier      : 0xa

```

```

Local Node descriptor:
  Autonomous System : 0.0.253.232
  Link State Id      : 10
  IGP Router Id      : 0x38120048184
Remote Node descriptor:
  Autonomous System : 0.0.253.232
  Link State Id      : 10
  IGP Router Id      : 0x38120048223
Link descriptor:
  IPV4 Interface Addr: 10.0.14.184
  IPV4 Neighbor Addr : 10.0.14.223
Nexthop      : 38.120.48.221
From         : 38.120.48.221
Res. Nexthop  : 0.0.0.0
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : None
Connector     : None
Community     : No Community Members
Cluster       : No Cluster Members
Originator Id : None
Flags         : Valid IGP
TieBreakReason : OriginatorID
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0
Add Paths Send : Default
Last Modified  : 02h27m54s
Interface Name : NotAvailable
Aggregator     : None
MED            : None
Peer Router Id : 38.120.48.221
Dest Class     : 0

-----
Link State Attribute TLVs :
Administrative group (color) : 0x0
Maximum link bandwidth : 100000 Kbps
Max. reservable link bandwidth : 100000 Kbps
Unreserved bandwidth0 : 100000 Kbps
Unreserved bandwidth1 : 100000 Kbps
Unreserved bandwidth2 : 100000 Kbps
Unreserved bandwidth3 : 100000 Kbps
Unreserved bandwidth4 : 100000 Kbps
Unreserved bandwidth5 : 100000 Kbps
Unreserved bandwidth6 : 100000 Kbps
Unreserved bandwidth7 : 100000 Kbps
TE Default Metric : 100
IGP Metric : 100
Adjacency Segment Identifier (Adj-SID) : flags 0x30 weight 0 sid 262136
-----

*A:Dut-A>config>router>bgp# show router bgp routes
=====
BGP Router ID:10.20.1.1      AS:1      Local AS:1
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes

```

```

=====
Flag   Network                               LocalPref   MED
      Nexthop                               Path-Id     Label
      As-Path
-----
u*>i   20.0.0.1/32                           100         2010
      10.20.1.2                             None        131057
      2
ub*i   20.0.0.1/32                           100         2010
      10.20.1.3                             None        131067
      2
-----
Routes : 2
=====
*A:DUT-A>config>router>bgp#

*A:Dut-A# show router bgp routes evpn mac mac-address 00:00:01:00:01:02 hunt
=====
BGP Router ID:10.20.1.1      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP EVPN Mac Routes
=====
-----
RIB In Entries
-----
Network      : N/A
Nexthop      : 10.20.1.2
From         : 10.20.1.2
Res. Nexthop : N/A
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : target:100:1 bgp-tunnel-encap:VXLAN
               mac-mobility:Seq:0/Static
Cluster      : No Cluster Members
Originator Id : None
Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : 111
EVPN type     : MAC
ESI           : 0:0:0:0:0:0:0:0:0      Tag           : 1
IP Address    : N/A                    Route Dist.    : 10.20.1.2:1
Mac Address   : 00:00:01:00:01:02
MPLS Label1  : X                      MPLS Label2    : Y
Route Tag     : Z
Neighbor-AS   : 111
Orig Validation: N/A
Source Class  : 0                      Dest Class     : 0

-----
RIB Out Entries
-----

```

```
-----
Routes : 1
=====
```

```
*A:Dut-A# show router bgp routes evpn ip-prefix prefix 3.0.1.6/32 detail
```

```
=====
BGP Router ID:10.20.1.1      AS:100      Local AS:100
=====
```

```
Legend -
```

```
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
```

```
=====
BGP EVPN IP-Prefix Routes
=====
```

```
-----
Original Attributes
```

```
Network      : N/A
Nexthop      : 10.20.1.2
From         : 10.20.1.2
Res. Nexthop : N/A
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : target:100:1 mac-nh:00:00:01:00:01:02
               bgp-tunnel-encap:VXLAN
Cluster      : No Cluster Members
Originator Id : None
Flags        : Used Valid Best IGP
Route Source : Internal
AS-Path      : No As-Path
EVPN type    : IP-PREFIX
ESI          : N/A
Gateway Address: 00:00:01:00:01:02
Prefix       : 3.0.1.6/32
MPLS Label   : X
Route Tag    : Z
Neighbor-AS  : N/A
Orig Validation: N/A
Source Class : 0

Interface Name : NotAvailable
Aggregator     : None
MED            : 0
Peer Router Id : 10.20.1.2
Tag            : 1
Route Dist.    : 10.20.1.2:1
Dest Class     : 0
```

```
Modified Attributes
```

```
Network      : N/A
Nexthop      : 10.20.1.2
From         : 10.20.1.2
Res. Nexthop : N/A
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : target:100:1 mac-nh:00:00:01:00:01:02
               bgp-tunnel-encap:VXLAN
Cluster      : No Cluster Members
Originator Id : None

Interface Name : NotAvailable
Aggregator     : None
MED            : 0
Peer Router Id : 10.20.1.2
```

```

Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 111
EVPN type      : IP-PREFIX
ESI            : N/A
Gateway Address: 00:00:01:00:01:02
Prefix         : 3.0.1.6/32
MPLS Label     : X
Route Tag      : W
Neighbor-AS    : 111
Orig Validation: N/A
Source Class   : 0
Tag            : 1
Route Dist.    : 10.20.1.2:1
Dest Class     : 0

```

```

-----
Routes : 1
=====

```

```

*A:7750SR7-PE# show router bgp routes bgp-ls

```

```

=====
BGP-LS NLRIs
=====
-----

```

```

Type: Node
Protocol: ISIS | OSPFLevel: L1 | L2 | NA
Identifier:
Network : 215.0.0.0/24
Nextthop : 202.50.0.2
Path Id : None
From : 10.255.105.141
AS Path: 65534
Local Node
  Name:<= Do we populate this?
  Autonomous System:
  Area ID:
  Rtr ID:
  Flags:

```

```

Type: Link
Protocol: ISIS | OSPFLevel: L1 | L2 | NA
Identifier:
Network : 215.0.0.0/24
Nextthop : 202.50.0.2
Path Id : None
From : 10.255.105.141
AS Path: 65534
Local Node
  Name:
  Autonomous System:
  Area ID:
  Rtr ID:

```

```

Remote Node
  Name:
  Autonomous System:
  Area ID:
  Rtr ID:

```

```

Descriptor
  IPv4 Address: 10.10.1.1

```

```

Attributes
  Flags:
  IGP Metric: 1000

```

```
Admin-Grp: 0
Maximum Link BW (kbps): 100000
Maximum Reservable BW (kbps): 100000
TE Default Metric: 1000
Unreserved BW (kbps): 10000

Type: IPv4 Prefix
Protocol: ISIS | OSPFLevel: L1 | L2 | NA
Identifier:
Network : 215.0.0.0/24
Nexthop : 202.50.0.2
Path Id : None
From : 10.255.105.141
AS Path: 65534
Local Node
  Name:<= Do we populate this?
  Autonomous System:
  Area ID: 0.0.0.0
  Rtr ID: 10.20.1.6
  Flags:
Prefix descriptor
  IPv4 Address: 10.20.1.6
  Prefix-len: 32
  MT_ID :0
  Route Type :Unknown
```

summary

Syntax	summary [all] summary [family <i>family</i>] [neighbor <i>ip-address</i>]
Context	show>router>bgp
Description	This command displays a summary of BGP neighbor information. If confederations are not configured, that portion of the output will not display. The “State” field displays the global BGP operational state. The valid values are: • Up — BGP global process is configured and running. • Down — BGP global process is administratively shutdown and not running. • Disabled — BGP global process is operationally disabled. The process must be restarted by the operator. For example, if a BGP peer is operationally disabled, then the state in the summary table shows the state ‘Disabled’
Parameters	family — Specify the type of routing information to be distributed by the BGP instance. Values ipv4 — Displays only those BGP peers that have the IPv4 family enabled. vpn-ipv4 — Displays the BGP peers that are IP-VPN capable. ipv6 — Displays the BGP peers that are IPv6 capable.

mcast-ipv4 — Displays the BGP peers that are mcast-ipv4 capable.

neighbor *ip-address* — Clears damping information for entries received from the BGP neighbor.

Values

ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

Output BGP Summary Output

The following table describes the command output fields for a BGP summary.

Table 64 BGP Summary Output Fields

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
BGP Admin State	Down — BGP is administratively disabled. Up — BGP is administratively enabled.
BGP Oper State	Down — BGP is operationally disabled. Up — BGP is operationally enabled.
Bfd	Yes — BFD is enabled. No — BFD is disabled.
Confederation AS	The configured confederation AS.
Member Confederations	The configured members of the BGP confederation.
Number of Peer Groups	The total number of configured BGP peer groups.
Number of Peers	The total number of configured BGP peers.
Total BGP Active Routes	The total number of BGP routes used in the forwarding table.
Total BGP Routes	The total number of BGP routes learned from BGP peers.

Table 64 BGP Summary Output Fields (Continued)

Label	Description
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers.
Total Path Memory	Total amount of memory used to store the path attributes.
Total Suppressed Routes	Total number of suppressed routes due to route damping.
Total History Routes	Total number of routes with history due to route damping.
Total Decayed Routes	Total number of decayed routes due to route damping.
Total VPN Peer Groups	The total number of configured VPN peer groups.
Total VPN Peers	The total number of configured VPN peers.
Total VPN Local Rts	The total number of configured local VPN routes.
Total VPN Remote Rts	The total number of configured remote VPN routes.
Total VPN Remote Active Rts.	The total number of active remote VPN routes used in the forwarding table.
Total VPN Supp.Rts.	Total number of suppressed VPN routes due to route damping.
Total VPN Hist. Rts.	Total number of VPN routes with history due to route damping.
Total VPN Decay Rts.	Total number of decayed routes due to route damping.
Neighbor	BGP neighbor address.
AS (Neighbor)	BGP neighbor autonomous system number.
PktRcvd	Total number of packets received from the BGP neighbor.
PktSent	Total number of packets sent to the BGP neighbor.
InQ	The number of BGP messages to be processed.
OutQ	The number of BGP messages to be transmitted.
Up/Down	The amount of time that the BGP neighbor has either been established or not established depending on its current state.
State Recv/Actv/Sent	The BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established).

Sample Output

```
A:Dut-C# show router bgp summary neighbor 3FFE::A0A:1064
=====
BGP Router ID : 10.20.1.3          AS : 100      Local AS : 100
=====
BGP Admin State      : Up          BGP Oper State      : Up
Number of Peer Groups : 4          Number of Peers      : 5
Total BGP Paths       : 8          Total Path Memory    : 1212
Total BGP Active Rts. : 0          Total BGP Rts.       : 0
Total Suppressed Rts. : 0          Total Hist. Rts.     : 0
Total Decay Rts.      : 0

Total VPN Peer Groups : 0          Total VPN Peers      : 0
Total VPN Local Rts.  : 0
Total VPN Remote Rts. : 0          Total VPN Remote Active Rts.: 0
Total VPN Supp. Rts.  : 0          Total VPN Hist. Rts.  : 0
Total VPN Decay Rts.  : 0

Total IPv6 Remote Rts. : 5          Total IPv6 Rem. Active Rts. : 4
=====
BGP Summary
=====
Neighbor
      AS      PktRcvd InQ  Up/Down  State| Rcv/Act/Sent (IPv4)
              PktSent OutQ              Rcv/Act/Sent (VpnIPv4)
                                      Rcv/Act/Sent (IPv6)
                                      Rcv/Act/Sent (MCastIPv4)
-----
3FFE::A0A:1064
          103      489    0 00h40m28s IPv4 Incapable
                              569    0      VPN-IPv4 Incapable
                              1/1/3
                              MCAST-IPv4 Incapable
=====
A:Dut-C#

A:Dut-C# show router bgp summary neighbor 10.20.1.4 family ipv6
=====
BGP Router ID : 10.20.1.3          AS : 100      Local AS : 100
=====
BGP Admin State      : Up          BGP Oper State      : Up
Number of Peer Groups : 4          Number of Peers      : 5
Total BGP Paths       : 8          Total Path Memory    : 1212
Total BGP Active Rts. : 0          Total BGP Rts.       : 0
Total Suppressed Rts. : 0          Total Hist. Rts.     : 0
Total Decay Rts.      : 0

Total VPN Peer Groups : 0          Total VPN Peers      : 0
Total VPN Local Rts.  : 0
Total VPN Remote Rts. : 0          Total VPN Remote Active Rts.: 0
Total VPN Supp. Rts.  : 0          Total VPN Hist. Rts.  : 0
Total VPN Decay Rts.  : 0

Total IPv6 Remote Rts. : 5          Total IPv6 Rem. Active Rts. : 4
=====
BGP IPv6 Summary
=====
```

```

Neighbor
      AS PktRcvd PktSent  InQ OutQ Up/Down  State|Rcv/Actv/Sent
-----
10.20.1.4
      100      554      572   0   0 00h41m27s 1/0/3
=====
A:Dut-C#

A:SetupCLI>show>router# bgp summary
=====
  BGP Router ID : 21.3.4.5          AS : 35012    Local AS : 100
=====
BGP Admin State      : Up          BGP Oper State      : Up
Confederation AS     : 40000
Member Confederations : 35012 65205 65206 65207 65208
Rapid Withdrawal     : Disabled
Bfd Enabled          : Yes

Number of Peer Groups : 1          Number of Peers      : 1
Total BGP Paths       : 3          Total Path Memory    : 396
Total BGP Active Rts. : 0          Total BGP Rts.       : 0
Total Supressed Rts.  : 0          Total Hist. Rts.     : 0
Total Decay Rts.      : 0

Total VPN Peer Groups : 1          Total VPN Peers      : 1
Total VPN Local Rts.  : 0
Total VPN Remote Rts. : 0          Total VPN Remote Active Rts.: 0
Total VPN Supp. Rts.  : 0          Total VPN Hist. Rts.  : 0
Total VPN Decay Rts.  : 0

Total IPv6 Remote Rts. : 0          Total IPv6 Rem. Active Rts. : 0
=====
BGP Summary
=====
Legend : D - Dynamic Neighbor
=====
Neighbor
      AS  PktRcvd InQ  Up/Down  State|Rcv/Act/Sent (IPv4)
      PktSent OutQ
      Rcv/Act/Sent (VpnIPv4)
      Rcv/Act/Sent (IPv6)
      Rcv/Act/Sent (MCastIPv4)
-----
3.3.3.3      20      0   0   01h55m56s Active
              0   0
=====
A:SetupCLI>show>router#

```

fib

Syntax **fib slot-number** [*family*] [*ip-prefix/prefix-length* [**longer**]] [**secondary**] [**qos**] [**accounting-class**] [**all**]
fib slot-number extensive [*ip-prefix/prefix-length*] [*family*] [**all**]
fib slot-number [*family*] **summary**
fib slot-number nh-table-usage

fib all summary

Context	show>router>fib
Description	This command displays FIB information for a specific IOM.
Parameters	<i>slot-number</i> — Specifies the slot number. Values 1 to 10 <i>family</i> — Specifies information about the type of routing information to be distributed by the instance. Values ipv4 — Displays only those peers that have the IPv4 family enabled. ipv6 — Displays the peers that are IPv6 capable. <i>ip-prefix</i> — Specifies to display IP prefix for prefix list entry in dotted decimal notation information. Values ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D <i>prefix-length</i> — Specifies prefix length. ipv4-prefix-length: 0 to 32 ipv6-prefix-length: 0 to 128 longer — Specifies the prefix list entry matches any route that matches the specified ip-prefix and prefix mask length values equal to or greater than the specified mask. secondary — Specifies a secondary FIB. summary — Displays a summary of the FIB information. nh-table-usage — Displays next-hop table usage. qos — Displays QoS information. accounting-class — Displays accounting class information. extensive — Displays detailed output. all — Displays output for all IP prefixes.
Output	The following output is an example of router FIB information.

Sample Output

```
*A:pe1# show router fib 1
=====
FIB Display
```

```

=====
Prefix [Flags] Protocol
NextHop
Src-Class
-----
180.10.0.1/32 [S] BGP
10.10.10.1 Indirect (lag1-to-server1)
10.10.10.5 Indirect (lag2-to-server2)
10.10.10.10 Indirect (lag3-to-server3)
-----
Total Entries : 1
Flags: S = sticky ECMP supported
-----
show router fib 1 nh-table-usage
=====
FIB Next-Hop Summary
=====
IPv4/IPv6 Active Available
-----
IP Next-Hop 4 16383
Tunnel Next-Hop 16519 993279
ECMP Next-Hop 511998 512000
ECMP Tunnel Next-Hop 33030 261120
=====

A:pe1# show router fib 1 summary
=====
FIB Summary
=====
Active
-----
Static 11
Direct 5
HOST 0
BGP 579643
BGP VPN 0
BGP EVPN 0
OSPF 0
ISIS 0
RIP 0
RIP_NG 0
LDP 0
Aggregate 1
Sub Mgmt 0
VPN Leak 0
NAT 0
Managed 0
Periodic 0
dhcpv6-pd 0
dhcpv6-bh 0
dhcpv6-na 0
dhcpv6-ta 0
-----
Total Installed 579660
-----
Current Occupancy 7%
Overflow Count 0
Suppressed by Selective FIB 0

```

```
Occupancy Threshold Alerts
  Alert Raised 0 Times;
=====
```

mvpn

Syntax	mvpn
Context	show>router
Description	This command displays Multicast VPN related information.
Output	The following output displays MVPN information.

Sample Output

```
*A:praragon-sim1# show router 100 mvpn
=====
MVPN 100 configuration data
=====
i-pmsi          : 224.100.201.101 ssm  admin status      : Up
hello-interval   : 30 seconds           hello-multiplier   : 35 * 0.1
three-way-hello  : Disabled             tracking support    : Disabled

s-pmsi range     : 0.0.0.0/0             data-delay-interval: 3 seconds
join-tlv-packing : N/A

signaling        : Bgp
vrf-import       : N/A
vrf-export       : N/A
vrf-target       : N/A
=====
*A:praragon-sim1#
```

route-table

Syntax	route-table [<i>family</i>] [<i>ip-prefix/prefix-length</i>] [longer exact protocol <i>protocol-name</i>] [all] [next-hop-type <i>type</i>] [qos] [alternative] [accounting-class] route-table [<i>family</i>] summary route-table <i>tunnel-endpoints</i> [<i>ip-prefix/prefix-length</i>] [longer exact] [detail] route-table [<i>family</i>] [<i>ip-prefix/prefix-length</i>] [longer exact protocol <i>protocol-name</i>] extensive [all]
Context	show>router>route-table
Description	This command displays route-table information.
Parameters	<i>family</i> — Specifies the type of routing information to be distributed by the BGP instance. ipv4 — Displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes.

- mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.
- ipv6** — Displays the BGP peers that are IPv6 capable.
- mcast-ipv6** — Displays the BGP peers that are mcast-ipv6 capable.

ip-prefix — Displays information about the IP prefix for prefix list entry in dotted decimal notation.

- Values**
- ipv4-prefix:
 - a.b.c.d (host bits must be 0)
 - ipv6-address:
 - x:x:x:x:x:x:x (eight 16-bit pieces)
 - x:x:x:x:x:d.d.d.d
 - x: [0 to FFFF]H
 - d: [0 to 255]D

prefix-length — Specifies prefix length.

- ipv4-prefix-length: 0 to 32
- ipv6-prefix-length: 0 to 128

longer — Specifies the prefix list entry matches any route that matches the specified ip-prefix and prefix mask length values equal to or greater than the specified mask.

exact — Specifies the prefix list entry matches any route that matches the specified ip-prefix and prefix mask length values of the specified mask.

protocol-name — Specifies the protocol name.

- Values**
- local, sub-mgmt, managed, static, ldp, ospf, ospf3, isis, rip, bgp, bgp-vpn, bgp-evpn, aggregate, vpn-leak, nat, periodic, dhcpv6-pd, dhcpv6-na, dhcpv6-ta, dhcpv6-pd-excl, ripng, ipsec, bgp-label, host, arp-nd

all — Displays information about all inactive routes.

alternative — Displays LFA and backup route details.

detail — Displays detailed information.

extensive — Displays extensive route table information.

Output The following output displays route table information.

Sample Output

```
show router route-table 180.10.0.1/32
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                       Metric
-----
180.10.0.1/32 [B] [S]                             Remote BGP      00h00m49s    170
      10.10.10.1                                     0
180.10.0.1/32 [B] [S]                             Remote BGP      00h00m49s    170
```



```

10.10.10.5
180.10.0.1/32 [B] [S] Remote BGP 00h00m49s 170
10.10.10.10 0
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = sticky ECMP requested
=====
*A:VRR>config>router# show router route-table extensive 200.200.200.200/32

```

Examples of the **show router route-table extensive** command output with unequal-cost ECMP BGP routes are shown below.

```

=====
Route Table (Router: Base)
=====
Dest Prefix      : 200.200.200.200/32
Protocol         : BGP
Age              : 01h20m41s
Preference       : 170
Indirect Next-Hop : 10.0.0.2
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  ECMP-Weight     : 9
  Resolving Next-Hop : 10.0.0.2
    Interface     : to_bridge_br2
    Metric        : 0
    ECMP-Weight   : N/A
Indirect Next-Hop : 192.0.2.2
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  ECMP-Weight     : 5
  Resolving Next-Hop : 192.0.2.2
    Interface     : to_bridge_br3
    Metric        : 0
    ECMP-Weight   : N/A
-----
No. of Destinations: 1
=====
*A:VRR>config>router# show router fib 1 extensive 200.200.200.200/32

```

```

=====
FIB Display (Router: Base)
=====
Dest Prefix      : 200.200.200.200/32
Protocol         : BGP
Indirect Next-Hop : 10.0.0.2
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  ECMP-Weight     : 9
  Resolving Next-Hop : 10.0.0.2
    Interface     : to_bridge_br2
    ECMP-Weight   : 1

```

```
Indirect Next-Hop      : 192.0.2.2
  QoS                  : Priority=n/c, FC=n/c
  Source-Class         : 0
  Dest-Class           : 0
  ECMP-Weight          : 5
  Resolving Next-Hop   : 192.0.2.2
    Interface          : to_bridge_br3
    ECMP-Weight        : 1
=====
Total Entries : 1
=====
*A:Dut-C>config>router>mpls>lsp# show router route-table 5.3.0.1/32 extensive
=====
Route Table (Router: Base)
=====
Dest Prefix            : 5.3.0.1/32
  Protocol              : BGP
  Age                   : 00h00m59s
  Preference            : 170
  Indirect Next-Hop     : 10.0.0.1
    QoS                 : Priority=n/c, FC=n/c
    Source-Class        : 0
    Dest-Class          : 0
    ECMP-Weight         : 1
    Resolving Next-Hop  : 1.0.0.2 (RSVP tunnel:115)
      Metric            : 10
      ECMP-Weight       : 1
    Resolving Next-Hop  : 1.0.0.2 (RSVP tunnel:61443)
      Metric            : 10
      ECMP-Weight       : 1
  Indirect Next-Hop     : 10.0.0.2
    QoS                 : Priority=n/c, FC=n/c
    Source-Class        : 0
    Dest-Class          : 0
    ECMP-Weight         : 30
    Resolving Next-Hop  : 1.0.0.3 (RSVP tunnel:94)
      Metric            : 10
      ECMP-Weight       : 20
    Resolving Next-Hop  : 1.0.0.3 (RSVP tunnel:61442)
      Metric            : 10
      ECMP-Weight       : 1
-----
No. of Destinations: 1
=====
```

station

Syntax	station [station-name]
Context	show>router>bmp
Description	This command displays BGP monitor (BMP) information.
Parameters	station-name — Specifies the station name of a BMP monitoring station, up to 32 characters.

Output The following is an example of BMP) monitor information.

Sample Output

```
*A:Dut-C# show router bmp station
=====
BMP Stations (monitoring router "Base")
=====
Station Name                Session State Since                Msgs Sent
-----
lys                          ESTABLISHED    05/08/2018 15:32:12    103054
myr                          ESTABLISHED    05/08/2018 15:32:13    100007
-----
No. of BMP Stations: 2
=====
*A:Dut-C# show router bmp station "lys"
=====
BMP Station "lys" (monitoring router "Base")
=====
Admin State      : enabled                Global BMP State : enabled
Station Address  : 100.3.7.7              Station Port     : 12344
Via Router       : Base                  TCP Auth Keychain:
Stats Report     : 15 seconds            Routes Report    : 5 seconds
Connect Interval : 120 seconds           Local Routes     : reporting
Initiation Msg   : Hello Station Lys
Reported families: ipv4
Session State    : ESTABLISHED           Last Change      : 05/08/2018 15:32:12
Local Address    : 100.3.7.3             Local Port       : 51041
Stat Report Msgs : 3028                  Route Report Msgs: 100021
Peer Up Msgs     : 4                     Peer Down Msgs   : 0
Initiation Msgs  : 1                     Goodbye Msgs     : 0
Bytes sent       : 10926521
Stats Timer      : 1 seconds left         Routes Timer     : 3 seconds left
Connect Timer    : not running            Monitored Peers  : 3 of 4
Output queue     : 0/50                   Last Msg Sent    : 05/15/2018 09:27:31
=====
*A:Dut-C#
```

5.14.2.2 Clear Commands

damping

Syntax	damping [{ <i>ip-prefix</i> / <i>ip-prefix-length</i>] [<i>neighbor ip-address</i> }] group name
Context	clear>router>bgp
Description	This command clears or resets the route damping information for received routes.
Parameters	<i>ip-prefix/ip-prefix-length</i> — Clears damping information for entries that match the IP prefix and prefix length.
Values	ipv4-prefix:

- a.b.c.d (host bits must be 0)
- ipv4-prefix-length: [0 to 32]
- ipv6-prefix:
- x:x:x:x:x:x:x (eight 16-bit pieces)
 - x:x:x:x:x:d.d.d.d
 - x: [0 to FFFF]H
 - d: [0 to 255]D
- ipv6-prefix-length: [0 to 128]

neighbor *ip-address* — Clears damping information for entries received from the BGP neighbor.

Values ipv4-address:

- a.b.c.d

ipv6-prefix:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

interface: 32 characters maximum, mandatory for link local addresses

group *name* — Clears damping information for entries received from any BGP neighbors in the peer group.

Values 32 characters maximum

flap-statistics

Syntax **flap-statistics** [{*ip-prefix/mask* [**neighbor** *ip-address*]}] | **group** *group-name* | [**regex** *reg-exp*] | [**policy** *policy-name*]

Context clear>router>bgp

Description This command clears route flap statistics.

Parameters *ip-prefix/mask* — Clears route flap statistics for entries that match the specified IP prefix and mask length.

Values ip-prefix: a.b.c.d (host bits must be 0)
mask: 0 to 32

neighbor *ip-address* — Clears route flap statistics for entries received from the specified BGP neighbor.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

group *group-name* — Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.

regex *reg-exp* — Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression.

policy *policy-name* — Clears route flap statistics for entries that match the specified route policy.

neighbor

Syntax	neighbor { <i>ip-address</i> as <i>as-number</i> external all } [soft soft-inbound hard] neighbor { <i>ip-address</i> as <i>as-number</i> external all } statistics neighbor <i>ip-address</i> end-of-rib
Context	clear>router>bgp
Description	This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shutdown and restarted.
Parameters	<i>ip-address</i> — Resets the BGP neighbor with the specified IP address. Values ipv4-address: • a.b.c.d ipv6-address: • x:x:x:x:x:x:x [-interface] • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: up to 32 characters, mandatory for link local addresses as <i>as-number</i> — Resets all BGP neighbors with the specified peer AS. Values 1 to 65535 external — Resets all EBGp neighbors. all — Resets all BGP neighbors. soft — The specified BGP neighbor(s) re-evaluates all routes in the Local-RIB against the configured export policies. soft-inbound — The specified BGP neighbor(s) re-evaluates all routes in the RIB-In against the configured import policies.

hard — Initiates a hard reconfiguration.

statistics — The BGP neighbor statistics.

end-of-rib — Clears the routing information base (RIB). This command applies when the router is helping the BGP neighbor through a BGP graceful restart. When the **clear router bgp neighbor** command is issued without the end-of-rib option and the neighbor is in the process of restarting, stale routes from the neighbor will be retained until the stale-routes-time is reached or else the neighbor exits graceful restart. When the command is issued with the end-of-rib option, stale routes from the neighbor are deleted immediately and graceful restart procedures are aborted.

protocol

Syntax	protocol
Context	clear>router>bgp
Description	Resets the entire BGP protocol.

station

Syntax	station <i>station-name</i> station all
Context	clear>router>bgp
Description	This command resets the specified station connection if one is specified or all connections if the all keyword is specified.
Parameters	<i>station-name</i> — Specifies up to station names, up to 32 characters, association with the BMP monitoring station. all — Specifies all configured stations.

5.14.2.3 Debug Commands

events

Syntax	events [neighbor <i>ip-address</i> group <i>name</i>] no events
Context	debug>router>bgp
Description	This command logs all events changing the state of a BGP peer.

Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

graceful-restart

Syntax	graceful-restart [neighbor <i>ip-address</i> group <i>name</i>] no graceful-restart
Context	debug>router>bgp
Description	This command enables debugging for BGP graceful-restart. The no form of the command disables the debugging.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor. Values ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

keepalive

Syntax	keepalive [neighbor <i>ip-addr</i> group <i>name</i>] no keepalive
Context	debug>router>bgp

Description	This command decodes and logs all sent and received keepalive messages in the debug log.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

notification

Syntax	notification [neighbor <i>ip-address</i> group <i>name</i>] no notification
Context	debug>router>bgp
Description	This command decodes and logs all sent and received notification messages in the debug log.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

open

Syntax	open [neighbor <i>ip-address</i> group <i>name</i>] no open

Context	debug>router>bgp
Description	This command decodes and logs all sent and received open messages in the debug log.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

outbound-route-filtering

Syntax	[no] outbound-route-filtering
Context	debug>router>bgp
Description	This command enables debugging for all BGP outbound route filtering (ORF) packets. ORF is used to inform a neighbor of targets (using target-list) that it is willing to receive.

packets

Syntax	packets [neighbor <i>ip-address</i> group <i>name</i>] packets
Context	debug>router>bgp
Description	This command decodes and logs all sent and received BGP packets in the debug log.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses

group name — Debugs only events affecting the specified peer group and associated neighbors.

route-refresh

Syntax	route-refresh [neighbor <i>ip-address</i> group name] no route-refresh
Context	debug>router>bgp
Description	This command enables and disables debugging for BGP route-refresh.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group name — Debugs only events affecting the specified peer group and associated neighbors.

rtm

Syntax	rtm [neighbor <i>ip-address</i> group name] no rtm
Context	debug>router>bgp
Description	This command logs RTM changes in the debug log.
Parameters	neighbor <i>ip-address</i> — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses

group name — Debugs only events affecting the specified peer group and associated neighbors.

socket

Syntax	socket [<i>neighbor ip-address</i> group name] no socket
Context	debug>router>bgp
Description	This command logs all TCP socket events to the debug log.
Parameters	neighbor ip-address — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses
	group name — Debugs only events affecting the specified peer group and associated neighbors.

timers

Syntax	timers [<i>neighbor ip-address</i> group name] no timers
Context	debug>router>bgp
Description	This command logs all BGP timer events to the debug log.
Parameters	neighbor ip-address — Debugs only events affecting the specified BGP neighbor.
Values	ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses

group name — Debugs only events affecting the specified peer group and associated neighbors.

update

Syntax	update [<i>neighbor ip-address</i> group name] no update
Context	debug>router>bgp
Description	This command decodes and logs all sent and received update messages in the debug log.
Parameters	neighbor ip-address — Debugs only events affecting the specified BGP neighbor. Values ipv4-address: • a.b.c.d (host bits must be 0) ipv6-address: • x:x:x:x:x:x:x [-interface] (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d [-interface] • x: [0 to FFFF]H • d: [0 to 255]D • interface: 32 characters maximum for link local addresses group name — Debugs only events affecting the specified peer group and associated neighbors.

bmp

Syntax	bmp
Context	debug>router
Description	This command enables the context to debug BMP information.

events

Syntax	events [<i>station station-name</i>] no events
Context	debug>router>bmp
Description	This command enables debugging for all BMP events. The no form of the command disables debugging for all BMP events.

Parameters	<i>station-name</i> — Specifies the station name of the BMP monitoring station, up to 32 characters.

packets

Syntax	packets [<i>station station-name</i>] no packets
Context	debug>router>bmp
Description	This command enables debugging for all BMP packets. The no form of the command disables debugging for all BMP packets.
Parameters	<i>station-name</i> — Specifies the station name of the BMP monitoring station, up to 32 characters.

6 Route Policies

6.1 Configuring Route Policies

Nokia's router supports two databases for routing information. The routing database is composed of the routing information learned by the routing protocols. The forwarding database is composed of the routes actually used to forward traffic through a router. In addition, link state databases are maintained by interior gateway protocols (IGPs) such as IS-IS and OSPF.

Routing protocols calculate the best route to each destination and place these routes in a forwarding table. The routes in the forwarding table are used to forward routing protocol traffic, sending advertisements to neighbors and peers.

A routing policy can be configured that will not place routes associated with a specific origin in the routing table. Those routes will not be used to forward data packets to the intended destinations and the routes are not advertised by the routing protocol to neighbors and peers.

Routing policies control the size and content of the routing tables, the routes that are advertised, and the best route to take to reach a destination. Careful planning is essential to implement route policies that can affect the flow of routing information or packets in and traversing through the router. Before configuring and applying a route policy, develop an overall plan and strategy to accomplish your intended routing actions.

There are no default route policies. Each policy must be created explicitly and applied to a routing protocol or to the forwarding table. Policy parameters are modifiable.

6.1.1 Policy Statements

Route policies contain policy statements containing ordered entries containing match conditions and actions you specify. The entries should be sequenced from the most explicit to least explicit. Packet forwarding and routing can be implemented according to your defined policies. Policy-based routing allows you to dictate where traffic can be routed, through specific paths, or whether to forward or drop the traffic. Route policies can match a given route policy entry and continue searching for other matches within either the same route policy or the next route policy.

The process can stop when the first complete match is found and executes the action defined in the entry, either to accept or reject packets that match the criteria or proceed to the next entry or the next policy. You can specify matching criteria based on source, destination, or particular properties of a route. Route policies can be constructed to support multiple stages to the evaluation and setting various route attributes. You can also provide more matching conditions by specifying criteria such as:

- Autonomous system (AS) path policy options — A combination of AS numbers and regular expression operators.
- Community list — A group sharing a common property.
- Prefix list — A named list of prefixes.
- To and From criteria — A route's source and destination.

6.1.1.1 Policy Statement Chaining and Logical Expressions

Multiple policy-statement names can be specified in the CLI commands that attach route policies to specific functions. A chain of routing policies is created if a list of two or more policy names is specified in the CLI command. Route policy chaining allows complex route processing logic to be broken into smaller components. This enables the reuse of common functions and facilitates the process of amending and updating route control logic as required.

Each route is evaluated against a policy chain as follows.

- The route is evaluated against the first listed policy. If the route matches an entry with **action next-policy**, or the route matches no entry and the **default-action** is **next-policy** (or there is no **default-action** specified at all), the evaluation continues into the second policy.
- The route is evaluated against the second listed policy. If the route matches an entry with **action next-policy** or a **default-action** of **next-policy** applies, the evaluation continues to the third policy.
- The sequential evaluation of policies continues until the route is accepted or rejected by an entry or explicit **default-action**.

In addition to policy chaining, the SR OS also supports policy logical expressions that enable applications to use complex multiple policy statements. A policy logical expression can be used as a standalone expression or as part of a policy chain. Each policy chain supports a maximum of one logical expression. The logical expression is usually the first element of the policy chain; however, it can appear in a noninitial position as long as its length does not exceed 64 characters.

A route policy logical expression is a string composed of logical operators (keywords AND, OR and NOT), up to 16 route policy names (each up to 64 characters in length and delimited by brackets () and square brackets [] to group sub-expressions (with up to three levels of nesting)). The total length of the route policy expression cannot exceed 900 characters.

The following are examples of valid logical expressions in the SR OS CLI syntax:

- “NOT [policyA]”
- “[policyA] AND [policyB] OR [policyC]”
- “NOT ([policyA] OR [policyB] OR [policyC])”

The final result of a route policy evaluation against a logical expression is TRUE or FALSE. The SR OS rules for policy evaluation are as follows.

- If the expression includes a NOT operator followed by a sub-expression in brackets, the expression is flattened using De Morgan's rule.
For example, the expression “NOT ([policyA] OR [policyB] OR [policyC])” is flattened to: “(NOT [policyA]) AND (NOT [policyB]) AND (NOT [policyC])”.
- The usual rules of precedence apply: NOT is the highest priority, then AND, then OR.
- The use of “NOT <expression>” negates the TRUE/FALSE result of the expression.
- Where the logic is “(<expression1>) AND (<expression2>)”, if *expression1* is FALSE, the result is FALSE and *expression2* is not evaluated. If *expression1* is TRUE, *expression2* is then evaluated. The final result is TRUE only if both expressions are TRUE.
- Where the logic is “(<expression1>) OR (<expression2>)”, if *expression1* is TRUE, the result is TRUE and *expression2* is not evaluated. If *expression1* is FALSE, *expression2* is then evaluated. The final result is TRUE if either expression is TRUE.
- If the evaluation of a policy terminates with an **action accept**, **action next-entry** or **action next-policy**, the result is TRUE. If the evaluation of a policy terminates with an **action reject** or **action drop**, then the result is FALSE. If a route matches no entry in a policy and there is no specified **default-action**, the implied default action is **next-policy**; if there is no next policy, this translates to the default action for the protocol.

- When a route is evaluated against a policy contained in a logical expression, the route property changes (such as MED, local preference, communities) made by the matching entry or default action apply cumulatively to the route. The result of a cumulative change is that a policy evaluated later in the logical expression (or later in the entire policy chain) may undo or reverse prior changes. A later policy in the logical expression (or policy chain) may also match a route on the basis of route properties that were modified by earlier policies.

When evaluation of the logical expression is complete, the final TRUE or FALSE result is translated back to a traditional action. The FALSE value is translated to action reject; the TRUE value is translated to action accept, action next-policy or action next-entry to match the action of the last policy that produced the TRUE result.

6.1.1.2 Routing Policy Subroutines

It is possible to reference a routing policy from within another routing policy to construct powerful subroutine based policies.

Up to the three levels of subroutine calls are supported. Policy subroutines produce a final result of TRUE or FALSE through matching and policy entry actions. A policy entry action of 'accept' will evaluate to TRUE, and a policy entry action of 'reject' will evaluate to FALSE.

When using next-policy action state in the subroutine, the match value is defined by the default action behavior. The action is protocol-dependent. See [Default Action Behavior](#) for information about the default actions that are applied during packet processing.



Note: When subroutines are configured to reject routes, the accept action state can be used as a possible configuration in the subroutine match criteria to return a true-match, and the reject action state can be applied in the main policy entry that has called the subroutine.

If a match is not found during the evaluation of one or more routing policies, the final evaluation will return the accept or the reject provided by the default behavior based on the policy type (import/export) and the destination and/or source protocol.

6.1.1.3 Policy Evaluation Command

Operators can evaluate a routing policy against a BGP neighbor, routing context, or individual prefix before applying the policy to the neighbor or routing context. This command will display prefixes that are rejected by a policy and what modifications are made by a policy.

6.1.1.4 Exclusive Editing for Policy Configuration

Operators can set an exclusive lock on policy edit sessions. When the exclusive flag is set by an operator that is editing policy, other users (console or SNMP) are restricted from being able to begin, edit, commit, or abort policy. An administrative override is made available to reset the exclusive flag in the event of a session failure.

6.1.1.5 Default Action Behavior

The default action specifies how packets are to be processed when a policy related to the route is not explicitly configured. The following default actions are applied in the event that:

- A route policy does not specify a matching condition, all the routes being compared with the route policy are considered to be matches.
- A packet does not match any policy entries, then the next policy is evaluated. If a match does not occur then the last entry in the last policy is evaluated.
- If no default action is specified, the default behavior of the protocol controls whether the routes match or not.

If a default action is defined for one or more of the configured route policies, then the default action is handled as follows:

- The default action can be set to all available action states including accept, reject, next-entry, and next-policy.
- If the action states accept or reject, then the policy evaluation terminates and the appropriate result is returned.
- If a default action is defined and no matches occurred with the entries in the policy, then the default action is used.
- If a default action is defined and one or more matches occurred with the entries of the policy, then the default action is not used.

6.1.1.6 Denied IP Unicast Prefixes

The Route Table Manager does not inherently restrict any IP prefixes from the forwarding table, but individual routing protocols may not accept prefixes that are not globally routable.

6.1.1.7 Controlling Route Flapping

Route damping is a controlled acceptance of unstable routes from BGP peers so that any ripple effect caused by route flapping across BGP AS border routers is minimized. The motive is to delay the use of unstable routes (flapping routes) to forward data and advertisements until the route stabilizes.

Nokia's implementation of route damping is based on the following parameters:

- **Figure of Merit** — A route is assigned a Figure of Merit (FoM), which is proportional to the frequency of flaps. FoM should be able to characterize a route's behavior over a period of time.
- **Route flap** — A route flap is not limited to the withdrawn route. It also applies to any change in the AS path or the next hop of a reachable route. A change in AS path or next hop indicates that the intermediate AS or the route-advertising peer is not suppressing flapping routes at the source or during the propagation. Even if the route is accepted as a stable route, the data packets destined to the route could experience unstable routing due to the unstable AS path or next hop.
- **Suppress threshold** — The threshold is a configured value that, when exceeded, the route is suppressed and not advertised to other peers. The state is considered to be down from the perspective of the routing protocol.
- **Reuse threshold** — When FoM value falls below a configured reuse threshold and the route is still reachable, the route is advertised to other peers. The FoM value decays exponentially after a route is suppressed. This requires the BGP implementation to decay thousands of routes from a misbehaving peer.

The two events that could trigger the route flapping algorithm are:

- **Route flapping** — If a route flap is detected within a configured maximum route flap history time, the route's FoM is initialized and the route is marked as a potentially unstable route. Every time a route flaps, the FoM is increased and the route is suppressed if the FoM crosses the suppress threshold.
- **Route reuse timer trigger** — A suppressed route's FoM decays exponentially. When it crosses the reuse threshold, the route is eligible for advertisement if it is still reachable.

If the route continues to flap, the FoM, with respect to time scale, looks like a sawtooth waveform with the exponential rise and decay of FoM. To control flapping, the following parameters can be configured:

- **half-life** — The half-life value is the time, expressed in minutes, required for a route to remain stable in order for one half of the FoM value to be reduced. For example, if the half-life value is 6 (minutes) and the route remains stable for 6 minutes, then the new FoM value is 3. After another 6 minutes passes and the route remains stable, the new FoM value is 1.5.
- **max-suppress** — The maximum suppression time, expressed in minutes, is the maximum amount of time that a route can remain suppressed.
- **suppress** — If the FoM value exceeds the configured integer value, the route is suppressed for use or inclusion in advertisements.
- **reuse** — If the suppress value falls below the configured **reuse** value, then the route can be reused.

6.1.2 Regular Expressions

The ability to perform a filter match on confederations in the AS path and/or communities is supported. This filter allows customers to configure match criteria for specific confederation sets and sequences within the AS path so that they can be filtered out before cluttering the service provider's routing information base (RIB). When matching communities, the filter allows customers to configure match criteria within the community value.

SR OS uses regular expression strings to specify match criteria for:

- an AS path string; for example, "100 200 300"
- a community string; for example, "100:200" where 100 is the AS number, and 200 is the community-value
- any AS path beginning with a confederation SET or SEQ containing 65001 and 65002 only: for example "< 65001 65002 >.*"
- any AS path containing a confederation SET or SEQ, regardless of the contents: for example, ".* <.*> .*"

A regular expression is expressed in the form of terms and operators.

A term for an AS path regular expression must meet the following criteria:

- Regular expressions should always be enclosed in quotes.
- An elementary term; for example, an AS number "200"

- A range term composed of two elementary terms separated by the “-” character, such as “200-300”.
- The “.” dot wild-card character which matches any elementary term.
- A regular expression enclosed in parenthesis “()”.
- A regular expression enclosed in square brackets used to specify a set of choices of elementary or range terms; for example, [100-300 400] matches any AS number between 100 and 300 or the AS number 400.

A term for a community string regular expression is a string that is evaluated character by character and is composed of:

- an elementary term, which for a community string, is any single digit, such as “4”.
- a range term composed of two elementary terms separated by the ‘-’ character like “2-3”.
- a colon ':' to delimit the AS number from the community value
- the '.' dot wild-card character which matches any elementary term or ':'.
- a regular expression enclosed in parenthesis “()”.
- a regular expression enclosed in square brackets used to specify a set of choices of elementary or range terms; for example, [51-37] matches digit 5 or any single digit between 1 and 3 or the digit 7.
- extended communities of type route-target and route-origin may embed two regular expressions separated by ONE ampersand (“&”) character. The first expression is applied to the **as-value** of the community string and the second to the local administrative value.
- large communities may embed three regular expressions separated by two ampersand (“&”) characters. The first expression is applied to the Global Administrator part of the large community, the second expression is applied to the Local Data part 1 element of the large community, and the third expression is applied to the Local Data part 2 element of the large community.

A raw hex format can be keyed to represent all extended communities. For example, "ext:0102:dc0000020032" is the same as "target:65530:20". Hex values “ext:” and “&” can also be used to filter extended communities. The first expression is applied to the type/subtype of the extended community and the second expression to the value. In this case, hex values can also be used in the operands; for example, the value {3,f} matches a minimum 3 and a maximum 15 repetitions of the term.

The regular expression OPERATORS are listed in [Table 65](#).

Table 65 Regular Expression Operators

Operator	Description
	Matches the term on alternate sides of the pipe.
*	Matches multiple occurrences of the term.
?	Matches 0 or 1 occurrence of the term.
+	Matches 1 or more occurrence of the term.
()	Used to parenthesize so a regular expression is considered as one term.
[]	Used to demarcate a set of elementary or range terms.
-	Used between the start and end of a range.
{m,n}	Matches least m and at most n repetitions of the term.
{m}	Matches exactly m repetitions of the term.
{m,}	Matches m or more repetitions of the term.
^	Matches the beginning of the string - only allowed for communities.
\$	Matches the end of the string - only allowed for communities.
\	An escape character to indicate that the following character is a match criteria and not a grouping delimiter.
<>	Matches any AS path numbers containing a confederation SET or SEQ.
&	Matches “.” between terms of a community string (applicable to extended communities origin, target, bandwidth, ext only).

Examples of “target:”, “origin:” and “ext:” community strings are listed in [Table 66](#).

Table 66 Community Strings Examples

Example Expression	Example Matches
"ext:...&f(.*)[af]\$"	Matches the community "ext:0002:ffa0000001a".
'target:1&22	Matches community target:100:221.
target:^1&^22	Matches community target:100:221.
target:(.*)0\$&(.*)1\$	Matches community target:100:221.
origin:^1&(.*)1\$	Matches community origin:100:221.

Examples of AS path and community string regular expressions are listed in [Table 67](#).

Table 67 AS Path and Community Regular Expression Examples

AS Path to Match Criteria	Regular Expression	Example Matches
Null AS path	null ¹	Null AS path
AS path is 11	11	11
AS path is 11 22 33	11 22 33	11 22 33
Zero or more occurrences of AS number 11	11*	Null AS path 11 11 11 11 11 11 11 ... 11
Path of any length that begins with AS numbers 11, 22, 33	11 22 33 .*	11 22 33 11 22 33 400 500 600
Path of any length that ends with AS numbers 44, 55, 66	.* 44 55 66	44 55 66 100 44 55 66 100 200 44 55 66 100 200 300 44 55 66 100 200 300 ... 44 55 66
One occurrence of the AS numbers 100 and 200, followed by one or more occurrences of the number 33	100 200 33+	100 200 33 100 200 33 33 100 200 33 33 33 100 200 33 33 33 ... 33
One or more occurrences of AS number 11, followed by one or more occurrences of AS number 22, followed by one or more occurrences of AS number 33	11+ 22+ 33+	11 22 33 11 11 22 33 11 11 22 22 33 11 11 22 22 33 33 11 ... 11 22 ... 22 33 ...33
Path whose second AS number must be 11 or 22	(. 11) (. 22) .* or .(11 22) .*	100 11 200 22 300 400 ...
Path of length one or two whose second AS number might be 11 or 22	.(11 22)?	100 200 11 300 22

Table 67 AS Path and Community Regular Expression Examples (Continued)

AS Path to Match Criteria	Regular Expression	Example Matches
Path whose first AS number is 100 and second AS number is either 11 or 22	100 (11 22) . *	100 11 100 22 200 300
Either AS path 11, 22, or 33	[11 22 33]	11 22 33
Range of AS numbers to match a single AS number	10-14	10 or 11 or 12 or 13 or 14
	[10-12]*	Null AS path 10 or 11 or 12 10 10 or 10 11 or 10 12 11 10 or 11 11 or 11 12 12 10 or 12 11 or 12 12 ...
Zero or one occurrence of AS number 11	11? or 11{0,1}	Null AS path 11
One through four occurrences of AS number 11	11{1,4}	11 11 11 11 11 11 11 11 11 11
One through four occurrences of AS number 11 followed by one occurrence of AS number 22	11{1,4} 22	11 22 11 11 22 11 11 11 22 11 11 11 11 22
Path of any length, except nonexistent, whose second AS number can be anything, including nonexistent	..* or ..{0,}	100 100 200 11 22 33 44 55
AS number is 100. Community value is 200.	^100:200\$	100:200
AS number is 11 or 22. Community value is any number.	^((11) (22)):(.*)\$	11:100 22:100 11:200 ...
AS number is 11. Community value is any number that starts with 1.	^11:(1.*)\$	11:1 11:100 11:1100 ...

Table 67 AS Path and Community Regular Expression Examples (Continued)

AS Path to Match Criteria	Regular Expression	Example Matches
AS number is any number. Community value is any number that ends with 1, 2, or 3.	<code>^(.*):(.?[1-3])\$</code>	11:1 100:2002 333:55553 ...
AS number is 11 or 22. Community value is any number that starts with 3 and ends with 4, 5 or 9.	<code>^((11) (22)):(3.*[459])\$</code>	11:34 22:3335 11:3777779 ...
AS number is 11 or 22. Community value ends in 33 or 44.	<code>[^((11 22)):(.*((33) (44))))\$</code>	11:33 22:99944 22:555533 ...
Range of Community values	<code>100&^([1-9][0-9] [1-9][0-9][0-9] 1[0-9][0-9][0-9] 2000)\$</code>	100:10 100:11 100: ... 100:2000


Note:

1. The **null** keyword matches an empty AS path.

6.1.3 BGP and OSPF Route Policy Support

OSPF and BGP requires route policy support. [Figure 39](#) and [Figure 40](#) display where route policies are evaluated in the protocol. [Figure 39](#) depicts BGP which applies a route policy as an internal part of the BGP route selection process. [Figure 40](#) depicts OSPF which applies routing policies at the edge of the protocol, to control only the routes that are announced to or accepted from the Route Table Manager (RTM).

Figure 39 BGP Route Policy Diagram




0968

6.1.3.1 BGP Route Policies

Nokia's implementation of BGP uses route policies extensively. The implied or default route policies can be overridden by customized route policies. The default BGP properties, with no route policies configured, behave as follows:

- Accept all BGP routes into the RTM for consideration.
- Announce all used BGP learned routes to other BGP peers
- Announce none of the IGP, static or local routes to BGP peers.

Figure 40 OSPF Route Policy Diagram

6.1.3.2 Re-advertised Route Policies

Occasionally, BGP routes may be readvertised from BGP into OSPF, IS-IS, and RIP. OSPF export policies (policies control which routes are exported to OSPF) are not handled by the main OSPF task but are handled by a separate task or an RTM task that filters the routes before they are presented to the main OSPF task.

6.1.3.3 Triggered Policies

With triggered policy enabled, deletion and re-addition of a peer after making changes to export policy causes the new updates sent out to all peers

Triggered policy is not honored if a new peer added to BGP. Update with the old policy is sent to the newly added peer. New policy does not get applied to the new peer until the peer is flapped.

With triggered policy enabled, if a new BGP/static route comes in, new addition or modification of an export policy causes the updates to send out dynamically to all peers with the new/modified export policy

When multiple peers, say P1, P2 and P3 share the same export policy, any modifications to export policy followed by clear soft on one of the peer P1, will send out routes to P1 only according to newly modified policy.

Though routes with newly modified policy are not sent to other peers (P2, and P3) as no clear soft issues on these peers, RIB-OUT will show that new routes with modified policy are sent to all the peers. RIB-IN on peers P2 and P3 are shown correctly.



Note: With the triggered policy enabled, deletion and re-addition of a peer after making changes to export policy causes the new updates sent out to all peers.

The triggered policy is not honored if a new peer is added to BGP. An update with the old policy is sent to the newly added peer. The new policy is not applied to the new peer until the peer is flapped.

With the triggered policy enabled, if a new BGP/static route comes in, the new addition or modification of an export policy causes the updates to be sent out dynamically to all peers with the new/modified export policy.

When multiple peers, such as P1, P2 and P3, share the same export policy, any modifications to the export policy followed by **clear soft** on one of the peer P1s, will send out routes to P1 only according to the newly modified policy. Though routes with the newly modified policy are not sent to other peers (P2 and P3) as there are no clear soft issues on these peers, RIB-OUT will show that the new routes with the modified policy are sent to all the peers. RIB-IN on peers P2 and P3 are shown correctly.

6.1.3.4 Set MED to IGP Cost using Route Policies

This feature sets MED to the IGP cost of a route exported into BGP as an action in route policies. The **med-out** command in the **bgp**, **group**, and **neighbor** configuration context supports this option, but this method lacks per-prefix granularity. The enhanced **metric** command supported as a route policy action supports setting MED to a fixed number, or adding, or subtracting a fixed number from the received MED, and sets IGP cost option. The enhanced **metric {set {igp | number1} | {add | subtract} number2}** command is under **config>router>policy-options>policy-statement>entry>action**.

The **metric set igp** command, when used in a BGP export policy, have the same effect as the current **med-out igp** command, except that it applies only to the routes matched by the policy entry.

The effect of the **metric set igp** command depends on the route type and policy type as summarized in [Table 68](#).

Table 68 **Metric Set IGP Effect**

BGP Policy Type	Matched Route Type	Set Metric IGP Effect
Export	Non-BGP route (static, OSPF, ISIS, etc.)	Add MED attribute. Set value to M.
Export	BGP route w/o MED	Add MED attribute. Set value to D.
Export	BGP route with MED (value A)	Overwrite MED attribute with value D.

6.1.3.5 BGP Policy Subroutines

A BGP policy can call another policy (subroutine) and that subroutine can call another subroutine, and so on. This facilitates re-usability of common logic and a structured approach to writing BGP policies. Up to three levels of subroutine are supported.

6.1.3.6 Route Policies for BGP Next-Hop Resolution and Peer Tracking


This feature adds the flexibility to attach a route policy to the BGP next-hop resolution process; it also allows a route policy to be associated with the optional BGP peer-tracking function. BGP next-hop resolution is a fundamental part of BGP protocol operation; it determines the best matching route (or tunnel) for the BGP next-hop address and uses information about this resolving route in the best path selection algorithm and to program the forwarding table. Attaching a policy to BGP next-hop resolution provides more control over which IP routes in the routing table can become resolving routes. Similar flexibility is also available for BGP peer-tracking, which is an optional feature that allows the session with a BGP neighbor to be taken down if there is no IP route to the neighbor address or if the best matching IP route is rejected by the policy.

6.1.4 Routing Policy Parameterization

Routing policy parameterization allows operators a powerful and flexible configuration approach to routing policies for policies are often reused across BGP peers of a common type (transit; peer; customer; etc).

In current modes of operation as shown in [Figure 41](#), an operator must create individual routing policies, prefix-lists, AS-Path lists, community lists, etc for each peer despite many times the policy ultimately being the same. In this case, should an operator with 100 peers with a common policy behavior but unique policies have to make a change to entry 135 in the policy, they must do it on all policies – a significant amount of work that can result in incorrect/inconsistent policy behavior.

Figure 41 Route Policy Past Mode of Operation



Using a parameter based system allows an operator to have a single policy that is consistent across all peers of a type, while retaining the flexibility to reference different policy functions (prefixes, prefix-lists, community lists, etc) with unique names if required, by defining variables and the variable value. This feature will be able to inherit some of these variables directly from the routing process (for example, “@peerAS@” would be the configured BGP peer-AS).

Additionally, rather than policies being fixed and requiring many statements, the use of parameters and variables may be passed to simplify policy configuration. This reduces the number of policies required on a peering edge router with large numbers of peers where only small amounts of configuration changes between peers, such as the ASN and prefix-list name.

Two types of policy variables are supported:


- Global policy variables are configured using the **global-variables** command and can be referenced from any policy. Use global policy variables when the variables are used by multiple policies.

- Local policy variables are configured using the **policy-variables** command in each policy, and can be referenced in a sub-policy. Use local policy variables when the variables are specific to a policy and are not used in any other policies.



Note: Local policy variables have precedence over global policy variables, if policy variables have the same name.

Figure 42 Route Policy Parameterization using sub-policies



Variables expansion can use two formats, with the variable name delimited by at-signs (@) at the beginning and the end.

- Standard variables must start and end with at-signs (@), for example: @variable@
- Midstring variables must be enclosed with at-signs (@) and may be midstring, for example: @variable@, start@variable@end, @variable@end, start@variable@

Table 69 lists route policy variables supported in policy parameters.

Table 69 Route Policy Variable Support in Policy Parameters

Parameter Name	Variable in Policy "from" Statement	Variable in Policy "action" Statement	Variable Format	Standard Format Release	Midstring Format Release
aigp-metric	No	Yes	Standard	13.0.R4	N/A

Table 69 **Route Policy Variable Support in Policy Parameters (Continued)**

Parameter Name	Variable in Policy “from” Statement	Variable in Policy “action” Statement	Variable Format	Standard Format Release	Midstring Format Release
as-path	Yes	Yes	Midstring	12.0.R1	13.0.R1
as-path expression	Yes	No	Midstring	12.0.R4	13.0.R1
as-path-group	Yes	No	Midstring	12.0.R1	13.0.R1
as-path-group expression	Yes	No	Midstring	12.0.R4	13.0.R1
as-path-length	Yes	N/A	Standard	15.0.R1	N/A
as-path-prepend	N/A	Yes	Standard	13.0.R4	N/A
cluster-id	No	N/A	N/A	N/A	N/A
community	Yes	Yes	Midstring	12.0.R1	13.0.R1
community-count	Yes	N/A	Standard	15.0.R1	N/A
damping	No	Yes	Midstring	13.0.R4	13.0.R4
local-preference	Yes (15.0.R1)	Yes (13.0.R4)	Standard	13.0.R4	N/A
metric	Yes (15.0.R1)	Yes (13.0.R4)	Standard	13.0.R4	N/A
next-hop	No	Yes	Standard	13.0.R4	N/A
origin	No	Yes	Standard	13.0.R4	N/A
path-type	No	N/A	N/A	N/A	N/A
preference	No	Yes	Standard	13.0.R4	N/A
prefix-list	Yes	No	Midstring	12.0.R1	13.0.R1
tag	No	Yes	Standard	13.0.R4	N/A
type	No	Yes	Standard	13.0.R4	N/A



Note: Midstring variables are only supported in the object name. Standard variables are supported in policy parameters inside as-path or as-path-group expression objects; for example: as-path “as” expression “. {65001,@variable@}65022”.

For the definition of the variables, there are three different possible types:

- **name** *name-string* **value** *value-string*
- **name** *name-string* **address** *ip-address*
- **name** *name-string* **number** *value-number*

Depending on the parameter referenced, specify the correct type as follows:

- value-string: **as-path**, **as-path-group**, **community**, **prefix-list**, **damping**
- ip-address: **next-hop**
- value-number: **aigp-metric**, **as-path-length**, **as-path-prepend**, **community-count**, **local-preference**, **metric**, **origin**, **origin-validation**, **preference**, **tag**, **type**

The logical flow of this is to configure a per-peer policy in which the variable names and values are defined. Using [Figure 42](#) as the example, the following configuration would be applied:

```
prefix-list "pfx-as63535"
    prefix 6.3.5.0/24 through 32
exit
prefix-list "pfx-as64535"
    prefix 6.4.5.0/24 through 32
exit
prefix-list "pfx-as65535"
    prefix 6.5.5.0/24 through 32
exit

community "cm-as63535" members "7750:63535"
community "cm-as65535" members "7750:65535"
community "cm-as64535-rejects" members "64535:14"
community "cm-as65535-rejects" members "65535:14"
community "cm-as64535-highpref" members "7777:64535"

as-path "as63535" expression "^63535$"
as-path "as64535" expression "^64535$"
as-path "as65535" expression "^65535$"

policy-statement "peer1"
    entry 5
        from
            policy-variables
                name "@cm-reject@" value "cm-as65535-rejects"
            exit
        policy "std-peering-inbound-drop"
    exit
    action reject
    entry 10
        from
            policy-variables
                name "@localcm@" value "cm-as65535"
                name "@peer-as@" value "as65535"
                name "@cm-highpref@" value "cm-as65535-highpref"
                name "@peer-prefix@" value "pfx-as65535"
            exit
        policy "std-peering-inbound-main"
```

```

        exit
        action accept
        exit
    exit
exit

policy-statement "peer2"
    entry 5
        from
            policy-variables
                name "@cm-reject@" value "cm-as64535-rejects"
            exit
            policy "std-peering-inbound-drop"
        exit
        action reject
    entry 10
        from
            policy-variables
                name "@localcm@" value "cm-as64535"
                name "@peer-as@" value "as64535"
                name "@cm-highpref@" value "cm-as64535-highpref"
                name "@peer-prefix@" value "pfx-as64535"
            exit
            policy "std-peering-inbound-main"
        exit
        action accept
        exit
    exit
exit

policy-statement "peer3"
    entry 5
        from
            policy-variables
                name "@cm-reject@" value "cm-as63535-rejects"
            exit
            policy "std-peering-inbound-drop"
        exit
        action reject
    entry 10
        from
            policy-variables
                name "@localcm@" value "cm-as63535"
                name "@peer-as@" value "as63535"
                name "@cm-highpref@" value "cm-as63535-highpref"
                name "@peer-prefix@" value "pfx-as63535"
            exit
            policy "std-peering-inbound-main"
        exit
        action accept
        exit
    exit
exit

policy-statement "std-peering-inbound-drop"
    default-action reject
    entry 10
        from
            community "@cm-reject@"

```

```
        exit
        action accept

policy-statement "std-peering-inbound-main"
  default-action reject
  entry 10
    from
      prefix-list "@peer-prefix@"
      as-path "@peer-as@"
    exit
    action accept
      community add "@localcm@"
      local-preference 400
    exit
  exit
  entry 20
    from
      community "@cm-highpref@"
    exit
    action accept
      community add "@localcm@"
      local-preference 4000
    exit
  exit
exit
```

This configuration would take slightly different actions depending on the peer.

Peer 1

- Routes that have a community matching “cm-as65535-rejects” are dropped.
- Routes matching the prefix list “pfx-as65535” that originated in the peer AS=65535 are accepted with a local preference of 400.
- Community “7750:65535” is added to accepted prefixes.
- As community-list “cm-65535-highpref” does not exist, no routes are modified with a local preference of 4000.

Peer 2

- Routes that have a community matching “cm-as64535-rejects” are dropped.
- Routes matching the prefix list “pfx-as65535” that originated in the peer AS=65535 are accepted with a local preference of 400.
- Prefixes matching “cm-as64535-highpref” are set to a local-preference of 4000.

Peer 3

- As community-list “cm-as63535-rejects” does not exist, no routes are dropped by the first entry.
- Routes matching the prefix list “pfx-as63535” that originated in the peer AS=63535 are accepted with a local preference of 400.

- Community “7750:63535” is added to accepted prefixes.
- As community-list “cm-63535-highpref” does not exist, no routes are modified with a local preference of 4000.

6.1.5 When to Use Route Policies


The following are examples of circumstances of when to configure and apply unique route policies.

- When you want to control the protocol to allow all routes to be imported into the routing table. This enables the routing table to learn about particular routes to enable packet forwarding and redistributing packets into other routing protocols.
- When you want to control the exporting of a protocol’s learned active routes.
- When you want a routing protocol to announce active routes learned from another routing protocol, which is sometimes called *route redistribution*.
- When you want unique behaviors to control route characteristics. For example, change the route preference.
- When you want unique behaviors to control route characteristics. For example, change the route preference, AS path, or community values to manipulate the control the route selection.
- When you want to control BGP route flapping (damping).

6.2 Route Policy Configuration Process Overview

Figure 43 displays the process to provision basic route policy parameters.

Figure 43 Route Policy Configuration and Implementation Flow



Router_Config_37

6.3 Configuration Notes

This section describes route policy configuration caveats.

6.3.1 General

- When configuring policy statements, the policy statement name must be unique.

6.3.2 Policy Reference Checks

The policy reference checks functionality can be enabled with **config>router>policy-reference-checks** to indicate policies and policy objects in policy statements that are referenced but do not exist. An INFO or WARNING message is displayed in the CLI.



Note: The policy reference checks feature is available only in the classic CLI and not via SNMP or model-driven interfaces. Model-driven interfaces and **mixed** management interface configuration mode use strict policy reference checks starting in Release 16.0.R3, where the checks cannot be disabled.

If a policy is referenced which has not yet been configured, and **no policy-reference-checks** is set, the configuration line succeeds with the CLI message:

```
WARNING: CLI Policy "bar" does not exist
```

If a policy is referenced which has not yet been configured, and **policy-reference-checks** is set, the configuration line errors and fails with the CLI message:

```
MINOR: CLI Policy "foo" does not exist
```

Policy reference checks are required in router configuration contexts that reference a policy name, including the following:

- **configure router bgp export**
- **configure router bgp group export**
- **configure router bgp group import**

- **configure router bgp group neighbor export**
- **configure router bgp group neighbor import**
- **configure router bgp import**
- **configure service vprn bgp export**
- **configure service vprn bgp group export**
- **configure service vprn bgp group import**
- **configure service vprn bgp group neighbor export**
- **configure service vprn bgp group neighbor import**
- **configure service vprn bgp import**
- **configure service vprn vrf-export**
- **configure service vprn vrf-import**

Policy reference checks are required in policy configuration contexts that reference a policy element, including the following:

- **entry from flow-spec-dest** *prefix-list-name*
- **entry from flow-spec-source** *prefix-list-name*
- **entry from group-address** *prefix-list-name*
- **entry from host-ip** *prefix-list-name*
- **entry from neighbor** {*ip-address* | *prefix-list-name*}
- **entry from next-hop** *prefix-list-name*
- **entry from prefix-list** *name* [*name* ... (up to 28 max)]
- **entry from source-address** **prefix-list** *prefix-list-name*
- **entry to neighbor** {*ip-address* | *prefix-list-name*}
- **entry to prefix-list** *name* [*name* ... (up to 28 max)]
- **as-path**
- **as-path-group**
- **community**
- **damping**
- policy statements used in sub-policies

6.3.2.1 Known Limitations

Policy elements must be entered in the correct order to resolve references. For example, if a prefix-list is referenced in a policy statement, the prefix-list must be entered first. This feature does not order policy statements so that references to policy elements exist.

Variable expansion is not supported with the policy reference checks feature. Variables that reference policy elements that do not exist are not checked and are permitted.

This feature is available only for configuring policies. If a policy is deleted, the command succeeds with no info or warning messages displayed. Information from the **show router policy-edits** CLI command can be used to see warnings about which elements do not exist.

The policy reference checks functionality is not supported with policy expressions.

6.4 Configuring Route Policies with CLI

This section provides information to configure route policies using the command line interface.

6.4.1 Route Policy Configuration Overview

Route policies allow you to configure routing according to specifically defined policies. You can create policies and entries to allow or deny paths based on various parameters such as destination address, protocol, packet size, and community list.

Policies can be as simple or complex as required. A simple policy can block routes for a specific location or IP address. More complex policies can be configured using numerous policy statement entries containing matching conditions to specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

6.4.1.1 When to Create Routing Policies

Route policies are created in the **config>router** context. There are no default route policies. Each route policy must be explicitly created and applied. Applying route policies can introduce more efficiency as well as more complexity to routers.

A route policy impacts the flow of routing information or packets within and through the router. A routing policy can be specified to prevent a particular customer's routes to be placed in the route table which causes those routes to not forward traffic to various destinations and the routes are not advertised by the routing protocol to neighbors.

Route policies can be created to control:

- A protocol to export all the active routes learned by that protocol.
- Route characteristics to control which route is selected to act as the active route to reach a destination and advertise the route to neighbors.
- Protocol to import all routes into the routing table. A routing table must learn about particular routes to be able to forward packets and redistribute to other routing protocols.
- Damping.

Before a route policy is applied, analyze the policy's purpose and be aware of the results (and consequences) when packets match the specified criteria and the associated actions and default actions, if specified, are executed. Membership reports can be filtered based on a specific source address.

6.4.1.2 Default Route Policy Actions

Each routing protocol has default behaviors for the import and export of routing information. [Table 70](#) shows the default behavior for each routing protocol.

Table 70 Default Route Policy Actions

Protocol	Import	Export
OSPF	Not applicable. All OSPF routes are accepted from OSPF neighbors and cannot be controlled via route policies.	• Internal routes: All OSPF routes are automatically advertised to all neighbors. • External routes: By default all non-OSPF learned routes are not advertised to OSPF neighbors
IS-IS	Not applicable. All IS-IS routes are accepted from IS-IS neighbors and can not be controlled via route policies	• Internal routes: All IS-IS routes are automatically advertised to all neighbors. • External routes: By default all non-IS-IS learned routes are not advertised to IS-IS peers.
RIP	By default, all RIP-learned routes are accepted.	• External routes: By default all non-RIP learned routes are not advertised to RIP peers.
BGP	By default, all routes from BGP peers are accepted and passed to the BGP route selection process.	• Internal routes: By default all active BGP routes are advertised to BGP peers • External routes: By default all non-BGP learned routes are not advertised to BGP peers.

6.4.1.3 Policy Evaluation

Routing policy statements can consist of as few as one or several entries. The entries specify the matching criteria. A route is compared to the first entry in the policy statement. If it matches, the specified entry action is taken, either accepted or rejected. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends.

If the route does not match the first entry, the route is compared to the next entry (if more than one is configured) in the policy statement. If there is a match with the second entry, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends, and so on.


Each route policy statement can have a default-action clause defined. If a default-action is defined for one or more of the configured route policies, then the default actions should be handled in the following ways:

- The process stops when the first complete match is found and executes the action defined in the entry.
- If the packet does not match any of the entries, the system executes the default action specified in the policy statement.

Figure 44 depicts an example of the route policy process.


Route policies can also match a given route policy entry and continue to search for other entries within either the same route policy or the next route policy by specifying the *next-entry* or *next-policy* option in the entry's **action** command. Policies can be constructed to support multiple states to the evaluation and setting of various route attributes.

Figure 45 depicts the next-policy and next-entry route processes.

Figure 44 **Route Policy Process Example**

Note:

1. See [Table 70](#)

Figure 45 Next Policy Logic Example


6.4.1.4 Damping

Damping initiates controls when routes flap. Route flapping can occur when an advertised route between nodes alternates (flaps) back and forth between two paths due to network problems which cause intermittent route failures. It is necessary to reduce the amount of routing state change updates propagated in order to limit processing requirements. Thus, when a route flaps beyond a configured value (the suppress value), then that route is removed from the routing tables and routing protocols until the value falls below the reuse value.

A route can be suppressed according to the Figure of Merit (FoM) value. The FoM is a value that is added to a route each time it flaps. A new route begins with an FoM value of 0.


Damping is optional. If damping is configured, the following parameter values must be explicitly specified as there are no default values:

- [suppress](#)
- [half-life](#)
- [reuse](#)

- **max-suppress**

When a route's FoM value exceeds the suppress value, then the route is removed from the routing table. The route is considered to be stable when the FoM drops below the reuse value by means of the specified half-life parameter. The route is returned to the routing tables. When routes have higher FoM and half-life values, they are suppressed for longer periods of time. [Figure 46](#) depicts an example of a flapping route, the suppress threshold, the half-life decay (time), and reuse threshold. The peaks represent route flaps, the slopes represent half-life decay.

Figure 46 Damping Example



6.4.2 Basic Configurations

This section provides information to configure route policies and configuration examples of common tasks. The minimal route policy parameters that need to be configured are:

- Policy statement with the following parameters specified:
 - At least one entry
 - Entry action

Following is a sample route policy configuration:

```
A:ALA-B>config>router>policy-options# info
-----
      community "all-types" members "5000:[1-6] [1-9] [0-9]"
      community "all-normal" members "5000:[1-5] [1-9] [0-9]"
. . .
as-path "Outside madeup paths" ".* 5001 .*"
as-path "Outside Internet paths" ".* 5002 .*"
policy-statement "RejectOutsideASPaths"
  entry 1
    from
```



```

        protocol bgpospf
        as-path "Outside madeup paths"
    exit
    action reject
    exit
exit
entry 2
    from
        protocol bgpospf
        as-path "Outside Internet paths"
    exit
    action reject
    exit
exit
entry 3
    from
        protocol ospf
    exit
    to
        protocol bgpospf
    exit
    action reject
    exit
exit
entry 4
    from
        protocol isis
    exit
    to
        protocol bgpospf
    exit
    action reject
    exit
exit
default-action accept
exit
exit
policy-statement "aggregate-customer-peer-only"
    entry 1
        from
            community "all-customer-announce"
        exit
        action accept
        exit
    exit
    default-action reject
    exit
exit
-----
A:ALA-B>config>router>policy-options#

```

6.4.3 Configuring Route Policy Components

This section describes configuring route policy components.

6.4.3.1 Beginning the Policy Statement

Use the following CLI syntax to begin a policy statement configuration. In order for a policy statement to be complete an entry must be specified (see [Configuring an Entry](#)).

CLI Syntax:

```
config>router>policy-options
    begin
        policy-statement name
        description text
```

The following error message displays when the you try to modify a policy options command without entering **begin** first.

```
A:ALA-B>config>router>policy-options# policy-statement "allow all"
MINOR: CLI The policy-
options must be in edit mode by calling begin before any changes can be made.
```

The following example displays policy statement configuration command usage. These commands are configured in the **config>router** context.

Example:

```
config>router# policy-options
policy-options# begin
```

There are no default policy statement options. All parameters must be explicitly configured.

6.4.3.2 Creating a Route Policy

To enter the mode to create or edit route policies, you must enter the **begin** keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

The following error message displays when the you try to modify a policy options command without entering **begin** first.

```
A:ALA-B>config>router>policy-options# policy-statement "allow all"
MINOR: CLI The policy-options must be in edit mode by calling begin before any
changes can be made.
```

```
A:ALA-B>config>router>policy-options# info
#-----
# Policy
```

```
#-----
        policy-options
        begin
            policy-statement "allow all"
description "General Policy"
        ...
            exit
exit
-----
A:ALA-B>config>router>policy-options#
```

6.4.3.3 Configuring a Default Action

Specifying a default action is optional. The default action controls those packets not matching any policy statement entries. If no default action is specified for the policy, then the action associated with the protocol to which the routing policy was applied is performed. The default action is applied only to those routes that do not match any policy entries.

A policy statement must include at least one entry (see [Configuring an Entry](#)).

To enter the mode to create or edit route policies, you must enter the **begin** keyword at the **config>router>policy-options** prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

The following example displays the default action configuration:

```
A:ALA-B>config>router>policy-options# info
-----
        policy-statement "1"
            default-action accept
                as-path add "test"
                community add "365"
                damping "flaptest"
                next-hop 10.10.10.104
            exit
        exit
-----
A:ALA-B>config>router>policy-options#
```

6.4.3.4 Configuring an Entry

An entry action must be specified. The other parameters in the **entry action** context are optional. Refer to the [Route Policy Command Reference](#) for the commands and syntax.

The following example displays entry parameters and includes the default action parameters which were displayed in the previous section.

```
A:ALA-B>config>router>policy-options# info
-----
      policy-statement "1"
        entry 1
          to
neighbor 10.10.10.104
          exit
          action accept
          exit
        exit
        entry 2
          from
            protocol ospf 1
          exit
          to
            protocol ospf
            neighbor 10.10.0.91
          exit
          action accept
          exit
        exit
        default-action accept
        . . .
      exit
    exit
```

The following example displays entry parameters and includes the default action parameters which were displayed in the previous section.

```
A:ALA-B>config>router>policy-options# info
-----
      policy-statement "1"
        entry 1
          to
            protocol bgp
            neighbor 10.10.10.104
          exit
          action accept
          exit
        exit
        entry 2
          from
            protocol ospf 1
          exit
          to
            protocol ospf
```

```

        neighbor 10.10.0.91
    exit
    action accept
    exit
exit
default-action accept
. . .
exit
exit

```

6.4.3.5 Configuring a Community List

Community lists are composed of a group of destinations which share a common property. Community lists allow you to administer actions on a configured group instead of having to execute identical commands for each member.

The following example displays a community list configuration:

```

A:ALA-B>config>router>policy-options# info
-----
community "eastern" members "100:200"
community "western" members "100:300"
community "northern" members "100:400"
community "southern" members "100:500"
community "headquarters" members "100:1000"
policy-statement "1"
    entry 1
        to
            protocol bgp
            neighbor 10.10.10.104
        exit
        action accept
    . . .
-----
A:ALA-B>config>router>policy-options#

```

6.4.3.6 Configuring Damping

The following considerations apply.

- For each damping profile, all parameters must be configured.
- The **suppress** value must be greater than the reuse value (see [Damping Example](#)).
- Damping can be enabled in the **config>router>bgp** context on the BGP global, group, and neighbor levels. If damping is enabled, but route policy does not specify a damping profile, the default damping profile will be used. This profile is always present and consists of the following parameters:

- half-life: 15 minutes
- max-suppress: 60 minutes
- suppress: 3000
- reuse: 750

The following example displays a damping configuration:

```
*A:cses-A13>config>router>policy-options# info
-----
      damping "damptest123"
        half-life 15
        max-suppress 60
        reuse 750
        suppress 1000
      exit
-----
*A:cses-A13>config>router>policy-options#
```

6.4.3.7 Configuring a Prefix List

The following example displays a prefix list configuration:

```
A:ALA-B>config>router>policy-options# info
-----
      prefix-list "western"
        prefix 10.10.0.1/32 exact
        prefix 10.10.0.2/32 exact
        prefix 10.10.0.3/32 exact
        prefix 10.10.0.4/32 exact
      exit
      damping "damptest123"
        half-life 15
        max-suppress 60
        reuse 750
      exit
-----
A:ALA-B>config>router>policy-options#
```

6.5 Route Policy Configuration Management Tasks

This section discusses route policy configuration management tasks.

6.5.1 Editing Policy Statements and Parameters

Route policy statements can be edited to modify, add, or delete parameters. To enter the mode to edit route policies, you must enter the `begin` keyword at the `config>router> policy-options` prompt. Other editing commands include:

- The **commit** command saves changes made to route policies during a session.
- The **abort** command discards changes that have been made to route policies during a session.

The following example displays a changed configuration:

```
A:ALA-B>config>router>policy-options>policy-statement# info
-----
description "Level 1"
entry 1
  to
    protocol bgp
    neighbor 10.10.10.104
  exit
  action accept
  exit
exit
entry 2
  from
    protocol ospf
  exit
  to
    protocol ospf
    neighbor 10.10.0.91
  exit
  action accept
  exit
exit
entry 4
  description "new entry"
  from
    protocol isis
    area 0.0.0.20
  exit
  action reject
exit
default-action accept
as-path add "test"
```

```
community add "365"  
damping "flapper"  
next-hop 10.10.10.104  
exit  
-----
```

6.5.2 Deleting an Entry

Use the following CLI syntax to delete a policy statement entry:

CLI Syntax:

```
config>router>policy-options  
begin  
commit  
abort  
policy-statement name  
no entry entry-id
```

The following example displays the commands required to delete a policy statement entry.

Example:

```
config>router>policy-options# begin  
policy-options# policy-statement "1"  
policy-options>policy-statement# no entry 4  
policy-options>policy-statement# commit
```

6.5.3 Deleting a Policy Statement

Use the following CLI syntax to delete a policy statement:

CLI Syntax:

```
config>router>policy-options  
begin  
commit  
abort  
no policy-statement name
```

The following example displays the commands required to delete a policy statement.

Example:

```
config>router>policy-options# begin  
policy-options# no policy-statement 1  
policy-options# commit
```


6.6 Route Policy Command Reference

6.6.1 Command Hierarchies

- [Route Policy Configuration Commands](#)
- [Segment Routing Commands](#)

6.6.1.1 Route Policy Configuration Commands

```

config
  — [no] router
    — [no] policy-reference-checks
    — [no] triggered-policy
    — [no] weighted-ecmp
    — [no] policy-options
      — abort
      — [no] as-path name
        — expression regular-expression
        — no expression
      — [no] as-path-group name
        — entry entry-id expression reg-exp
        — no entry entry-id
      — begin [exclusive]
      — commit
      — [no] community name
        — expression expression [exact]
        — no expression
        — members comm-id [comm-id]
        — no members comm-id [comm-id]
      — [no] damping name
        — half-life minutes
        — no half-life
        — max-suppress minutes
        — no max-suppress
        — reuse integer
        — no reuse
        — suppress integer
        — no suppress
      — [no] exclusive-lock-time seconds
      — [no] global-variables
        — name name-string value value-string
        — name name-string address ip-address
        — name name-string number value-number
        — no name name-string
      — [no] policy-statement name
        — default-action {accept | next-entry | next-policy | drop | reject}

```

-
- **no default-action**
 - **add-paths-send-limit** *send-limit*
 - **no add-paths-send-limit**
 - **admin-tag-policy** *policy-name*
 - **no admin-tag-policy**
 - **advertise-label** {*per-prefix* | *pop*}
 - **no advertise-label**
 - **aigp-metric** *metric*
 - **aigp-metric add**
 - **aigp-metric ipg**
 - **no aigp-metric**
 - **as-path** {*add* | *replace*} *name*
 - **no as-path**
 - **as-path-prepend** *as-path* [*repeat*]
 - **as-path-prepend most-recent** [*repeat*]
 - **no as-path-prepend**
 - **bgp-leak**
 - **no bgp-leak**
 - **community add** *name* [*name*]
 - **community remove** *name* [*name*]
 - **community replace** *name* [*name*]
 - **no community**
 - **create-udp-tunnel**
 - **no create-udp-tunnel**
 - **damping** {*name* | *none*}
 - **no damping**
 - **install-backup-path**
 - **no install-backup-path**
 - **local-preference** *local-preference*
 - **no local-preference**
 - **metric** {*add* | *subtract*} *metric*
 - **metric set** [*igp* | *metric-value*]
 - **no metric**
 - **multicast-redirection** [*fwd-service service-id*] *ip-int-name*
 - **no multicast-redirection**
 - **next-hop** *ip-address*
 - **no next-hop**
 - **[no] next-hop-self**
 - **origin** {*igp* | *egp* | *incomplete* | *param-name*}
 - **no origin**
 - **origin-validation-state** {*state* | *param-name*}
 - **no origin-validation-state**
 - **preference** *preference*
 - **no preference**
 - **[no] resolve-static**
 - **source-class** [*value*]
 - **no source-class**
 - **[no] sticky-ecmp**
 - **tag** *tag*
 - **no tag**
 - **type** {*type* | *param-name*}
 - **no type**
 - **description** *description-string*
 - **no description**

-
- [no] **entry** *entry-id*
 - **action** {accept | next-entry | next-policy | drop | reject}
 - **no action**
 - **add-paths-send-limit** *send-limit*
 - **no add-paths-send-limit**
 - **admin-tag-policy** *policy-name*
 - **advertise-label** {per-prefix | pop}
 - **no advertise-label**
 - **aigp-metric** *metric*
 - **aigp-metric** *metric* add
 - **aigp-metric** *ipg*
 - **no aigp-metric**
 - **as-path** {add | replace} *name*
 - **no as-path**
 - **as-path-prepend** *as-number* [*repeat*]
 - **as-path-prepend** **most-recent** [*repeat*]
 - **no as-path-prepend**
 - **bgp-leak**
 - **no bgp-leak**
 - **community** **add** *name* [*name*]
 - **community** **remove** *name* [*name*]
 - **community** **replace** *name* [*name*]
 - **no community**
 - **create-udp-tunnel**
 - **no create-udp-tunnel**
 - **damping** {*name* | none}
 - **no damping**
 - **fc** *fc* [priority {low | high}]
 - **no fc**
 - **install-backup-path**
 - **no install-backup-path**
 - **local-preference** *local-preference*
 - **no local-preference**
 - **metric** {add | subtract} *metric*
 - **metric** **set** {igp | *metric-value*}
 - **no metric**
 - **next-hop** *ip-address*
 - **multicast-redirection** [fwd-service *service-id*] *ip-int-name*
 - **no multicast-redirection**
 - **no next-hop**
 - [no] **next-hop-self**
 - **origin** {igp | egp | incomplete | *param-name*}
 - **no origin**
 - **origin-validation-state** {*state* | *param-name*}
 - **no origin-validation-state**
 - **policy** *name*
 - **preference** *preference*
 - **no preference**
 - [no] **resolve-static**
 - [no] **sticky-ecmp**
 - **tag** {no-tag | *tag*}
 - **no tag**
 - **type** *type*

-
- **no type**
 - **conditional-expression**
 - **route-exists** *expression*
 - **no route-exists**
 - **description** *description-string*
 - **no description**
 - **[no] from**
 - **area** *area-id*
 - **no area**
 - **as-path** *name*
 - **no as-path**
 - **as-path-group** *name*
 - **no as-path-group**
 - **as-path-length** *length* [equal | or-higher | or-lower] [unique]
 - **no as-path-length**
 - **cluster-id** *ip-address/mask* [*ip-address/mask*]
 - **cluster-id none**
 - **no cluster-id**
 - **community** *comm-name*
 - **community** *expression* *expression*
 - **no community**
 - **community-count** *count* [equal | or-higher | or-lower] [standard | extended | large]
 - **no community-count**
 - **evpn-type** *type*
 - **no evpn-type**
 - **[no] external**
 - **family** [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [ms-pw] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] [bgp-ls] [mcast-ipv6] [mcast-vpn-ipv6] [sr-policy-ipv4]
 - **no family**
 - **flow-spec-dest** *prefix-list-name*
 - **no flow-spec-dest**
 - **flow-spec-source** *prefix-list-name*
 - **no flow-spec-source**
 - **group-address** *prefix-list-name*
 - **no group-address**
 - **host-ip** *prefix-list-name*
 - **no host-ip**
 - **interface** *interface-name*
 - **no interface**
 - **level** {1 | 2}
 - **no level**
 - **local-preference** *preference* [equal | or-higher | or-lower]
 - **no local-preference**
 - **metric** *metric* [equal | or-higher | or-lower]
 - **no metric**
 - **mvpn-type** {1 | 2 | 3 | 4 | 5 | 6 | 7 }
 - **no mvpn-type**
 - **neighbor** {*ip-address* | **prefix-list** *name*}

```

— no neighbor
— next-hop ip-address
— next-hop prefix-list name
— no next-hop
— origin {igp | egp | incomplete | any | aaa | dynamic |
    static | bonding}
— no origin
— origin-validation-state state
— no origin-validation-state
— path-type {ibgp | ebgp}
— no path-type
— policy plcy-or-long-expr
— no policy
— policy-variables
    — name name-string value value-string
    — name name-string address ip-address
    — name name-string number value-number
    — no name name-string
— no policy name name-string
— prefix-list name [name]
— no prefix-list
— protocol protocol [all | instance instance]
— protocol bgp bgp-label
— no protocol
— source-address ip-address
— source-address prefix-list prefix-list-name
— no source-address
— state state
— no state
— tag {no-tag | tag}
— no tag
— type type
— no type
— [no] to
    — level {1 | 2}
    — no level
    — neighbor {ip-address | prefix-list name}
    — no neighbor
    — [no] prefix-list name [name]
    — protocol protocol [all | {instance instance}]
    — protocol bgp bgp-label
    — no protocol
— renumber from entry-id to entry-id

config
— [no] router
    — [no] policy-options
        — [no] prefix-list name
            — prefix ip-prefix/prefix-length [exact | longer | through length | prefix-
                length-range length1-length2]
            — no prefix [ipv-prefix/prefix-length] [exact | longer | through length |
                prefix-length-range length1-length2]

```

6.6.1.2 Segment Routing Commands

```

config
  — [no] router
    — segment-routing
      — sr-policies
        — [no] ingress-statistics
          — [no] shutdown
        — reserved-label-block name
        — no reserved-label-block
        — [no] shutdown
        — static-policy name [create]
        — no static-policy name
          — binding-sid number
          — no binding-sid
          — color color
          — no color
          — distinguisher id
          — no distinguisher
          — endpoint ipv4-address
          — no endpoint
          — head-end local
          — head-end ipv4-address
          — no head-end
          — preference preference
          — segment-list [1..32] [create]
          — no segment-list [1..32]
            — segment [1..11] [create]
            — no segment [1..11]
              — mpls-label value
              — no mpls-label
            — [no] shutdown
            — weight weight
            — no weight
          — [no] shutdown

```

6.6.2 Command Descriptions

6.6.2.1 Generic Commands

abort

Syntax abort

Context config>router>policy-options

Description This command is required to discard changes made to a route policy.

begin

Syntax **begin {exclusive}**

Context config>router>policy-options

Description This command is required in order to enter the mode to create or edit route policies.

Parameters **exclusive** — Specifies an exclusive lock on the policy configuration. Other CLI and SNMP users will be unable to edit the policy configuration until the lock is removed (via commit, abort, a timeout occurring, or a forced override).

commit

Syntax **commit**

Context config>router>policy-options

Description This command is required to save changes made to a route policy.

description

Syntax **description** *string*
no description

Context config>router>policy-options>policy-statement
config>router>policy-options>policy-statement>entry

Description This command creates a text description which is stored in the configuration file to help identify the content of the entity.

The **no** form of the command removes the string from the configuration.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

6.6.2.2 Segment Routing

segment-routing

Syntax	segment-routing
Context	config>router>
Description	This command creates a context to configure protocol-independent parameters relating to segment routing.

sr-policies

Syntax	sr-policies
Context	config>router>segment-routing
Description	This command creates the context to configure segment routing policies. A segment routing policy specifies traffic to be matched by the policy and actions to take on the matched traffic by applying the instructions encoded in one or more segment lists.

ingress-statistics

Syntax	[no] ingress-statistics
Context	config>router>segment-routing>sr-policies
Description	This command administratively enables the collection of ingress traffic statistics for all segment routing policies. The statistics provide counts for the number of incoming packets and bytes corresponding to each (color, endpoint) combination. If there are any SR-MPLS interfaces on an FP2 or older line-cards, an attempt to enable this command will fail. The no form of the command disables ingress stats collection for all segment routing policies.
Default	no ingress-statistics

shutdown

Syntax	[no] shutdown
Context	config>router>segment-routing>sr-policies>ingress-statistics

- Description** This command administratively disables the collection of ingress statistics for all segment routing policies.
- The **no** form of the command administratively enables the collection of ingress statistics for all segment routing policies.

reserved-label-block

- Syntax** **reserved-label-block** *name*
no reserved-label-block
- Context** config>router>segment-routing>sr-policies
- Description** This command associates a reserved label block with segment routing policies. The *name* must already exist. Reserved label blocks are configured under the **config>router>mpls-labels** hierarchy.
- A locally-targeted segment routing policy (statically configured or BGP signaled) cannot be activated if its binding SID (BSID) is not an available label between the start-label and end-label of the referenced reserved label block.
- The **no** form of the command removes any association of segment routing policies with a reserved label block.
- Default** no reserved-label-block
- Parameters** *name* — Specifies the name of a **reserved-label-block** that has already been configured, up to 64 characters.

shutdown

- Syntax** [**no**] **shutdown**
- Context** config>router>segment-routing>sr-policies
- Description** This command deactivates all segment routing policies and removes the associated entries from the forwarding plane of the router.
- It is necessary to execute this shutdown if you want to make a change to the reserved-label-block reference.
- The **no** form of the command enables all segment routing policies so that they can be revalidated and reinstalled as necessary.

static-policy

- Syntax** **static-policy** *name* [**create**]

no static-policy *name*

Context	config>router>segment-routing>sr-policies
Description	This command creates a context to configure a segment routing policy. The resulting segment routing policy is targeted for local installation or propagation by BGP to another router. The no form of the command deletes the statically defined segment routing policy.
Default	no static-policy
Parameters	<i>name</i> — Specifies the name assigned to the statically defined segment routing policy, up to 64 characters. create — Keyword used to create the policy.

binding-sid

Syntax	binding-sid <i>number</i> no binding-sid
Context	config>router>segment-routing>sr-policies>static-policy
Description	This command associates a binding SID with a statically defined segment routing policy. This is a mandatory parameter in order to enable the segment routing policy. The BSID label should be an available label in the reserved-label-block range. The no form of the command removes the BSID association.
Default	no binding-sid
Parameters	<i>number</i> — Specifies the binding SID label value. Values 32 to 1048575

color

Syntax	color <i>color</i> no color
Context	config>router>segment-routing>sr-policies>static-policy
Description	This command associates a color value with a statically defined segment routing policy. This is a mandatory parameter in order to enable the segment routing policy. The no form of the command removes the color association.
Default	no color

Parameters *color* — Specifies the color ID.
Values 0 to 4294967295

distinguisher

Syntax **distinguisher** *id*
 no distinguisher

Context config>router>segment-routing>sr-policies>static-policy

Description This command associates a distinguisher value with a statically defined segment routing policy. This is a mandatory parameter for non-local segment routing policies (for which head-end is not local). Every non-local segment routing policy must have a unique distinguisher value. When a non-local static segment routing policy is imported into BGP and originated as a BGP route, the configured distinguisher value is copied into the NLRI of the route.

 The **no** form of the command removes the distinguisher association.

Default no distinguisher

Parameters *id* — Specifies the distinguisher ID.
Values 1 to 4294967295

endpoint

Syntax **endpoint** *ipv4-address*
 no endpoint

Context config>router>segment-routing>sr-policies>static-policy

Description This command associates an IPv4 endpoint address with a statically-defined segment routing policy. This is a mandatory parameter in order to enable the segment routing policy.

 An endpoint address of 0.0.0.0 is a special value that matches all BGP next-hops, however in order to use this policy the BGP route must have a color extended community with the color-only bits set to '01' or '10'.

 The **no** form of the command removes the endpoint association.

Default no endpoint

Parameters *ipv4-address* — Specifies the endpoint address of the segment routing policy.
Values *a.b.c.d*

head-end

Syntax	head-end local head-end <i>ipv4-address</i> no head-end
Context	config>router>segment-routing>sr-policies>static-policy
Description	This command associates a head-end location with a statically-defined segment routing policy. The head-end is an identification of the router which is the target to install the policy. This is a mandatory parameter for enabling the segment routing policy. In order to associate a static policy with the local router as head-end, the keyword local must be specified. The static policy is associated with another (non-local) router if the head-end parameter is set to any IPv4 address. When a non-local static segment routing policy is imported into BGP and originated as a BGP route, the configured head-end address is converted into an IPv4-address specific route-target extended community that is automatically added to the route. The no form of the command removes the head-end association.
Default	no head-end
Parameters	local — Keyword indicating that the policy is intended to be used by the local router and not advertised to other BGP routers. <i>ipv4-address</i> — Specifies the head-end address of the (non-local) segment routing policy. Values <i>a.b.c.d</i>

preference

Syntax	preference <i>preference</i>
Context	config>router>segment-routing>sr-policies>static-policy
Description	This command associates a preference value with a statically defined-segment routing policy. This is an optional parameter. When there are multiple policies for the same (color, endpoint) combination that are targeted for local installation, only one is selected as the active path for the (color, endpoint). In this selection process (which considers both static local policies and BGP signaled policies), the policy with the highest preference value is preferred over all policies with a lower preference value. The no form of the command reverts to the default value.
Default	preference 100

Parameters	<i>preference</i> — Specifies the preference ID.
Values	0 to 4294967295

segment-list

Syntax	segment-list [1..32] [create] no segment-list <i>list</i>
Context	config>router>segment-routing>sr-policies
Description	This command creates the context to configure a segment list for the statically-defined segment routing policy. Up to 32 segment lists are supported per policy. The no form of the command deletes the segment list.
Parameters	create — Keyword used to create the segment list.

segment

Syntax	segment [1..11] [create] no segment
Context	config>router>segment-routing>sr-policies>static-policy>segment-list
Description	This command creates the context to configure a segment inside a segment-list of a statically-defined segment routing policy. Each segment list can have up to 11 segments. The no form of the command deletes the segment context.
Default	no segment
Parameters	create — Keyword used to create the list.

mpls-label

Syntax	mpls-label <i>value</i> no mpls-label
Context	config>router>segment-routing>sr-policies>static-policy>segment-list>segment
Description	This command configures the MPLS label value this is associated with a segment. The no form of the command removes the label value.

Default	no mpls-label
Parameters	<i>value</i> — Specifies the MPLS label value.
Values	0 to 1048575

shutdown

Syntax	[no] shutdown
Context	config>router>segment-routing>sr-policies>static-policy>segment-list
Description	This command deactivates a segment-list. If this is done on an active policy with more than one segment list, then traffic forwarded by the policy will be diverted to the remaining segment-lists. The no form of the command enables the segment list so that it can be validated and installed as necessary.
Default	shutdown

weight

Syntax	weight <i>weight</i> no weight
Context	config>router>segment-routing>sr-policies>static-policy>segment-list
Description	This command associates a weight value with a segment list of a statically-defined segment routing policy in order to achieve weighted ECMP behavior. Weight is an optional parameter. When any segment-list in the active policy has a weight greater than 1, traffic matching the policy will be load-balanced across the segment lists according to their relative weight values. The no form of the command reverts to the default value.
Default	weight 1
Parameters	<i>weight</i> — Specifies the weight value.
Values	1 to 4294967295

shutdown

Syntax	[no] shutdown
Context	config>router>segment-routing>sr-policies>static-policy

Description	This command deactivates the associated static policy and causes another policy for the same (color, endpoint) combination to be promoted as the active path, assuming there is another valid policy. It is necessary to execute this shutdown if you want to make critical configuration changes to the static policy. The no form of the command enables the static policy so that it can be validated and installed as necessary.

6.6.2.3 Route Policy Options

as-path

Syntax	[no] as-path <i>name</i>
Context	config>router>policy-options
Description	This command creates a route policy AS path to use in route policy entries. The no form of the command deletes the AS path.
Default	no as-path
Parameters	<i>name</i> — The AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

expression

Syntax	expression <i>regular-expression</i> no expression
Context	config>router>policy-options>as-path
Description	This command creates a route policy AS path regular expression statement to use in route policy entries. The no form of the command deletes the AS path regular expression statement.
Parameters	<i>reg-exp</i> — The AS path regular expression. Allowed values are any string up to 256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

null — The AS path expressed as an empty regular expression string.

as-path-group

Syntax	[no] as-path-group <i>name</i>
Context	config>router>policy-options
Description	This command creates a route policy AS path regular expression statement to use in route policy entries. The no form of the command deletes the AS path regular expression statement.
Default	no as-path-group
Parameters	<i>name</i> — The AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

entry

Syntax	entry <i>entry-id</i> <i>expression</i> <i>reg-exp</i> no entry
Context	config>router>policy-options>as-path-group
Description	This command creates the context to edit route policy entries within an autonomous system path group. Multiple entries can be created using unique entries. The router exits the filter when the first match is found and executes the action specified. For this reason, entries must be sequenced correctly from most to least explicit. An entry does not require matching criteria defined (in which case, everything matches) but must have at least define an action in order to be considered complete. Entries without an action are considered incomplete and will be rendered inactive. The no form of the command removes the specified entry from the autonomous system path group.
Parameters	<i>entry-id</i> — The entry ID expressed as a decimal integer. An <i>entry-id</i> uniquely identifies match criteria and the corresponding action. It is recommended that multiple entries be given <i>entry-ids</i> in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.
Values	1 to 4294967295

reg-exp — The AS path group regular expression. Allowed values are any string up to 256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

community

Syntax	[no] community <i>name</i>
Context	config>router>policy-options
Description	This command creates a route policy community to use in route policy entries. The no form of the command deletes the community list or the provided community ID.
Default	no community
Parameters	<i>name</i> — The community list name. Allowed values are any string up to 64 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on). The entire string must be enclosed within double quotes.

expression

Syntax	expression <i>expression</i> [exact] no expression
Context	config>router>policy-options>community
Description	This command creates a logical expression to match a route policy community. The no form of the command deletes the logical expression.
Default	no expression
Parameters	expression <i>expression</i> — Specifies a logical expression containing terms and operators. It can contain sub-expressions enclosed in round brackets. Values up to 900 characters <expression> is one of the following: <expression> {AND OR} <expression> [NOT] (<expression>) [NOT] <comm-id> For example: from community expression "[community list A] OR ([community list B] AND [community list C])" exact — All the communities indicated by the expression must be present in the route in order for a match to occur.

members

Syntax	[no] members <i>comm-id</i> [<i>comm-id</i>]
Context	config>router>policy-options>community
Description	This command add members to a route policy community list to use in route policy entries. The no version of this command deletes route policy community members.
Parameters	<i>comm-id</i> — Specifies a BGP community value, up to 72 characters. A community ID can be specified in different forms. Values [<i>as-num:comm-val</i> <i>reg-ex</i> <i>ext-comm</i> <i>well-known-comm</i> <i>large-comm</i>] where: • <i>as-num</i> — 0 to 65535 • <i>comm-val</i> — 0 to 65535 • <i>reg-ex</i> — A regular expression string. Allowed values are any string up to 72 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (such as “#”, “\$”, or spaces), the entire string must be enclosed within double quotes. • <i>ext-comm</i> — The extended community, defined as one of the following: – {<i>target</i> <i>origin</i>}:<i>ip-address:comm-val</i> – {<i>target</i> <i>origin</i>}:<i>reg-ex1&reg-ex2</i> – {<i>target</i> <i>origin</i>}:<i>ip-address:reg-ex2</i> – {<i>target</i> <i>origin</i>}:<i>asnum:ext-comm-val</i> – {<i>target</i> <i>origin</i>}:<i>ext-asnum:comm-val</i> – bandwidth:<i>asnum:val-in-mbps</i> – ext:4300:<i>ovstate</i> – ext:value1:value2 – flowspec-set:<i>ext-asnum:group-id</i> – flowspec-set-trans:<i>ext-asnum:group-id</i> – color:<i>co-bits:color-value</i> where: – <i>target</i> — route target – <i>origin</i> — route origin – <i>ip-address</i> — a.b.c.d – <i>ext-comm-val</i> — 0 to 4294967295 – <i>ext-asnum</i> — 0 to 4294967295

- *reg-ex1*, *reg-ex2* — A regular expression string. Allowed values are any string up to 63 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
 - **bandwidth** — bandwidth
 - *val-in-mbps* — 0 to 16777215
 - **ext:4300** — origin verification
 - *ovstate* — 0, 1, or 2 (0 for valid, 1 for not found, 2 for invalid)
 - **ext** — extended
 - *value1* — 0000 to FFFF
 - *value2* — 0 to FFFFFFFFFF
 - **flowspec-set** — flow-spec set
 - **flowspec-set-trans** — flow-spec set transitive
 - *group-id* — 0 to 16383
 - *co-bits* — 00, 01, 10 or 11
 - *color-value* — 0 to 4294967295
 - *well-known-comm* — **null**, **no-export**, **no-export-subconfed**, **no-advertise**, **llgr-stale**, **no-llgr**, **blackhole**
 - *large-comm* — large community, defined as one of the following:
 - *ext-asnum:ext-comm-val:ext-comm-val*
 - *reg-ex3&reg-ex3&reg-ex3*
- where:
- *reg-ex3* — A regular expression string. Allowed values are any string up to 68 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

exclusive-lock-time

Syntax	exclusive-lock-time <i>seconds</i> no exclusive-lock
Context	config>router>policy-options
Description	This command specifies the inactivity timer for the exclusive lock time for policy editing. When a session is idle for greater than this time, the lock is removed and the configuration changes is aborted.
Default	exclusive-lock-time 300

Parameters	<i>seconds</i> — Specifies the duration the session with exclusive lock may be inactive.
Values	Values: 1 to 3600

global-variables

Syntax	global-variables no global-variables
Context	config>router>policy-options
Description	This command enables the global-variables configuration context. The no form of the command removes all global variables.

name

Syntax	name <i>name-string</i> value <i>value-string</i> name <i>name-string</i> address <i>ip-address</i> name <i>name-string</i> number <i>value-number</i> no name <i>name-string</i>
Context	config>router>policy-options>global-variables config>router>policy-options>policy-statement>from>policy-variables
Description	This command allows operators to configure routing policies that are often reused across BGP peers of a common type (transit, peer, customer, and so on). Using global variables allows an operator to have a single variable that is consistent across all peers of a type, while retaining the flexibility to reference different policy functions (prefixes, prefix-lists, community lists, and so on) with unique names. Depending on the parameter referenced, the correct type should be specified as follows: • value-string: as-path, as-path-group, community, prefix-list, damping • ip-address: next-hop • value-number: aigp-metric, as-path-prepend, local-preference, metric, origin, origin-validation, preference, tag, type The no form of the command removes the global variable.
Parameters	<i>name-string</i> — The name of the global variable, with the variable delimited by at-signs (@) at the beginning and the end of the name. Allowed values are any string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

value-string — The value of the policy variable. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

value-number — Specifies the numerical value of the policy variable.

Values 0 to 4294967295

ip-address — Specifies the IP address of the policy variable.

Values ipv4-address:

- a.b.c.d

ipv6-address:

- x:x:x:x:x:x:x [-interface]
- x:x:x:x:x:d.d.d.d [-interface]
- x: [0 to FFFF]H
- d: [0 to 255]D

policy-options

Syntax	[no] policy-options
Context	config>router
Description	This command enables the context to configure route policies. Route policies are applied to the routing protocol. The no form of the command deletes the route policy configuration.

policy-reference-checks

Syntax	[no] policy-reference-checks
Context	config>router
Description	This command checks policy references to ensure that a policy exists and displays a CLI error if the policy does not exist. Enabling this option protects against accidentally referencing a missing or misspelled policy, that can lead to unexpected results when the policy is evaluated. The no version of this command disables policy reference checks and allows policies that do not exist to be referenced.
Default	no policy-reference-checks

triggered-policy

Syntax	[no] triggered-policy
Context	config>router
Description	This command triggers route policy re-evaluation.

By default, when a change is made to a policy in the **config router policy options** context and then committed, the change is effective immediately. There may be circumstances when the changes should or must be delayed; for example, if a policy change is implemented that would affect every BGP peer on a router, the consequences could be dramatic. It is more effective to control changes on a peer by peer basis.

If the **triggered-policy** command is enabled, and a given peer is established, and you want the peer to remain up, then, in order for a change to a route policy to take effect, a **clear** command with the *soft* or *soft-inbound* option must be used. In other words, when a **triggered-policy** is enabled, any routine policy change or policy assignment change within the protocol will not take effect until the protocol is reset or a clear command is issued to re-evaluate route policies; for example, **clear router bgp neighbor x.x.x.x soft**. This keeps the peer up and the change made to a route policy is applied only to that peer, or group of peers.

Default	no triggered-policy

weighted-ecmp

Syntax	[no] weighted-ecmp
Context	config>router
Description	This command enables weighted load balancing in the base router instance for certain types of OSPF, IS-IS, and static routes with equal-cost multipath (ECMP) next hops.

For OSPF and static routes, this command only applies to IPv4 routes where all the next hops are tunnel next hops corresponding to MPLS LSPs with configured load-balancing weights. Weighted load balancing over MPLS LSPs is supported in the following cases:

- an IGP prefix resolved to IGP shortcuts in the RTM (**igp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance)
- a BGP prefix with the BGP next hop resolved to IGP shortcuts in RTM (**igp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance)
- a static route prefix resolved to an indirect next hop, which is resolved to a set of equal-metric MPLS LSPs in the TTM. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set.
- a static route prefix resolved to an indirect next-hop, which is resolved to IGP shortcuts in the RTM
- a BGP prefix with a BGP next hop resolved to a static route, which resolves to a set of tunnel next hops towards an indirect next hop in the RTM or TTM

- a BGP prefix resolved to another BGP prefix whose next hop is resolved to a set of ECMP tunnel next hops with a static route in the RTM or TTM or to IGP shortcuts in the RTM

For IS-IS routes, in addition to enabling the behavior described for OSPF and static routes, this command also allows weighted load balancing when all the ECMP next hops are interfaces with configured load-balancing weights. The interface-level weighted ECMP support for IS-IS applies to both IPv4 and IPv6.

If one or more LSPs or interfaces in the ECMP set of a prefix do not have a load-balancing weight configured, the regular ECMP spraying for the prefix will be performed.

The **no** form of the command restores regular ECMP spraying of packets to static and IGP route destinations.

Default no weighted-ecmp

6.6.2.4 Route Policy Damping Commands

damping

Syntax	[no] damping <i>name</i>
Context	config>router>policy-options
Description	This command creates a context to configure a route damping profile to use in route policy entries. The no form of the command deletes the named route damping profile.
Default	no damping
Parameters	<i>name</i> — Specifies the damping profile name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

half-life

Syntax	half-life <i>minutes</i> no half-life
Context	config>router>policy-options>damping
Description	This command configures the half-life parameter for the route damping profile.

The half life value is the time, expressed in minutes, required for a route to remain stable in order for the Figure of Merit (FoM) value to be reduced by one half; for example, if the half life value is 6 (minutes) and the route remains stable for 6 minutes, then the new FoM value is 3 (minutes). After another 3 minutes pass and the route remains stable, the new FoM value is 1.5 (minutes).

When the FoM value falls below the [reuse](#) threshold, the route is once again considered valid and can be reused or included in route advertisements.

The **no** form of the command removes the half life parameter from the damping profile.

Default	no half-life
Parameters	<i>minutes</i> — Specifies the half life in minutes expressed as a decimal integer.
Values	1 to 45

max-suppress

Syntax	max-suppress <i>minutes</i> no max-suppress
Context	config>router>policy-options>damping
Description	This command configures the maximum suppression parameter for the route damping profile. This value indicates the maximum time, expressed in minutes, that a route can remain suppressed. The no form of the command removes the maximum suppression parameter from the damping profile.
Default	no max-suppress
Parameters	<i>minutes</i> — Specifies the maximum suppression time, in minutes, expressed as a decimal integer.
Values	1 to 720

reuse

Syntax	reuse <i>integer</i> no reuse
Context	config>router>policy-options>damping
Description	This command configures the reuse parameter for the route damping profile.

When the Figure of Merit (FoM) value falls below the **reuse** threshold, the route is once again considered valid and can be reused or included in route advertisements.

The **no** form of the command removes the reuse parameter from the damping profile.

Default	no reuse
Parameters	<i>integer</i> — Specifies the reuse value expressed as a decimal integer.
Values	1 to 20000

suppress

Syntax	suppress <i>integer</i> no suppress
Context	config>router>policy-options>damping
Description	This command configures the suppression parameter for the route policy damping profile. A route is suppressed when it has flapped frequently enough to increase the Figure of Merit (FoM) value to exceed the suppress threshold limit. When the FoM value exceeds the suppress threshold limit, the route is removed from the route table or inclusion in advertisements. The no form of the command removes the suppress parameter from the damping profile.
Default	no suppress
Parameters	<i>integer</i> — Specifies the suppress value expressed as a decimal integer.
Values	1 to 20000

6.6.2.5 Route Policy Prefix Commands

prefix-list

Syntax	[no] prefix-list <i>name</i>
Context	config>router>policy-options
Description	This command creates a context to configure a prefix list to use in route policy entries. The no form of the command deletes the named prefix list.

Parameters *name* — Specifies the prefix list name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".

An empty prefix list can be configured for pre-provisioning. This empty prefix list will not find a match when referred to by a policy. When removing member prefixes from a prefix list, the prefix list will not be automatically removed when the last member is removed. If required, an empty prefix list must be explicitly removed using the **no** form of the command.

prefix

Syntax **[no] prefix** *ip-prefix/prefix-length* { **[exact | longer | through** *length* **]** | **[prefix-length-range** *length1-length2* **]** }

no prefix [*ipv-prefix/prefix-length*] **[exact | longer | through** *length* **| prefix-length-range** *length1-length2* **]**

Context config>router>policy-options>prefix-list

Description This command creates a prefix entry in the route policy prefix list.

The **no** form of the command deletes the prefix entry from the prefix list.

Parameters *ip-prefix* — Specifies the IP prefix for prefix list entry in dotted decimal notation.

Values ipv4-prefix:

- a.b.c.d (host bits must be 0)

ipv4-prefix-length: [0 to 32]

ipv6-prefix:

- x:x:x:x:x:x:x (eight 16-bit pieces)
- x:x:x:x:x:d.d.d.d
- x: [0 to FFFF]H
- d: [0 to 255]D

ipv6-prefix-length: [0 to 128]

exact — Specifies the prefix list entry only matches the route with the specified *ip-prefix* and prefix *mask* (length) values.

longer — Specifies the prefix list entry matches any route that matches the specified *ip-prefix* and prefix *mask* length values equal to or greater than the specified mask.

through *length* — Specifies the prefix list entry matches any route that matches the specified ip-prefix and has a prefix length between the specified *length* values inclusive.

Values 0 to 32

prefix-length-range *length1 - length2* — Specifies a route must match the most significant bits and have a prefix length with the given range. The range is inclusive of start and end values.

Values 0 to 32, *length2* > *length1*

6.6.2.6 Route Policy Entry Match Commands

entry

Syntax	entry <i>entry-id</i> no entry
Context	config>router>policy-options>policy-statement
Description	This command creates the context to edit route policy entries within the route policy statement. Multiple entries can be created using unique entries. The router exits the filter when the first match is found and executes the action specified. For this reason, entries must be sequenced correctly from most to least explicit. An entry does not require matching criteria defined (in which case, everything matches) but must have at least define an action in order to be considered complete. Entries without an action are considered incomplete and will be rendered inactive. The no form of the command removes the specified entry from the route policy statement.
Parameters	<i>entry-id</i> — Specifies the entry ID expressed as a decimal integer. An <i>entry-id</i> uniquely identifies match criteria and the corresponding action. It is recommended that multiple entries be given <i>entry-ids</i> in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries. Values 1 to 4294967295

from

Syntax	[no] from
Context	config>router>policy-options>policy-statement>entry
Description	This command creates the context to configure policy match criteria based on a route's source or the protocol from which the route is received. If no condition is specified, all route sources are considered to match.

The **no** form of the command deletes the source match criteria for the route policy statement entry.

area

Syntax	area <i>area-id</i> no area
Context	config>router>policy-options>policy-statement>entry>from
Description	This command configures an OSPF area as a route policy match criterion. This match criterion is only used in export policies. All OSPF routes (internal and external) are matched using this criterion if the best path for the route is by the specified area. The no form of the command removes the OSPF area match criterion.
Default	no area
Parameters	<i>area-id</i> — Specifies the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer. Values 0.0.0.0 to 255.255.255.255 (dotted decimal), 0 to 4294967295 (decimal)

as-path

Syntax	as-path <i>name</i> no as-path
Context	config>router>policy-options>policy-statement>entry>from
Description	This command configures an AS path regular expression statement as a match criterion for the route policy entry. If no AS path criterion is specified, any AS path is considered to match. AS path regular expression statements are configured at the global route policy level (config>router>policy-options>as-path <i>name</i>). The no form of the command removes the AS path regular expression statement as a match criterion.
Default	no as-path

Parameters *name* — Specifies the AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".

as-path-group

Syntax **as-path-group** *name*
no as-path-group *name*

Context config>router>policy-options>policy-statement>entry>from

Description This command creates a route policy AS path regular expression statement to use in route policy entries.

 The **no** form of the command deletes the AS path regular expression statement.

Default no as-path-group

Parameters *name* — Specifies the AS path regular expression name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".

as-path-length

Syntax **as-path-length** *length* [equal | or-higher | or-lower] [unique]
no as-path-length

Context config>router>policy-options>policy-statement>entry>from

Description This command matches BGP routes based on their AS path length (the number of AS numbers in the AS_PATH).

 If no comparison qualifiers are present (**equal**, **or-higher**, **or-lower**), then **equal** is the implied default.

 Confederation member AS numbers in the AS_PATH do not count towards the total. An AS_SET element is considered to have a length of 1.

 The **unique** option counts .

 A non-BGP route does not match a policy entry if it contains the **as-path-length** command.

Default	no as-path-length
Parameters	<i>length</i> — Specifies the length of the AS path. Values 0 to 255, or a parameter name delimited by starting and ending at-sign (@) characters equal — Specifies that matched routes should have the same number of AS path elements as the value specified. or-higher — Specifies that matched routes should have the same or a greater number of AS path elements as the value specified. or-lower — Specifies that matched routes should have the same or a lower number of AS path elements as the value specified. unique — Specifies that only the unique AS numbers should be counted (that is, multiple occurrences of the same AS number in the sequence count as one).

cluster-id

Syntax	cluster-id <i>ip-address/mask</i> [<i>ip-address/mask</i>] cluster-id none no cluster-id
Context	config>router>policy-options>policy-statement>entry>from
Description	This command enables BGP routes to be matched based on the IP addresses encoded in the CLUSTER_LIST attribute. The first <i>ip-address/mask pair</i> is matched against the most recently added cluster ID. Each subsequent <i>ip-address/mask pair</i> is tested against the next most recent cluster ID. For example, to match all routes reflected by the RR with cluster ID 1.1.1.1 and then any other RR before reaching the router where the policy is applied, use the command cluster-id 0.0.0.0/0 1.1.1.1/32.



Note: The command matches routes with two or more cluster IDs; the third and older cluster IDs are not evaluated and are automatically considered matching.

The **cluster-id none** form of the command only matches BGP routes without any CLUSTER_LIST attribute.

A non-BGP route does not match a policy entry if it contains the **cluster-id** command.

Default	no cluster-id
Parameters	<i>ip-address</i> — Specifies the 32-bit cluster ID in dotted decimal notation. Values a.b.c.d

mask — Specifies a bit mask to apply to the *ip-address* parameter.

Values 0 to 32 (0 is only allowed if the *ip-address* is 0.0.0.0)

none — Specifies that only BGP routes without a CLUSTER_LIST attribute should be matched.

community-count

Syntax	community-count <i>count</i> [equal or-higher or-lower] [standard extended large] no community-count
Context	config>router>policy-options>policy-statement>entry>from
Description	This command matches BGP routes based on community length (that is, the number of community members in the COMMUNITY, EXTENDED_COMMUNITY, or LARGE_COMMUNITY the attributes). If no comparison qualifiers are present (equal, or-higher, or-lower), then equal is the implied default. Without the optional standard, extended, or large keyword, the community length applies to the total number of communities, of all types. If some keywords are present, then only the types specified are counted against the limit. A non-BGP route does not match a policy entry if it contains the community-count command.
Default	no community-count
Parameters	<i>count</i> — Specifies the number of community members. Values 0 to 1024, or a parameter, up to 32 characters, name delimited by a starting and ending at-sign (@) character equal — Specifies that matched routes should have the same number of AS path elements as the value specified. or-higher — Specifies that matched routes should have the same or a greater number of community members as the value specified. or-lower — Specifies that matched routes should have the same or a lower number of community members as the value specified. standard — Specifies that only communities in the COMMUNITY attribute should be counted. extended — Specifies that only communities in the EXTENDED_COMMUNITY attribute should be counted. large — Specifies that only communities in the LARGE_COMMUNITY attribute should be counted.

local-preference

Syntax	local-preference <i>preference</i> [equal or-higher or-lower] no local-preference
Context	config>router>policy-options>policy-statement>entry>from
Description	This command matches BGP routes based on local preference (the value in the LOCAL_PREF attribute). If no comparison qualifiers are present (equal, or-higher, or-lower), then equal is the implied default. A non-BGP route does not match a policy entry if it contains the local-preference command.
Default	no local-preference
Parameters	<i>preference</i> — Specifies the local preference value. Values 0 to 4294967295, or a parameter name delimited by starting and ending at-sign (@) characters equal — Specifies that matched routes should have the same local preference as the value specified. or-higher — Specifies that matched routes should have the same or a greater local preference as the value specified. or-lower — Specifies that matched routes should have the same or a lower local preference as the value specified.

metric

Syntax	metric <i>metric</i> [equal or-higher or-lower] no metric
Context	config>router>policy-options>policy-statement>entry>from
Description	This command matches BGP routes based on local preference (the value in the MULTI_EXIT_DISC attribute). If no comparison qualifiers are present (equal, or-higher, or-lower), then equal is the implied default. A non-BGP route does not match a policy entry if it contains the metric command. In addition, a BGP route without a MED attribute also does not match a policy entry if it contains a metric command.
Default	no metric

Parameters	<i>metric</i> — Specifies the MED value. Values 0 to 4294967295, or a parameter name delimited by starting and ending at-sign (@) characters equal — Specifies that matched routes should have the same MED as the value specified. or-higher — Specifies that matched routes should have the same or a greater MED as the value specified. or-lower — Specifies that matched routes should have the same or a lower MED as the value specified.

next-hop

Syntax	next-hop <i>ip-address</i> next-hop prefix-list <i>name</i> no next-hop
Context	config>router>policy-options>policy-statement>entry>from
Description	This command enables BGP routes to be matched based on the BGP next-hop address. The match condition is evaluated against the IPv4 or IPv6 address in the NEXT_HOP or MP_REACH_NLRI attribute. When the next-hop match is applied to VPN-IP routes, the Route Distinguisher (RD) is ignored. A non-BGP route does not match a policy entry if it contains the next-hop command.
Default	no next-hop
Parameters	<i>ip-address</i> — An IPv4 or IPv6 address. Values a.b.c.d or x:x:x:x:x:x:x or x:x:x:x:x:d.d.d.d <i>name</i> — Specifies the name of a prefix-list (up to 32 characters). prefix-list — Specifies that the BGP next hop should be matched against a prefix-list instead of an individual IP address.

path-type

Syntax	path-type {ibgp ebgp} no path-type
Context	config>router>policy-options>policy-statement>entry>from

Description	This command matches BGP routes based on their path type (EBGP or IBGP). A route learned from an EBGP peer has path-type ebgp. A route learned from an IBGP or confed-EBGP peer has path-type ibgp. A non-BGP route does not match a policy entry if it contains the path-type command.
Default	no path-type
Parameters	<i>ip-address</i> — An IPv4 or IPv6 address. Values a.b.c.d or x:x:x:x:x:x:x or x:x:x:x:x:d.d.d.d ibgp — Matches routes from internal BGP peers. ebgp — Matches routes from external BGP peers.

external

Syntax	[no] external
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies the external route matching criteria for the entry.
Default	no external

family

Syntax	family [ipv4] [label-ipv4] [vpn-ipv4] [ipv6] [label-ipv6] [vpn-ipv6] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [ms-pw] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] [bgp-ls] [mcast-ipv6] [mcast-vpn-ipv6] [sr-policy-ipv4] no family
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies address families as matching conditions.
Default	no family
Parameters	ipv4 — Matches routes belonging to the IPv4 unicast (unlabeled) address family. label-ipv4 — Matches routes belonging to the IPv4 unicast (labeled) address family. vpn-ipv4 — Matches routes belonging to the IPv4 VPN (SAFI 128) address family. ipv6 — Matches routes belonging to the IPv6 unicast (unlabeled) address family. label-ipv6 — Matches routes belonging to the IPv6 unicast (labeled) address family. vpn-ipv6 — Matches routes belonging to the IPv6 VPN (SAFI 128) address family. mcast-ipv4 — Matches routes belonging to the IPv4 multicast SAFI address family.

l2-vpn — Matches routes belonging to the L2 VPN address family.

mvpn-ipv4 — Matches routes belonging to the IPv4 multicast VPN address family.

mvpn-ipv6 — Matches routes belonging to the IPv6 multicast VPN address family.

mdt-safi — Matches routes belonging to the MDT SAFI address family.

ms-pw — Matches routes belonging to the multi-segment pseudowire address family.

flow-ipv4 — Matches routes belonging to the IPv4 flowspec address family.

flow-ipv6 — Matches routes belonging to the IPv6 flowspec address family.

route-target — Matches routes belonging to the address family for RT constrain routes.

mcast-vpn-ipv4 — Matches routes belonging to the IPv4 VPN multicast (SAFI 129) address family.

evpn — Matches routes belonging to the EVPN address family.

bgp-ls — Enables the advertisement of BGP-LS address family to the associated BGP neighbors.

mcast-ipv6 — Matches routes belonging to the IPv6 multicast SAFI address family.

mcast-vpn-ipv6 — Matches routes belonging to the IPv6 multicast routes from a VPRN over the provider network. This family is only applicable in the base BGP routing context.

sr-policy-ipv4 — Matches routes belonging to the segment routing policy address family (AFI1/SAFI73).

flow-spec-dest

Syntax	flow-spec-dest <i>prefix-list-name</i> no flow-spec-dest
Context	config>router>policy-options>policy-statement>entry>from
Description	This command is used to match BGP flow-spec routes on the basis of the destination IP prefix in the flow specification. An IPv4 flow-spec route is matched by this command if its NLRI contains a type 1 subcomponent encoding a prefix and prefix-length that is covered by an entry in the referenced prefix-list. An IPv6 flow-spec route is matched by this command if its NLRI contains a type 1 component encoding prefix-offset=0 and a prefix & prefix-length that is covered by an entry in the referenced prefix-list. The flow-spec-dest command has no effect when the policy is not applied as a BGP import or export policy.
Default	no flow-spec-dest
Parameters	<i>prefix-list-name</i> — Specifies the name of a prefix-list containing IPv4 and/or IPv6 prefix entries [up to 32 characters].

flow-spec-source

Syntax	flow-spec-source <i>prefix-list-name</i> no flow-spec-source
Context	config>router>policy-options>policy-statement>entry>from
Description	This command is used to match BGP flow-spec routes on the basis of the source IP prefix in the flow specification. An IPv4 flow-spec route is matched by this command if its NLRI contains a type 2 subcomponent encoding a prefix and prefix-length that is covered by an entry in the referenced prefix-list. An IPv6 flow-spec route is matched by this command if its NLRI contains a type 2 component encoding prefix-offset=0 and a prefix & prefix-length that is covered by an entry in the referenced prefix-list. The flow-spec-source command has no effect when the policy is not applied as a BGP import or export policy.
Default	no flow-spec-source
Parameters	<i>prefix-list-name</i> — Specifies the name of a prefix-list containing IPv4 and/or IPv6 prefix entries, up to 32 characters.

group-address

Syntax	group-address <i>prefix-list-name</i> no group-address
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies the multicast group-address prefix list containing multicast group-addresses that are embedded in the join or prune packet as a filter criterion. The prefix list must be configured prior to entering this command. Prefix lists are configured in the config>router>policy-options>prefix-list context. The no form of the command removes the criterion from the configuration.
Default	no group-address
Parameters	<i>prefix-list-name</i> — Specifies the prefix-list name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The <i>prefix-list-name</i> is defined in the config>router>policy-options>prefix-list context.

host-ip

Syntax	host-ip <i>prefix-list-name</i>
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies a prefix list host IP address as a match criterion for the route policy-statement entry.
Default	no host-ip
Parameters	<i>prefix-list-name</i> — Specifies the prefix-list name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The <i>prefix-list-name</i> is defined in the config>router>policy-options>prefix-list context.

interface

Syntax	interface <i>interface-name</i> no interface
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies the router interface, specified either by name or address, as a filter criterion. The no form of the command removes the criterion from the configuration.
Default	no interface
Parameters	<i>ip-int-name</i> — Specifies the name of the interface as a match criterion for this entry. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

level

Syntax	level {1 2} no level
Context	config>router>policy-options>policy-statement>entry>from config>router>policy-options>policy-statement>entry>to
Description	This command specifies the ISIS route level as a match criterion for the entry.
Default	no level

Parameters **1 | 2** — Matches the IS-IS route learned from level 1 or level 2.

mvpn-type

Syntax **mvpn-type {1 | 2 | 3 | 4 | 5 | 6 | 7}**
no mvpn-type

Context config>router>policy-options>policy-statement>entry>from

Description This command allows match on ng-MVPN BGP route type when the policy is used for VRF-import/VRF-export/BGP global export policy. The policy will only be applied to multicast routes.

 The **no** form of the command disables **mvpn-type** in the policy evaluation.

Default no mvpn-type

Parameters **1 | 2 | 3 | 4 | 5 | 6 | 7** — BGP MVPN route-type as per RFC 6514.

neighbor

Syntax **neighbor {ip-address | prefix-list name}**
no neighbor

Context config>router>policy-options>policy-statement>entry>to
config>router>policy-options>policy-statement>entry>from

Description This command specifies the neighbor address as found in the source address of the actual join and prune message as a filter criterion. If no neighbor is specified, any neighbor is considered a match.

 The **no** form of the of the command removes the neighbor IP match criterion from the configuration.

Default no neighbor

Parameters *ip-address* — Specifies the neighbor IP address in dotted decimal notation.

Values ipv4-address:
 • a.b.c.d
 ipv6-address:
 • x:x:x:x:x:x:x [-interface]
 • x:x:x:x:x:d.d.d.d [-interface]
 • x: [0 to FFFF]H
 • d: [0 to 255]D
 • interface: 32 characters maximum, mandatory for link local addresses

prefix-list *name* — Specifies the prefix-list name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The *name* specified must already be defined.

origin

Syntax	origin { igp egp incomplete any aaa dynamic static bonding } no origin
Context	config>router>policy-options>policy-statement>entry>from
Description	This command configures a match criteria for the origin attribute. Originally, the origin attribute was applicable only to BGP as a mandatory well-known BGP attribute. The functionality of the origin attribute has expanded to subscriber-management routes (/32 IPv4 host and IPv6 PD WAN host routes). By default, each subscriber-management route will internally (local to the node) carry the origin attribute with one of three values (aaa, dynamic, and static). The value of the attribute depends on the origin of the subscriber-management route. The aaa, dynamic or static values are never carried in BGP updates as part of the BGP origin attribute and are visible within the BGP process. The values for the origin attribute in the subscriber-management routes allows customized advertisement of the subscriber-management routes by the routing policy.
Default	no origin
Parameters	igp — Configures matching path information originating within the local AS. egp — Configures matching path information originating in another AS. incomplete — Configures matching path information learned by another method. any — Specifies to ignore this criteria. aaa — Specifies to use the subscriber host address originated from AAA. Values IPv4 — subscriber-management /32 host routes that are originated from the RADIUS framed-ip-address VSA other than 255.255.255.254. The 255.255.255.254 returned by the RADIUS indicates that the BNG (NAS) should assign an IP address from its own pool. IPv6 — subscriber-management routes that are originated through framed-ipv6-prefix (SLAAC), delegated-ipv6-prefix (IA_PD) or alc-ipv6-address (IA_NA) RADIUS attributes. It is also applicable to VSA Alc-IPv6-Sub-If-Prefix, where the subscriber interface prefix can be originated from RADIUS. This is valid for IPoE and PPPoE type hosts.

dynamic — Specifies to use the subscriber host address originated from DHCP, DHCPv6, or the local address server.

Values IPv4 — subscriber-management /32 host routes that originate from the DHCP server (local or remote) or RADIUS framed-ip-address=255.255.255.254 (RFC 2865).
IPv6 — subscriber-management routes that are assigned via local DHCPv6 server pools whose name is obtained through the Alc-Delegated-IPv6-Pool (PD pool) and Framed-IPv6-Pool (NA pool) RADIUS attributes, or the local address server whose name is obtained through the Alc-SLAAC-IPv6-Pool (SLAAC pool) RADIUS attribute. This is valid for IPoE and PPPoE type hosts.
In addition, for IPoEv6 only, the pool name can also be obtained from ipv6-delegated-prefix-pool (PD pool) and ipv6-wan-address-pool (NA pool) from the LUDB.

static — Specifies to use the subscriber host address originated from the local user database.

Values IPv4 — subscriber-management /32 host routes that are originated from the LUDB. This should also cover the RADIUS fallback category (RADIUS falls back to system defaults or to the LUDB).
IPv6 — subscriber-management routes obtained from the LUDB from ipv6-address (IA_NA) or ipv6-prefix (IA_PD), or ipv6-slaac-prefix (SLAAC).

origin-validation-state

Syntax	origin-validation-state state no origin-validation-state
Context	config>router>policy-options>policy-statement>entry>from
Description	This command is used to match BGP routes on the basis of origin validation state: • Valid (0) • Not-Found (1) • Invalid (2)
Default	no origin-validation-state
Parameters	valid — Marks the route as having an origin validation state of valid. notFound — Marks the route as having an origin validation state of Not Found. invalid — Marks the route as having an origin validation state of invalid.

policy-statement

Syntax	[no] policy-statement <i>name</i>
Context	config>router>policy-options
Description	This command creates the context to configure a route policy statement. Route policy statements control the flow of routing information to and from a specific protocol, set of protocols, or to a specific BGP neighbor. The policy-statement is a logical grouping of match and action criteria. A single policy-statement can affect routing in one or more protocols and/or one or more protocols peers/neighbors. A single policy-statement can also affect both the import and export of routing information. The no form of the command deletes the policy statement.
Default	no policy-statement
Parameters	<i>name</i> — Specifies the route policy statement name. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

policy

Syntax	policy <i>plcy-or-long-expr</i> no policy
Context	config>router>policy-options>policy-statement>entry>from
Description	This command is used to call another policy by name and evaluate it as a subroutine, or to evaluate a logical expression of subroutine policies. If the result of the subroutine evaluation is an 'accept', then the route is considered to match the entry in the parent policy that called the subroutine. If the result of the subroutine evaluation is a 'reject', then the route is considered a non-match of the entry in the parent policy that called the subroutine. Up to 3 levels of subroutine calls are supported. If a subroutine at maximum depth has this command, it is automatically considered a non-match of all routes. The no form of the command removes the policy statement as a match criterion.
Default	no policy

Parameters *plcy-or-long-expr* — Specifies the name of a single **policy-statement** (up to 64 characters in length) or a policy logical expression (up to 255 characters in length) consisting of **policy-statement** names (enclosed in square brackets), logical operations ('and', 'or', 'not'), and parentheses for grouping.

policy

Syntax **policy** *plcy-or-long-expr*
no policy

Context config>router>policy-options>policy-statement>entry>action

Description This command configures a nested policy statement as a match criterion for the route policy entry.

Policy statements are configured at the global route policy level (**config>router>policy-options policy-statement**).

The command is used to call another policy by name and evaluate it as a subroutine. If the result of the subroutine evaluation is an 'accept', then the route is considered to match the entry in the parent policy that called the subroutine. If the result of the subroutine evaluation is a 'reject', then the route is considered a non-match of the entry in the parent policy that called the subroutine. Up to 3 levels of subroutine calls are supported. If a subroutine at maximum depth has this command, it is automatically considered a non-match of all routes.

The **no** form of the command removes the policy statement as a match criterion.

Default no policy

Parameters *plcy-or-long-expr* — Specifies the route policy name (up to 64 characters long) or a policy logical expression (up to 255 characters). Allowed values are any string up to 255 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

policy-variables

Syntax **policy-variables**

Context config>router>policy-options>policy-statement>from

Description This command enables the **policy-variables** configuration context.

The **no** form of the command removes all policy variables.

prefix-list

Syntax	prefix-list <i>name</i> [<i>name</i>] no prefix-list
Context	config>router>policy-options>policy-statement>entry>from config>router>policy-options>policy-statement>entry>to
Description	This command configures a prefix list as a match criterion for a route policy statement entry. If no prefix list is specified, any network prefix is considered a match. An empty prefix list will evaluate as if 'no match' was found. The prefix lists specify the network prefix (this includes the prefix and length) a specific policy entry applies. A maximum of 28 prefix names can be specified. The no form of the command removes the prefix list match criterion.
Default	no prefix-list
Parameters	<i>name</i> — Specifies the prefix list name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".

protocol

Syntax	protocol <i>protocol</i> [all instance <i>instance</i>] protocol bgp <i>bgp-label</i> no protocol
Context	config>router>policy-options>policy-statement>entry>from
Description	This command configures a routing protocol as a match criterion for a route policy statement entry. This command is used for both import and export policies depending how it is used. If no protocol parameter is specified, any protocol is considered a match. The no form of the command removes the protocol match criterion.
Default	no protocol

Parameters	<i>protocol</i> — Specifies the protocol name to match on.
Values	direct, static, bgp, isis, ospf, rip, aggregate, bgp-vpn, igmp, pim, ospf3, ldp, sub-mgmt, mld, managed, vpn-leak, nat, periodic, ipsec, dhcpv6-pd, dhcpv6-na, dhcpv6-ta, dhcpv6-pd-excl, ripng, bgp-label, direct-interface, arp-nd
	<i>instance</i> — Specifies the OSPF, OSPF3, or IS-IS instance.
Values	isis-inst — 0 to 31 ospf-inst — 0 to 31 ospf3-inst — 0 to 31, 64 to 95
	all — Keyword used only with OSPF or IS-IS that specifies all instances.

protocol

Syntax	protocol <i>protocol</i> [all instance <i>instance</i>] protocol bgp <i>bgp-label</i> no protocol
Context	config>router>policy-options>policy-statement>entry>to
Description	This command configures a routing protocol as a match criterion for a route policy statement entry. This command is used for both import and export policies depending how it is used. If no protocol parameter is specified, any protocol is considered a match. The no form of the command removes the protocol match criterion.
Default	no protocol
Parameters	<i>protocol</i> — Specifies the protocol name to match on.
Values	bgp, isis, ospf, rip, bgp-vpn, ospf3, vpn-leak, ldp, ripng, bgp-label
	<i>instance</i> — Specifies the OSPF, OSPF3, or IS-IS instance.
Values	isis-inst — 0 to 31 ospf-inst — 0 to 31 ospf3-inst — 0 to 31, 64 to 95
	all — Keyword used only with OSPF or IS-IS that specifies all instances.

renumber

Syntax	renumber from <i>entry-id</i> to <i>entry-id</i>
Context	config>router>policy-options>policy-statement

Description This command allows the operator to renumber the existing entry ID to a new entry ID. When performing the renumbering action, the two entry IDs must be different. The existing (**from**) *entry-id* must exist. The new (**to**) *entry-id* must not exist.

Renumbering is not saved in the configuration because it is a performing action.

Parameters **from** *entry-id* — Specifies the existing entry ID to be renumbered.

Values 1 to 4294967295

to *entry-id* — Specifies the new entry ID to be assigned.

Values 1 to 4294967295

source-address

Syntax **source-address** *ip-address*
source-address **prefix-list** *prefix-list-name*
no source-address

Context config>router>policy-options>policy-statement>entry>from

Description This command specifies the source address that is embedded in the join or prune packet as a filter criterion.

The **no** form of the command removes the criterion from the configuration.

This command specifies a multicast data source address as a match criterion for this entry.

Default no source-address

Parameters *ip-address* — Specifies the IP prefix for the IP match criterion in dotted decimal notation.

Values ipv4-address:
• a.b.c.d
ipv6-address:
• x:x:x:x:x:x:x
• x:x:x:x:x:d.d.d.d
• x: [0 to FFFF]H
• d: [0 to 255]D

prefix-list-name — The prefix list name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

state

Syntax	state state no state
Context	config>router>policy-options>policy-statement>entry>from
Description	This command will configure a match criteria on the state attribute. The state attribute carries the state of an SRRP instance and it can be applied to: • subscriber-interface routes • subscriber-management routes (/32 IPv4 and IPv6 PD wan-host) • managed-routes (applicable only to IPv4). Based on the state attribute of the route we can manipulate the route advertisement into the network. We can enable or disable (in case there is no SRRP running) tracking of SRRP state by routes. This is done on a per subscriber-interface route basis, where a subscriber-interface route is tracking a single SRRP instance state (SRRP instance might be in a Fate Sharing Group). For subscriber-management and managed-routes, tracking is enabled per group interface under which SRRP is enabled. This command specifies a multicast data source address as a match criterion for this entry.
Default	no state
Parameters	srrp-master — Track routes with the state attribute carrying srrp-master state. srrp-non-master — Track routes with the state attribute carrying srrp-non-master state. ipsec-master-with-peer — Track routes with the state attribute carrying ipsec-master-with-peer state. ipsec-non-master — Track routes with the state attribute carrying ipsec-non-master state. ipsec-master-without-peer — Track routes with the state attribute carrying ipsec-master-without-peer state.

tag

Syntax	tag {no-tag tag} no tag
Context	config>router>policy-options>policy-statement>entry>from

Description	This command matches the tag value in static or IGP routes. A decimal or hexadecimal value of 4 octets can be entered. For IS-IS, OSPF, and static routes, all four octets can be used. For RIP and RIPng, only the two most significant octets are used if more than two octets are configured. The no form of the command removes the tag field match criterion.
Default	no tag
Parameters	<i>tag</i> — Matches the configured tag value. Values Accepts decimal or hexadecimal formats: • IS-IS, OSPF and static routes: 0x0 – 0xFFFFFFFF or 1 – 4294967295 • RIP and RIPng: 0x0 – 0xFFFF or 1 – 65535 no-tag — Specifies that no tag value is set.

evpn-type

Syntax	evpn-type <i>type</i> no evpn-type
Context	config>router>policy-options>policy-statement>entry>from
Description	This command matches BGP routes based on the EVPN route type. All route types, supported in SROS, are the following: • Type 1 or Auto-Discovery Ethernet Tag route; including both the AD per-ES and AD per-EVI routes • Type 2 or MAC/IP route • Type 3 or IMET (Inclusive Multicast Ethernet Tag) route • Type 4 or ES (Ethernet Segment) route • Type 5 of IP-prefix route (includes IPv4 and IPv6 prefixes) The no version of this command removes the evpn-type matching.
Parameters	<i>type</i> — Specifies the EVPN route type. Values 1 to 5

to

Syntax	[no] to
Context	config>router>policy-options>policy-statement>entry

Description	This command creates the context to configure export policy match criteria based on a route's destination or the protocol into which the route is being advertised. If no condition is specified, all route destinations are considered to match. The to command context only applies to export policies. If it is used for an import policy, match criteria is ignored. The no form of the command deletes export match criteria for the route policy statement entry.

type

Syntax	type {1 2} no type
Context	config>router>policy-options>policy-statement>entry>from
Description	This command configures an OSPF type metric as a match criterion in the route policy statement entry. If no type is specified, any OSPF type is considered a match. The no form of the command removes the OSPF type match criterion.
Default	no type
Parameters	1 — Matches OSPF routes with type 1 LSAs. 2 — Matches OSPF routes with type 2 LSAs.

6.6.2.7 Route Policy Action Commands

action

Syntax	action {accept next-entry next-policy drop reject} no action
Context	config>router>policy-options>policy-statement>entry
Description	This command creates the context to configure actions to take for routes matching a route policy statement entry. This command is required and must be entered for the entry to be active. Any route policy entry without the action command will be considered incomplete and will be inactive.

The **no** form of the command deletes the action context from the entry.

Default	no action
Parameters	accept — Specifies that routes matching the entry match criteria will be accepted and propagated. next-entry — Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next policy entry (if any others are specified). next-policy — Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next route policy (if any others are specified). drop — Specifies that routes matching the entry match criteria should be rejected. This parameter provides a context for modifying route properties. reject — Specifies that routes matching the entry match criteria should be rejected. This parameter does not provide a context for modifying route properties.

add-paths-send-limit

Syntax	add-paths-send-limit <i>send-limit</i> no add-paths-send-limit
Context	config>router>policy-options>policy-statement>entry config>router>policy-options>policy-statement>default-action
Description	This command sets the Add-Paths send-limit to a specific value for all routes matched by the policy entry or default action. Add-Paths allows a BGP router to send multiple paths for the same NLRI/prefix to a peer advertising the Add-Paths receive capability. The send-limit dictates the maximum number of paths that can be advertised. The default send-limit is controlled by the instance, group or neighbor level configuration and applies to all prefixes in a particular address family. Using route policies allows the default send-limit to be overridden to use a larger or smaller maximum value on a per-prefix basis. For example if, for most prefixes advertised to a peer, at most 1 path should be advertised but for a few exceptional prefixes up to 4 paths should be advertised then the neighbor-level send-limit can be set to a value of 1 and the add-paths-send-limit in the policy entry that matches the exceptional routes can be set to a value of 4.
Default	no add-paths-send-limit
Parameters	send-limit — Specify the maximum number of paths to advertise for matched routes to an Add-Paths peer. Values 1to16

admin-tag-policy

Syntax	admin-tag-policy <i>policy-name</i> no admin-tag-policy
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command assigns a route admin tag policy as an action in a route policy. The admin tag policy must exist under config>router>admin-tags. The no form of the command removes the admin tag policy.
Parameters	<i>policy-name</i> — Specifies the name of the admin tag policy, up to 64 characters.

advertise-label

Syntax	advertise-label { per-prefix pop } no advertise-label
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	The effect of the advertise-label command depends on the context where the associated policy is applied. When the per-prefix option is used and the command is configured as the default action or entry-specific action of a VRF export policy, every qualifying matched route is advertised with a per-prefix label in the resulting VPN-IP routes. In this situation, non-qualifying routes include local interface routes and BGP-VPN routes. The command overrides, for specific routes, the configured label-mode of the exporting VPRN service. When configured with the per-prefix option, the command also affects BGP import policies applied to a base router BGP peer. When a label-IPv4 route is matched and accepted by a BGP import policy entry or default action with this command, and it is the best path for the prefix in the label-IPv4 RIB, a per-prefix label is used in the advertised route if there is a BGP next-hop change. A label-IPv4 route advertised with a pre-prefix label supports ECMP forwarding across multiple BGP next-hops. When configured with the pop option, the command also serves a purpose in route-table-import policies. When a /32 IPv4 static, OSPF, or IS-IS route is matched and accepted by a label-IPv4 RIB route-table-import policy entry or default-action with this command, and the route is a candidate to be advertised as a label-IPv4 route (due to a BGP export policy), the advertised BGP label is programmed for a Pop operation. When a /32 static, OSPF, or IS-IS route is imported into the label-IPv4 RIB and then exported as a BGP route, the default behavior is to program a “swap” operation in the datapath, which swaps the BGP label with the tunnel label that takes traffic to the destination of the /32 route.

Default	no advertise-label
Parameters	per-prefix — Requests per-prefix label allocation for matched routes. This takes effect only in VRF export policies and BGP import policies, and only for certain types of routes. pop — Requests Pop label allocation for matched routes. This takes effect only in label-IPv4 route-table-import policies and only applies to /32 IPv4 routes that were learned through static configuration, OSPF, or IS-IS.

aigp-metric

Syntax	aigp-metric <i>metric</i> aigp-metric add aigp-metric igp no aigp-metric
Context	config>router>policy-options>policy-statement>entry>action config>router>policy-options>policy-statement>default-action
Description	This command assigns a BGP AIGP metric to routes matching the entry. The effect of this command on a route matched and accepted by a route policy entry depends on how the policy is applied (BGP import policy vs. BGP export policy), the type of route and the specific form of the command.

In a BGP import policy this command is used to:

- Associate an AIGP metric with an IBGP route received with an empty AS path and no AIGP attribute.
- Associate an AIGP metric with an EBGP route received without an AIGP attribute that has an AS path containing only AS numbers belonging to the local AIGP administrative domain.
- Modify the received AIGP metric value prior to BGP path selection

In a BGP export policy this command is used to:

- Add the AIGP attribute and set the AIGP metric value in a BGP route originated by exporting a direct, static or IGP route from the routing table
- Remove the AIGP attribute from a route advertisement to a particular peer
- Modify the AIGP metric value in a route advertisement to a particular peer

Default	no aigp-metric				
Parameters	<i>metric</i> — Administratively defined metric. <table> <tr> <td>Values</td><td>0 to 4294967295</td></tr> <tr> <td>Default</td><td>none</td></tr> </table>	Values	0 to 4294967295	Default	none
Values	0 to 4294967295				
Default	none				

name — The AIGP metric parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

add — Adds the AIGP attribute.

igp — Sets the AIGP metric to the IGP metric.

as-path

Syntax	as-path {add replace} <i>name</i> no as-path
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command assigns a BGP AS path list to routes matching the route policy statement entry. If no AS path list is specified, the AS path attribute is not changed. The no form of the command disables the AS path list editing action from the route policy entry.
Default	no as-path
Parameters	add — Specifies that the AS path list is to be prepended to an existing AS list. replace — Specifies AS path list replaces any existing as path attribute. <i>name</i> — Specifies the AS path list name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@". The <i>name</i> specified must already be defined.

as-path-prepend

Syntax	as-path-prepend <i>as-path</i> [<i>repeat</i>] as-path-prepend most-recent [<i>repeat</i>] no as-path-prepend
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action

Description	The command prepends a BGP AS number once or numerous times to the AS path attribute of routes matching the route policy statement entry. If an AS number is not configured, the AS path is not changed. If the optional <i>number</i> is specified, then the AS number is prepended as many times as indicated by the number. The no form of the command disables the AS path prepend action from the route policy entry.
Default	no as-path-prepend
Parameters	<i>as-path</i> — Specifies the AS number to prepend expressed as a decimal integer. Values 1 to 4294967295 param-name — Specifies the AS path parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@". <i>repeat</i> — Specifies the number of times to prepend the specified AS number expressed as a decimal integer. Values 1 to 50 param-name — Specifies the AS path parameter variable name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@". most-recent — Specifies that the most recent AS number must be prepended to the AS-Path attribute of the route.

bgp-leak

Syntax	bgp-leak no bgp-leak
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command causes qualifying matched BGP routes to be marked as leakable, meaning they are candidates to be leaked into other routing instances (copied with their complete set of path attributes). A BGP route is a qualifying route if the NLRI has an IPv4 or IPv6 prefix without a label.



Note: A leakable BGP route is not actually leaked into another routing instance unless it is accepted by a leak-import policy of that other routing instance.

The **bgp-leak** command has an effect only when the policy is applied as a BGP import policy in the base router or a VPRN context.

Default no bgp-leak

community

Syntax **community add** *name* [*name*]
community remove *name* [*name*]
community replace *name* [*name*]
no community

Context config>router>policy-options>policy-statement>default-action
 config>router>policy-options>policy-statement>entry>action

Description This command adds or removes a BGP community list to or from routes matching the route policy statement entry.

If no community list is specified, the community path attribute is not changed.

The community list changes the community path attribute according to the **add** and **remove** keywords.

The **no** form of the command disables the action to edit the community path attribute for the route policy entry.

Default no community

Parameters *name* — Specifies up to 28 names.

add — The specified community list is added to any existing list of communities.

remove — The specified community list is removed from the existing list of communities.

replace — The specified community list replaces any existing community attribute.

name — The community list name. Allowed values are any string up to 64 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".

community

Syntax	community <i>comm-name</i> community expression <i>expression</i> no community
Context	config>router>policy-options>policy-statement>entry>from
Description	This command adds or removes a BGP community list to or from routes matching the route policy statement entry. If no community list is specified, the community path attribute is not changed. The community list changes the community path attribute according to the add and remove keywords. The no form of the command disables the action to edit the community path attribute for the route policy entry.
Default	no community
Parameters	<i>comm-name</i> — Specifies up to 28 names. <i>expression</i> — Applies parameters to routes matching the entry. Values expression is one of the following up to 900 characters: <expression> {AND OR} <expression> [NOT] (<expression>) [NOT] "["<comm-name>"]

damping

Syntax	damping { <i>name</i> none } no damping
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command configures a damping profile used for routes matching the route policy statement entry. If no damping criteria is specified, the default damping profile is used. The no form of the command removes the damping profile associated with the route policy entry.
Default	no damping

- Parameters** *name* — The damping profile name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by at-signs (@) and may be midstring; for example, "@variable@", "start@variable@end", "@variable@end", or "start@variable@".
- The *name* specified must already be defined.
- none** — Disables route damping for the route policy.

create-udp-tunnel

- Syntax** **create-udp-tunnel**
no create-udp-tunnel
- Context** config>router>policy-options>policy-statement >default-action
config>router>policy-options>policy-statement>entry>action
- Description** This command instructs the router to create an MPLS-over-UDP tunnel upon receiving BGP routes that match the import policy.
- Default** no create-udp-tunnel

fc

- Syntax** **fc** *fc* [priority {low | high}]
no fc
- Context** config>router>policy-options>policy-statement>entry>action\$
- Description** This command associates a forwarding-class and optionally priority with the routes matched by a route policy entry. The command takes effect when the action of the route policy entry is accept, next-entry or next-policy. It has no effect except in route policies applied as VRF import policies, BGP import policies or RIP import policies.
- The **no** form of the command removes the QoS association of the routes matched by the route policy entry.
- Default** no fc
- Parameters** *fc* — Specify the name of one of the predefined forwarding classes in the system.
- Values** be, l2, af, l1, h2, ef, h1, nc
- Default** none (no QoS information is associated with matched routes)

priority {low | high} — This parameter associates an enqueueing priority with routes matched by the policy entry. Specifying a priority is optional.

Values **high** — Setting the enqueueing parameter to **high** for a packet increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

low — Setting the enqueueing parameter to **low** for a packet decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing, once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.

Default low

default-action

Syntax **default-action {accept | next-entry | next-policy | drop | reject}**
no default-action

Context config>router>policy-options>policy-statement

Description This command enables the context to configure actions to apply to routes that do not match any entries of a route policy statement.

The **no** form of the command deletes the **default-action** context for the policy statement.

Default no default-action

Parameters **accept** — Specifies that routes not matched by any entry should be allowed or accepted. This parameter provides a context for modifying route properties.

next-entry — Specifies that routes not matched by any entry should be evaluated by the next sequential entry in the policy chain, after route properties are possibly modified by the default action of the current policy.

next-policy — Specifies that routes not matched by any entry should be evaluated by the next sequential policy in the policy chain, after route properties are possibly modified by the default action of the current policy.

drop — Specifies that routes not matched by any entry should be disallowed or rejected. This parameter provides a context for modifying route properties.

reject — Specifies that routes not matched by any entry should be disallowed or rejected. This parameter does not provide a context for modifying route properties.

install-backup-path

Syntax **install-backup-path**

no install-backup-path

Context	config>router>policy-options>policy-statement>entry>action config>router>policy-options>policy-statement>default-action
Description	When the best BGP route for an IPv4 or IPv6 prefix is matched by a policy entry or policy default action with this command, BGP attempts to find and install a pre-programmed backup path for the prefix in order to provide BGP fast reroute protection. The install-backup-path command overrides and has no dependency on commands such as the BGP instance backup-path command or the VPRN-level enable-bgp-vpn-backup command, which enable BGP fast reroute for an entire address family. The install-backup-path command provides more precise control over which IP prefixes are supported with pre-programmed backup paths. If, within a VPRN, the best path for an IP prefix is provided by a VPRN BGP route, the backup path can be provided by another VPRN BGP route or an imported VPN-IP route. If, within a VPRN, the best path for an IP prefix is provided by an imported VPN-IP route, the backup path can be provided by another VPN-IP route. The install-backup-path command is supported only in BGP import policies and VRF import policies and has no effect on policy types other than BGP import policies and VRF import policies. The install-backup-path command applies only to the following types of matched routes: IPv4, IPv6, label-IPv4, label-IPv6, VPN-IPv4, and VPN-IPv6.
Default	no install-backup-path

local-preference

Syntax	local-preference <i>preference</i> no local-preference
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry
Description	This command assigns a BGP local preference to routes matching a route policy statement entry. If no local preference is specified, the BGP configured local preference is used. The no form of the command disables assigning a local preference in the route policy entry.
Default	no local-preference

Parameters	<i>preference</i> — The local preference expressed as a decimal integer.
Values	0 to 4294967295
	<i>name</i> — The local preference parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

metric

Syntax	metric {add subtract} <i>metric</i> metric set {[gp <i>metric-value</i>] no metric
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	In a BGP import or export policy, this command assigns a MED value to routes matched by the policy statement entry. The MED value may be set to a fixed value (overriding the received value), set to the routing table cost of the route used to resolve the NEXT_HOP of the BGP route (igp option), or modified by adding or subtracting a fixed value offset. The no form of the command removes the MED attribute from the matched routes.
Default	no metric
Parameters	add — Specifies that an integer is added to any existing metric. If the result of the addition results in a number greater than 4294967295, the value 4294967295 is used. subtract — Specified <i>integer</i> is subtracted from any existing metric. If the result of the subtraction results in a number less than 0, the value of 0 is used. set — Specified <i>integer</i> replaces any existing metric. igp — Sets the MED value to the routing table cost of the route used to resolve the NEXT_HOP of the BGP route. <i>metric</i> — The metric modifier expressed as a decimal integer. Values 0 to 4294967295 <i>name</i> —The metric parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@"

multicast-redirection

Syntax	multicast-redirection [fwd-service <i>service-id</i>] <i>ip-int-name</i> no multicast-redirection
Context	config>router>policy-options>policy-statement>default-action
Description	This command enables a redirection under a filtering policy. The filtering policy in this case becomes a redirection policy and it is defined under the router>policy-option hierarchy. After the redirection policy is applied to the subscriber, all IGMP messages will be processed per subscriber host before they get redirected to the referenced interface (and possibly service). However, multicast traffic will not be replicated directly per subscriber host but instead it will be forwarded on the interface that is referenced in the redirection policy. The redirected interface must have IGMP enabled. Currently all traffic is redirected and there is no ability to selectively redirect multicast traffic based on match conditions (multicast-groups, source IP address of IGMP messages, etc). Multicast redirection is supported between VPRN services and also between interfaces within the Global Routing Context. Multicast redirection is not supported between the VPRN services and the Global Routing Context. Multicast redirection is supported in the wholesale/retail VPRN context.  Note: Redirecting from a VPRN instance to the GRT is not supported. Redirecting from a VPRN to a different VPRN is supported and redirecting from an IES to another IES is also supported.
Default	no multicast-redirection
Parameters	fwd-service <i>service-id</i> — Specifies the service to which traffic should be redirected. This option is applied only in the VPRN context. It is possible to redirect the multicast group into another service instance routing interface. <i>ip-int-name</i> — specifies the alternate interface to which IGMP messages are redirected.

next-hop

Syntax	next-hop <i>ip-address</i> no next-hop
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command assigns the specified next hop IP address to routes matching the policy statement entry. If a next-hop IP address is not specified, the next-hop attribute is not changed.

The **no** form of the command disables assigning a next hop address in the route policy entry.

Default	no next-hop
Parameters	<i>ip-address</i> — The next hop IP address in dotted decimal notation.
Values	ipv4-prefix: • a.b.c.d (host bits must be 0) ipv4-prefix-length: [0 to 32] ipv6-prefix: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D name — The next-hop parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

next-hop-self

Syntax	[no] next-hop-self
Context	config>router>policy-options>policy-statement>entry>action config>router>policy-options>policy-statement>default-action
Description	This command configures BGP to advertise routes that match a policy entry (or that match no other policy entry and, therefore, to which the default action applies) using a local address of the BGP instance as the BGP next-hop address. The command applies to IPv4, IPv6, label-IPv4, and label-IPv6 routes. It also applies to VPN-IPv4 and VPN-IPv6 routes, but only when used in conjunction with the enable-rr-vpn-forwarding command. This command affects how routes are advertised to IBGP peers, regardless of whether or not they were learned from an IBGP or EBGP peer The no form of the command uses protocol standard behavior to decide whether or not to set next-hop-self in advertised routes.
Default	no next-hop-self

origin

Syntax	origin {igp egp incomplete param-name} no origin

Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command sets the BGP origin assigned to routes exported into BGP. If the routes are exported into protocols other than BGP, this option is ignored. The no form of the command disables setting the BGP origin for the route policy entry.
Default	no origin
Parameters	igp — Sets the path information as originating within the local AS. egp — Sets the path information as originating in another AS. incomplete — Sets the path information as learned by some other means. <i>param-name</i> — The origin parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

origin-validation-state

Syntax	origin-validation-state { <i>state</i> <i>param-name</i> } no origin-validation-state
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command is used to mark BGP IPv4 and IPv6 routes matching the default-action or a specific entry of a route policy with one of the 3 following origin validation states: • Valid (0) • Not-Found (1) • Invalid (2)
Default	no origin-validation-state
Parameters	<i>state</i> — Specifies the default operational origin validation state for this policy statement. Values valid — Marks the route as having an origin validation state of valid. notFound — Marks the route as having an origin validation state of Not Found. invalid — Marks the route as having an origin validation state of invalid.

preference

Syntax	preference <i>preference</i> no preference
Context	config>router>policy-options>policy-statement name>default-action config>router>policy-options>policy-statement>entry>action
Description	This command assigns a route preference to routes matching the route policy statement entry. If no preference is specified, the default Route Table Manager (RTM) preference for the protocol is used. The no form of the command disables setting an RTM preference in the route policy entry.
	Note: This command is supported with the following protocols: RIP import, BGP import, VPRN VRF import (vrf-import), and VPRN GRT lookup export (export-grt).
Default	no preference
Parameters	<i>preference</i> — The route preference expressed as a decimal integer. Values 1 to 255 (0 represents unset - MIB only) <i>name</i> — The preference parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

resolve-static

Syntax	[no] resolve-static
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command has an affect only in BGP route-table-import policies and applies only to BGP IPv4 and IPv6 routes created by importing static routes with indirect next-hops. When such a route matches a policy entry with this action, the BGP next-hop is the resolved next-hop of the static route. The no form of the command reverts to the default behavior, which copies the indirect next-hop of the static route into the BGP next-hop without resolving it further.
Default	no resolve-static

source-class

Syntax	source-class [<i>value</i>] no source-class
Context	config>router>policy-options>policy-statement>entry>action config>router>policy-options>policy-statement>default-action
Description	This command configures the default source-class for this policy-statement.
Parameters	<i>value</i> — Specifies the default operational source-class for this policy statement. Values 1 to 255

sticky-ecmp

Syntax	sticky-ecmp no sticky-ecmp
Context	config>router>policy-options>policy-statement>entry>action config>router>policy-options>policy-statement>default-action
Description	This command specifies that BGP routes matching an entry or default-action of a route policy should be tagged internally as requiring sticky ECMP behavior. When a BGP route with multiple equal-cost BGP next-hops is programmed for sticky ECMP the failure of one or more of its BGP next-hops causes only the affected traffic flows to be re-distributed to the remaining next-hops; by default (without sticky-ECMP) all flows are potentially affected, even those using a next-hop that did not fail.
Default	no sticky-ecmp

tag

Syntax	tag { no-tag <i>tag</i> } no tag
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command assigns a tag to routes matching the entry, which is then applied to IGP routes. A decimal or hexadecimal value of 4 octets can be entered. For IS-IS and OSPF, all four octets can be used. For RIP and RIPng, only the two most significant octets are used if more than two octets are configured. The no form of the command removes the tag.

Default	no tag
Parameters	<i>tag</i> — Assigns an IS-IS, OSPF, RIP or RIPv6 tag to routes matching the entry.
Values	Accepts decimal or hexadecimal formats: IS-IS and OSPF: 0x0–0xFFFFFFFF or 1–4294967295 RIP and RIPv6: 0x0–0xFFFF or 1–65535 <i>name</i> — The tag parameter variable name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Policy parameters must start and end with at-signs (@); for example, "@variable@".

type

Syntax	type <i>type</i> no type
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	This command sets the subtype for the Type 5 LSA (external LSA). The no form of the command disables assigning a type in the route policy entry.
Default	type 2
Parameters	<i>type</i> — Specifies the type metric.
Values	Subtype 1 — The external metric in the external LSA is comparable with the internal metric, and thus one can sum up all the metrics along the path (both internal and external) to get the total cost to the destination. Subtype 2 — The metric in the external LSA is much more important than the internal metric, so the internal metrics should only be considered when comparing two external routes that have the same external metric.

route-exists

Syntax	route-exists <i>expression</i> no route-exists
Context	config>router>policy-options>policy-statement>entry>cond-expr

Description	This command is used to specify a route existence expression to control evaluation of the policy entry. If the route existence expression evaluates to 'true' the matching and action commands of the policy entry are applied as normal. If the route existence expression evaluates to 'false' the entire policy entry is skipped and processing continues with the next entry; however, conditional expressions are only parsed when the route policy is used as a BGP export policy or VRF export policy.
Default	no route-exists
Parameters	<i>expression</i> — “[<prefix-list-name>]” [all none] If neither the all nor the none keyword are used the match logic is 'any' – that is, the route expression evaluates as 'true' if any exact match entry in the referenced prefix-list has an active route in the route table associated with the policy. all – the route expression evaluates as 'true' only if all the exact match entries in the referenced prefix-list have an active route in the route table associated with the policy. none – the route expression evaluates as 'true' only if none of the exact match entries in the referenced prefix-list have an active route in the route table associated with the policy.

6.7 Show, Clear, and Debug Command Reference

6.7.1 Command Hierarchies

- [Show Commands](#)

6.7.1.1 Show Commands

```

show
  — router router-name
    — policy [name [expand-variables] | damping [name] | prefix-list [name] | as-path
      [name] | as-path-group [name] | community [name] | admin | exclusive-access]
    — policy name sub-policy entry entry-id [name subpol-name]
    — policy name sub-policy entry entry-id name subpol-name entry-2 entry-id name-2
      subpol-name
    — policy name sub-policy entry entry-id name subpol-name entry-2 entry-id name-2
      subpol-name entry-3 entry-id name-3 subpol-name

show
  — router router-name
    — bgp
      — policy-test plcy-or-long-expr [plcy-or-expr [plcy-or-expr .. (up to 14 max)]]
        [family family] [rd rd] [prefix ip-prefixip-prefix-length] [longer] [neighbor
          ip-address] [display-rejects] [{detail | brief}]

```

6.7.2 Command Descriptions

6.7.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

policy

Syntax **policy** [*name* **expand-variables**] | **damping** [*name*] | **prefix-list** [*name*] | **as-path** [*name*] | **as-path-group** [*name*] | **community** [*name*] | **admin** | **exclusive-access**
policy *name* **sub-policy** **entry** *entry-id* [**name** *subpol-name*]
policy *name* **sub-policy** **entry** *entry-id* **name** *subpol-name* **entry-2** *entry-id* **name-2** *subpol-name*
policy *name* **sub-policy** **entry** *entry-id* **name** *subpol-name* **entry-2** *entry-id* **name-2** *subpol-name* **entry-3** *entry-id* **name-3** *subpol-name*

Context show>router

Description This command displays configured policy statement information.

Parameters **admin** — When this keyword is specified, the entire policy option configuration is displayed, including any uncommitted configuration changes. This command is similar to the **info** command.

as-path — When this keyword is specified, the parameter displays configured as-path regular expressions (optionally filtered by name).

as-path-group — When this keyword is specified, the command displays configured as-path-groups, optionally filtered by name.

community — When this keyword is specified, the command displays configured community lists, optionally filtered by name.

damping — When this keyword is specified, the information about the route policy damping configurations is displayed.

entry *entry-id* — Displays the configuration (with parameter value substitution) of the subroutine policy called by the entry with entry ID.

exclusive-access — Displays the policy that edits exclusive access and lock information.

expand-variables — Displays the variables that are replaced with values in main policies or in sub-policies.

prefix-list *name* — Displays the prefix lists configured in the route policy.

policy *name* — Displays information similar to the info command for a specific policy-statement. If a *name* is provided, the matching policy-statement displays. If no *statement* name is specified, a list of all policies statements and descriptions display.

sub-policy — When this keyword is specified, the configuration of a subroutine policy is displayed.

subpol-name — Specifies the policy name up to 64 characters.

entry-id — Specifies the numeric entry identifier of the policy.

Values 1 to 4294967295

Output [Table 71](#) describes route policy output fields.

Table 71 **Route Policy Output Fields**

Label	Description
Policy	Displays a list of route policy names.
Description	Displays the description of each route policy.
Policies	The total number of policies configured.

The following route policy commands are displayed with different command parameter options:

- show router policy
- show router policy admin
- show router policy "BGP to RIP"
- show router policy damping
- show router policy "pol1" sub-pol entry 10 name "pol2"

The following output is an example of route policy information.

Sample Output

The **show router policy** command displays all configured route policies.

```
A:ALA-1# show router policy
=====
Route Policies
=====
Policy                               Description
-----
Direct And Aggregate                 Policy Statement ABC
-----
Policies :
=====
A:ALA-1#
```

The **show router policy admin** command is similar to the **info** command which displays information about the route policies and parameters.

```
A:ALA-1# show router policy admin
  prefix-list "All-Routes"
    prefix 0.0.0.0/0 longer
    prefix 2.0.0.0/8 longer
    prefix 3.0.0.0/8 longer
    prefix 4.0.0.0/8 longer
    prefix 5.0.0.0/8 longer
    prefix 6.0.0.0/8 exact
    prefix 224.0.0.0/24 longer
  exit
...
A:ALA-1#
```

The **show router policy *name*** command displays information about a specific route policy.

show router policy “BGP To RIP”

```
description "Policy Statement For 'BGP To RIP'"
  entry 10
    description "Entry For Policy 'BGP To RIP'"
    from
      protocol bgp
    exit
    to
      protocol rip
    exit
    action accept
      metric set 1
      next-hop 10.0.18.200
      tag 805
    exit
  exit
  default-action reject
A:ALA-1#
```

The **show router policy damping** command displays information about the route policy damping configurations.

A:ALA-1# show router policy damping

```
=====
Route Damping Profiles
=====
  damping "TEST-LOW"
    half-life 22
    max-suppress 720
    reuse 10000
    suppress 15000
  exit
  damping "TEST-HIGH"
    half-life 22
    max-suppress 720
    reuse 1000
    suppress 5000
  exit
  damping "TEST-MEDIUM"
    half-life 22
    max-suppress 720
    reuse 5000
    suppress 11000
  exit
=====
A:ALA-1#
```

The **show router policy *name* sub-policy entry *entry-id* name *subpol-name*** command displays information about the subroutine called by an entry of the route policy. In the following sample output, the term “telnet” is the value of a variable.

show router policy “BGP To RIP”

```
*A:Dut-B# show router policy "pol1" sub-policy entry 10 name "pol2"
=====
pol1 entry 10 sub-policies
=====
```

```

-----
Level1 Sub-policy : pol2      Parent-policy : pol1, Parent-entry 10
-----
    entry 20
        from
            community "telnet"
        exit
    exit
=====

```

policy-test

Syntax	policy-test <i>plcy-or-long-expr</i> [<i>plcy-or-expr</i> [<i>plcy-or-expr</i> .. (up to 14 max)]] [family <i>family</i>] [rd <i>rd</i>] [prefix <i>ip-prefix</i> / <i>ip-prefix-length</i>] [longer] [neighbor <i>ip-address</i>] [display-rejects] [{ detail brief }]
Context	show>router>bgp
Description	This command allows an operator to evaluate an existing policy or chain or policies (possibly including a logical expression) against received BGP routes to identify the routes accepted or rejected by the full evaluation of the policy, policy chain, or logical expression.
Parameters	<i>plcy-or-long-expr</i> — Specifies the name of a policy, up to 64 characters or a policy logical expression up to 255 characters. <i>plcy-or-expr</i> — Specifies the name of a policy, up to 64 characters or a policy logical expression up to 64 characters. <i>family</i> — Specifies the BGP address family. Values ipv4, ipv6, vpn-ipv4, vpn-ipv6, label-ipv4, label-ipv6 <i>rd</i> — Specifies the route distinguisher value in one of the following formats: <i>ip-addr:comm-val</i> <i>2byte-asnumber:ext-comm-val</i> <i>4byte-asnumber:comm-val</i> <i>ip-prefix/ip-prefix-length</i> — Specifies the IPv4 or IPv6 prefix to be evaluated. longer — Matches prefixes with a longer prefix-length. <i>ip-address</i> — Specifies the neighbor IP address. display-rejects — Displays routes rejected by the policy test as opposed to routes that are accepted. detail — Displays detailed information. brief — Displays abridged information.

7 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect (localised ID)*

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*
RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/
MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual
Private Networks (VPNs)*
RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge
Routers (6PE)*
RFC 4893, *BGP Support for Four-octet AS Number Space*
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*
RFC 5065, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5396, *Textual Representation of Autonomous System (AS) Numbers (asplain)*
RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6
Next Hop*
RFC 5575, *Dissemination of Flow Specification Rules*
RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 6996, *Autonomous System (AS) Reservation for Private Use*
RFC 7311, *The Accumulated IGP Metric Attribute for BGP*
RFC 7607, *Codification of AS 0 Processing*
RFC 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE)
Information Using BGP*
RFC 7854, *BGP Monitoring Protocol (BMP)*
RFC 7911, *Advertisement of Multiple Paths in BGP*
RFC 7999, *BLACKHOLE Community*
RFC 8092, *BGP Large Communities Attribute*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet
(SAToP)*
RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation
Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*

draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-04, Operational Aspects of Proxy-ARP/ND in EVPN Networks
draft-ietf-bess-evpn-vpls-seamless-integ-03, (PBB-)EVPN Seamless Integration with (PBB-)VPLS
draft-snr-bess-pbb-evpn-isid-cmacflush-01, PBB-EVPN ISID-based CMAC-Flush
RFC 7432, BGP MPLS-Based Ethernet VPN
RFC 7623, Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)
RFC 8214, Virtual Private Wire Service Support in Ethernet VPN
RFC 8317, Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)
RFC 8365, A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)

Frame Relay

ANSI T1.617 Annex D, DSS1 - Signalling Specification For Frame Relay Bearer Service
FRF.1.2, PVC User-to-Network Interface (UNI) Implementation Agreement
FRF.12, Frame Relay Fragmentation Implementation Agreement
FRF.16.1, Multilink Frame Relay UNI/NNI Implementation Agreement
FRF.5, Frame Relay/ATM PVC Network Interworking Implementation
FRF2.2, PVC Network-to-Network Interface (NNI) Implementation Agreement
ITU-T Q.933 Annex A, Additional procedures for Permanent Virtual Connection (PVC) status management

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information
RFC 3471, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description
RFC 3473, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions
RFC 4204, Link Management Protocol (LMP)
RFC 4208, Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model
RFC 4872, RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery
RFC 5063, Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

gRPC Network Management Interface (gNMI), Capabilities, Get, Set, Subscribe (ONCE, SAMPLE, ON_CHANGE)

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS (helper mode)*

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*
RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*
RFC 768, *User Datagram Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specifications*
RFC 1350, *The TFTP Protocol (revision 2)*
RFC 2347, *TFTP Option Extension*
RFC 2348, *TFTP Blocksize Option*
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*
RFC 2428, *FTP Extensions for IPv6 and NATs*
RFC 2784, *Generic Routing Encapsulation (GRE)*
RFC 2890, *Key and Sequence Number Extensions to GRE*
RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
RFC 4254, *The Secure Shell (SSH) Connection Protocol*
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
RFC 5925, *The TCP Authentication Option*
RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

- RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*
- RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*
- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*
- RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*
- RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*
- RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*
- RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*
- RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*
- RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*
- RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*
- RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks (MPLS encapsulation)*
- RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

- RFC 791, *Internet Protocol*
- RFC 792, *Internet Control Message Protocol*
- RFC 826, *An Ethernet Address Resolution Protocol*
- RFC 951, *Bootstrap Protocol (BOOTP)*
- RFC 1034, *Domain Names - Concepts and Facilities*
- RFC 1035, *Domain Names - Implementation and Specification*
- RFC 1191, *Path MTU Discovery (router specification)*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1534, *Interoperation between DHCP and BOOTP*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2003, *IP Encapsulation within IP*
RFC 2131, *Dynamic Host Configuration Protocol*
RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*
RFC 3046, *DHCP Relay Agent Information Option (Option 82)*
RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5722, *Handling of Overlapping IPv6 Fragments*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*
RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
RFC 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*
RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*
RFC 5903, *ECP Groups for IKE and IKEv2*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 6379, *Suite B Cryptographic Suites for IPsec*
RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*
RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*
RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*
RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, Controlling State Advertisements Of Non-negotiated LDP Applications
draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities
draft-pdutta-mpls-ldp-v2-00, LDP Version 2
draft-pdutta-mpls-mldp-up-redundancy-00, Upstream LSR Redundancy for Multipoint LDP Tunnels
draft-pdutta-mpls-multi-ldp-instance-00, Multiple LDP Instances
draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction
RFC 3037, LDP Applicability
RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (helper mode)
RFC 5036, LDP Specification
RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs)
RFC 5443, LDP IGP Synchronization
RFC 5561, LDP Capabilities
RFC 5919, Signaling LDP Label Advertisement Completion
RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
RFC 6512, Using Multipoint LDP When the Backbone Has No Route to the Root
RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
RFC 7032, LDP Downstream-on-Demand in Seamless MPLS
RFC 7552, Updates to LDP for IPv6

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions
RFC 2661, Layer Two Tunneling Protocol "L2TP"
RFC 2809, Implementation of L2TP Compulsory Tunneling via RADIUS
RFC 3438, Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update
RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3)
RFC 4719, Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)
RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module*

openconfig-isis.yang version 0.3.0, *IS-IS Module*
openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*
openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*
openconfig-lacp.yang version 1.1.0, *LACP Module*
openconfig-lldp.yang version 0.1.0, *LLDP Module*
openconfig-local-routing.yang version 1.0.1, *Local Routing Module*
openconfig-network-instance.yang version 0.8.0, *Network Instance Module*
openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*
openconfig-vlan.yang version 2.0.0, *VLAN Module*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1212, *Concise MIB Definitions*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1215, *A Convention for Defining Traps for use with the SNMP*
RFC 1724, *RIP Version 2 MIB Extension*
RFC 1901, *Introduction to Community-based SNMPv2*
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*
RFC 2206, *RSVP Management Information Base using SMIv2*
RFC 2213, *Integrated Services Management Information Base using SMIv2*
RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*
RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*
RFC 2515, *Definitions of Managed Objects for ATM Management*
RFC 2570, *SNMP Version 3 Framework*
RFC 2571, *An Architecture for Describing SNMP Management Frameworks*
RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
RFC 2573, *SNMP Applications*
RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*
RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
RFC 2579, *Textual Conventions for SMIv2*
RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*
RFC 4273, *Definitions of Managed Objects for BGP-4*
RFC 4292, *IP Forwarding Table MIB*
RFC 4293, *Management Information Base for the Internet Protocol (IP)*
RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
RFC 5102, *Information Model for IP Flow Information Export*
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2 (TLS client, RSA public key)*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*
RFC 6991, *Common YANG Data Types*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
RFC 7950, *The YANG 1.1 Data Modeling Language*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*
RFC 5921, *A Framework for MPLS in Transport Networks*
RFC 5960, *MPLS Transport Profile Data Plane Architecture*
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*
RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*
RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*
RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 5185, *OSPF Multi-Area Adjacency*
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*
RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*
RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*
RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*
RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*
RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*
RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*
RFC 6073, *Segmented Pseudowire*
RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
RFC 6718, *Pseudowire Redundancy*
RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
RFC 6870, *Pseudowire Preferential Forwarding Status bit*
RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing

draft-ietf-spring-conflict-resolution-05, Segment Routing MPLS Conflict Resolution

draft-ietf-spring-segment-routing-ldp-interop-09, Segment Routing interworking with LDP

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, Jitter Network Interfaces

ANSI T1.105.06, Physical Layer Specifications

ANSI T1.105.09, Network Timing and Synchronization

ITU-T G.703, Physical/electrical characteristics of hierarchical digital interfaces

ITU-T G.707, Network node interface for the synchronous digital hierarchy (SDH)

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC)

ITU-T G.823, The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy

ITU-T G.824, The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy

ITU-T G.825, The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)

ITU-T G.841, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002

ITU-T G.957, Optical interfaces for equipments and systems relating to the synchronous digital hierarchy

Time Division Multiplexing (TDM)

ANSI T1.403, DS1 Metallic Interface Specification

ANSI T1.404, DS3 Metallic Interface Specification

Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*, issued 09/2008

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*, issued 03/2003

ITU-T G.8261, *Timing and synchronization aspects in packet networks*, issued 04/2008

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*, issued 08/2007

ITU-T G.8264, *Distribution of timing information through packet networks*, issued 10/2008

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*, issued 10/2010

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*, issued 07/2014

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

