



Alcatel-Lucent 7450

ETHERNET SERVICE SWITCH | RELEASE 13.0.R1
ROUTING PROTOCOLS GUIDE

Alcatel-Lucent Proprietary
This document contains proprietary information of Alcatel-Lucent and is not to be disclosed or used except in
accordance with applicable agreements.
Copyright 2015 © Alcatel-Lucent. All rights reserved.

All specifications, procedures, and information in this document are subject to change and revision at any time without notice. The information contained herein is believed to be accurate as of the date of publication. Alcatel-Lucent provides no warranty, express or implied, regarding its contents. Users are fully responsible for application or use of the documentation.

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

Copyright 2015 Alcatel-Lucent.

All rights reserved.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

Table of Contents

Preface	17
About This Guide	17
Audience	17
List of Technical Publications	18
Technical Support	20
Getting Started	
In This Chapter	21
Alcatel-Lucent 7450 ESS-Series Router Configuration Process	21
Multicast	
In This Chapter	23
Introduction to Multicast	24
Multicast Models	25
Any-Source Multicast (ASM)	25
Source Specific Multicast (SSM)	25
Multicast in IP-VPN Networks	26
Multicast Features	27
Internet Group Management Protocol	27
IGMP Versions and Interoperability Requirements	28
IGMP Version Transition	28
Source-Specific Multicast Groups	28
Query Messages	29
Protocol Independent Multicast (PIM)	30
PIM-SM Functions	30
Encapsulating Data Packets in the Register Tunnel	33
PIM Bootstrap Router Mechanism	33
PIM-SM Routing Policies	33
Reverse Path Forwarding Checks	35
Anycast RP for PIM-SM	36
PIM Interface on IES Subscriber Group Interfaces	38
Multicast Only Fast Reroute (MoFRR)	40
Multicast Source Discovery Protocol (MSDP)	42
Anycast RP for MSDP	42
MSDP Procedure	42
MSDP Peer Groups	44
MSDP Mesh Groups	44
MSDP Routing Policies	44
Multicast in Virtual Private Networks	45
Dynamic Multicast Signaling over P2MP in GRT Instance	45
Multicast Extensions to MBGP	47
MBGP Multicast Topology Support	47
IPv6 Multicast	48
Multicast Listener Discovery (MLD v1 and v2)	48
PIM SSM	48

Table of Contents

IPv6 PIM ASM	49
Embedded RP	49
Multicast Connection Admission Control (MCAC)	50
MCAC Policy Overview	50
MCAC Algorithm	51
MCAC on Link Aggregation Group Interfaces	53
Multicast Debugging Tools	54
Mtrace	54
Mstat	56
Mrinfo	56
Configuring Multicast Parameters with CLI	57
Multicast Configuration Overview	58
Basic Configuration	59
Common Configuration Tasks	62
Configuring IGMP Parameters	62
Enabling IGMP	62
Configuring an IGMP Interface	64
Configuring Static Parameters	65
Configuring SSM Translation	67
Configuring PIM Parameters	68
Enabling PIM	68
Configuring PIM Interface Parameters	69
Importing PIM Join/Register Policies	74
Configuring Multicast Source Discovery Protocol (MSDP) Parameters	76
Configuring MCAC Parameters	77
Disabling IGMP or PIM	80
Multicast Command Reference	83
Command Hierarchies	83
Configuration Commands	99
Show Commands	165
Show Router PIM Commands	178
Clear Commands	229
Debug Commands	237

RIP

In This Chapter	247
RIP Overview	248
RIP Features	249
RIP Version Types	249
RIPv2 Authentication	249
RIP Packet Format	250
RIPng	252
RIPng Protocol	252
Common Attributes	253
Metrics	253
Timers	253
Import and Export Policies	254
Hierarchical Levels	254
RIP Configuration Process Overview	255

Configuration Notes	256
General	256
Configuring RIP with CLI	257
RIP & RIPng Configuration Overview	258
Preconfiguration Requirements	258
RIP Hierarchy	258
Basic RIP Configuration	259
Common Configuration Tasks	260
Configuring Interfaces	261
Configuring a Route Policy	262
Configuring RIP Parameters	264
Configuring Global-Level Parameters	266
Configuring Group-Level Parameters	267
Configuring Neighbor-Level Parameters	268
RIP Configuration Management Tasks	269
Modifying RIP Parameters	269
Deleting a Group	270
Deleting a Neighbor	270
RIP Command Reference	271
Command Hierarchies	271
RIP Configuration Commands	279
Show Commands	293
Clear Commands	304
Debug RIP Commands	305

OSPF

In This Chapter	307
Configuring OSPF	308
OSPF Areas	309
Backbone Area	309
Stub Area	310
Not-So-Stubby Area	311
OSPFv3 Graceful Restart Helper	317
Virtual Links	318
Neighbors and Adjacencies	319
Link-State Advertisements	320
Metrics	320
Authentication	321
IP Subnets	322
Preconfiguration Recommendations	322
Multiple OSPF Instances	323
Route Export Policies for OSPF	323
Preventing Route Redistribution Loops	324
Multi-Address Support for OSPFv3	325
IP Fast-reroute (IP FRR) For OSPF and IS-IS Prefixes	326
IP FRR Configuration	326
ECMP Considerations	327
IP FRR and RSVP Shortcut (IGP Shortcut)	328
IP FRR and BGP Next-Hop Resolution	328
OSPF and IS-IS Support for Loop-Free Alternate Calculation	328

Table of Contents

Loop-Free Alternate Shortest Path First (LFA SPF) Policies	335
Configuration of Route Next-Hop Policy Template	335
Configuring Affinity or Admin Group Constraint in Route Next-Hop Policy	335
Configuring SRLG Group Constraint in Route Next-Hop Policy	337
Interaction of IP and MPLS Admin Group and SRLG	338
Configuring Protection Type and Next-Hop Type Preference in Route next-hop policy template	339
Application of Route Next-Hop Policy Template to an Interface	339
Excluding Prefixes from LFA SPF	340
Modification to LFA Next-Hop Selection Algorithm	341
OSPF LSA Filtering	343
OSPF Configuration Process Overview	344
Configuration Notes	345
General	345
OSPF Defaults	345
Configuring OSPF with CLI	347
OSPF Configuration Guidelines	348
Basic OSPF Configuration	349
Configuring the Router ID	350
Configuring OSPF Components	351
Configuring OSPF Parameters	351
Configuring an OSPF Area	352
Configuring a Stub Area	353
Configuring a Not-So-Stubby Area	354
Configuring a Virtual Link	355
Configuring an Interface	357
Configuring Authentication	359
Overview	359
Configuring Authentication Keys and Algorithms	359
Configuring Authentication using Keychains	362
Assigning a Designated Router	363
Configuring Route Summaries	365
Configuring Route Preferences	367
OSPF Configuration Management Tasks	369
Modifying a Router ID	369
Deleting a Router ID	371
Modifying OSPF Parameters	372
OSPF Command Reference	375
Command Hierarchies	375
Configuration Commands	381
Show Commands	419
Clear Commands	467
OSPF Debug Commands	469

IS-IS

In This Chapter	475
Configuring IS-IS	476
Routing	477
IS-IS Frequently Used Terms	479
ISO Network Addressing	480

IS-IS PDU Configuration	481
IS-IS Operations	482
IS-IS Route Summarization	483
Partial SPF Calculation	483
IS-IS Administrative Tags	484
Setting Route Tags	484
Using Route Tags	485
Unnumbered Interface Support	485
Segment Routing in Shortest Path Forwarding	485
Configuring Segment Routing in Shortest Path	486
Segment Routing Operational Procedures	490
Segment Routing Tunnel Management	497
Remote LFA with Segment Routing	499
Data Path Support	502
IS-IS Control Protocol Changes	505
BGP Shortcut using Segment Routing Tunnel	508
BGP label Route Resolution using Segment Routing Tunnel	508
Service Packet Forwarding with Segment Routing	509
Mirror Services and Lawful Intercept	510
IS-IS Configuration Process Overview	512
Configuration Notes	513
General	513
Configuring IS-IS with CLI	515
IS-IS Configuration Overview	516
Router Levels	516
Area Address Attributes	516
Interface Level Capability	517
Route Leaking	518
Basic IS-IS Configuration	519
Common Configuration Tasks	521
Configuring IS-IS Components	522
Enabling IS-IS	522
Modifying Router-Level Parameters	522
Configuring ISO Area Addresses	524
Configuring Global IS-IS Parameters	525
Configuring Interface Parameters	526
Configuring IS-IS Link Groups	531
IS-IS Configuration Management Tasks	532
Disabling IS-IS	532
Removing IS-IS	532
Modifying Global IS-IS Parameters	533
Modifying IS-IS Interface Parameters	534
Configuring Authentication using Keychains	535
Configuring Leaking	536
Redistributing External IS-IS Routers	540
Specifying MAC Addresses for All IS-IS Routers	541
IS-IS Command Reference	543
Command Hierarchies	543
IS-IS Configuration Commands	549
Show Commands	595

Table of Contents

Clear Commands	635
Debug Commands	637

BGP

In This Chapter	641
BGP Overview	642
BGP Sessions	643
BGP Session States	644
Detecting BGP Session Failures	644
Peer Tracking	645
Bidirectional Forwarding Detection (BFD)	645
Fast External Failover	646
High Availability BGP Sessions	647
BGP Graceful Restart	647
BGP Session Security	649
TCP MD5 Authentication	649
TTL Security Mechanism	649
BGP Groups	650
BGP Design Concepts	651
Route Reflection	652
BGP Confederations	654
BGP Messages	656
Open Message	656
Changing the Autonomous System Number	657
Changing a Confederation Number	657
BGP Advertisement	657
Update Message	658
Keepalive Message	658
Notification Message	659
UPDATE Message Error Handling	659
Route Refresh Message	660
BGP Path Attributes	661
Origin	662
AS Path	662
AS Override	664
Using Local AS for ASN Migration	664
4-Octet Autonomous System Numbers	665
Next-Hop	666
Next-Hop IPv4 Address Family over IPv6	668
Next-Hop VPN-IPv4 Address Family over IPv6	669
Next-Hop VPN-IPv6 Address Family over IPv6	669
Next-Hop Resolution	670
Next-Hop Tracking	672
Next-Hop Indirection	672
Using Multiple Address Families over IPv6 BGP Sessions	672
MED	673
Deterministic MED	673
Local Preference	673
Route Aggregation Path Attributes	673
Community and Extended Community Attributes	675

Route Reflection Attributes	677
Multi-Protocol BGP Attributes	677
4-Octet AS Attributes	679
AIGP Metric	679
BGP Routing Information Base (RIB)	680
RIB-IN Features	680
BGP Import Policies	680
LOC-RIB Features	681
BGP Decision Process	681
BGP Route Installation in the Route Table	683
Weighted ECMP for BGP Routes	685
BGP Route Installation in the Tunnel Table	686
BGP Fast Reroute	686
QoS Policy Propagation via BGP (QPPB)	688
BGP Policy Accounting	689
Route Flap Damping (RFD)	690
RIB-OUT Features	692
BGP Export Policies	692
Outbound Route Filtering (ORF)	694
RT Constrained Route Distribution	695
Min Route Advertisement Interval (MRAI)	696
Advertise-Inactive	697
Best-External	698
Add-Paths	699
Split-Horizon	700
BGP Applications	701
Next-hop Resolution Using Tunnels	701
BGP Flow-Spec	705
Validating Received Flow Routes	707
Using Flow Routes to Create Dynamic Filter Entries	707
Configuration of TTL Propagation for BGP Label Routes	708
TTL Propagation for RFC 3107 Label Route at Ingress LER	708
TTL Propagation for RFC 3107 Label Routes at LSR	709
BGP Prefix Origin Validation	710
BGP Route Leaking	713
BGP Configuration Process Overview	715
Configuration Notes	716
General	716
BGP Defaults	716
BGP MIB Notes	717
Configuring BGP with CLI	719
BGP Configuration Overview	720
Preconfiguration Requirements	720
BGP Hierarchy	720
Internal and External BGP Configurations	720
Basic BGP Configuration	722
Common Configuration Tasks	724
Creating an Autonomous System	725
Configuring a Router ID	726
BGP Confederations	728

Table of Contents

BGP Route Reflectors	730
BGP Components	732
Configuring Group Attributes	732
Configuring Neighbor Attributes	733
Configuring Route Reflection	734
Configuring a Confederation	735
BGP Configuration Management Tasks	736
Modifying an AS Number	736
Modifying a Confederation Number	737
Modifying the BGP Router ID	737
Modifying the Router-Level Router ID	738
Deleting a Neighbor	739
Deleting Groups	740
BGP Command Reference	741
Command Hierarchies	741
Configuration Commands	751
Other BGP-Related Commands	799
Show Commands	805
Clear Commands	871
Debug Commands	873

Route Policies

In This Chapter	877
Configuring Route Policies	878
Policy Statements	879
Routing Policy Subroutines	880
Policy Evaluation Command	880
Exclusive Editing for Policy Configuration	880
Default Action Behavior	881
Denied IP Prefixes	881
Controlling Route Flapping	882
Regular Expressions	884
BGP and OSPF Route Policy Support	890
BGP Route Policies	891
Re-advised Route Policies	892
Triggered Policies	892
Set MED to IGP Cost using Route Policies	893
BGP Policy Subroutines	894
Route Policies for BGP Next-Hop Resolution and Peer Tracking	894
Routing Policy Parameterization	895
When to Use Route Policies	899
Route Policy Configuration Process Overview	900
Configuration Notes	901
General	901
Configuring Route Policies with CLI	903
Route Policy Configuration Overview	904
When to Create Routing Policies	904
Default Route Policy Actions	905
Policy Evaluation	906
Damping	909

Basic Configurations	910
Configuring Route Policy Components	912
Beginning the Policy Statement	913
Creating a Route Policy	914
Configuring a Default Action	915
Configuring an Entry	916
Configuring a Community List	917
Configuring Damping	918
Configuring a Prefix List	919
Configuring PIM Join/Register Policies	920
Configuring Bootstrap Message Import and Export Policies	922
Route Policy Configuration Management Tasks	923
Editing Policy Statements and Parameters	923
Deleting an Entry	925
Deleting a Policy Statement	925
Route Policy Command Reference	927
Command Hierarchies	927
Route Policy Command Reference	933
Show Commands	967
Standards and Protocol Support	983

List of Tables

Getting Started

Table 1: Configuration Process	21
--------------------------------------	----

Multicast

Table 2: Join Filter Policy Match Conditions	34
Table 3: Register Filter Policy Match Conditions	34

RIP

Table 4: Route Preference Defaults by Route Type	288
Table 5: RIP Neighbor Standard Output Fields	296

OSPF

Table 6: Route Preference Defaults by Route Type	367
Table 7: Route Preference Defaults by Route Type	388
Table 8: Route Preference Defaults by Route Type	394

IS-IS

Table 9: Data Path Support	503
Table 10: Potential Adjacency	517
Table 11: Potential Adjacency	571

BGP

Table 12: Multi-Protocol BGP support in SR-OS	678
Table 13: MED Comparison with always-compare-med	683
Table 14: BGP Fast Reroute Scenarios (Base Context)	687
Table 15: Subcomponents of IPv4 Flow Route NLRI	705
Table 16: Subcomponents of IPv6 Flow Route NLRI	706
Table 17: IPv4 Flowspec Actions	707
Table 18: SR OS and IETF MIB Variations	717
Table 19: MIB Variable with SNMP	717

Route Policies

Table 20: Regular Expression Operators	885
Table 21: Community Strings Examples	886
Table 22: AS Path and Community Regular Expression Examples	886
Table 23: Metric Set IGP Effect	893
Table 24: Default Route Policy Actions	905

List of Figures

Multicast

Figure 1: Anycast RP for PIM-SM Implementation Example	37
Figure 2: PIM Interface on IES Subscriber Group Interface	39
Figure 3: MoFRR Steady State No Failure	41
Figure 4: MoFRR Switch to Standby Stream on a Link Failure	41
Figure 5: Dynamic mLDP Signaling for IP Multicast in GRT	46

RIP

Figure 6: RIP Packet Format	250
Figure 7: RIPv1 Format	251
Figure 8: RIPv2 Format	251
Figure 9: RIP Configuration and Implementation Flow	255

OSPF

Figure 10: Backbone Area	310
Figure 11: PEs Connected to an MPLS-VPN Super Backbone	313
Figure 12: Sham Links	314
Figure 13: Example Topology with Primary and LFA Routes	329
Figure 14: Example Topology with Broadcast Interfaces	330
Figure 15: OSPF Configuration and Implementation Flow	344
Figure 16: OSPF Areas	417

IS-IS

Figure 17: IS-IS Routing Domain	476
Figure 18: Using Area Addresses to Form Adjacencies	481
Figure 19: Packet Label Encapsulation using Segment Routing Tunnel	487
Figure 20: Programming multiple tunnels to the same destination	492
Figure 21: Handling of Same Prefix and SID in different IS-IS Instances	495
Figure 22: Remote LFA Algorithm	500
Figure 23: Remote LFA Next-Hop in Segment Routing	502
Figure 24: Maximum Pushed Transport Label Stack in Shortest Path Forwarding with Segment Routing	504
Figure 25: IS-IS Configuration and Implementation Flow	512
Figure 26: Configuring a Level 1 Area	528
Figure 27: Configuring a Level 1/2 Area	530

BGP

Figure 28: BGP Sessions	643
Figure 29: Fully Meshed BGP Configuration	652
Figure 30: BGP Configuration with Route Reflectors	653
Figure 31: BGP Update Message with Path Identifier for IPv4 NLRI	699
Figure 32: BGP Configuration and Implementation Flow	715
Figure 33: Confederation Network Diagram Example	728
Figure 34: Route Reflection Network Diagram Example	730

List of Figures

Route Policies

Figure 35: BGP Route Policy Diagram	890
Figure 36: BGP Route Policy Diagram	891
Figure 37: OSPF Route Policy Diagram	892
Figure 38: Route Policy Past Mode of Operation	895
Figure 39: Route Policy Parameterization using sub-policies	896
Figure 40: Route Policy Configuration and Implementation Flow	900
Figure 41: Route Policy Process Example	907
Figure 42: Next Policy Logic Example	908
Figure 43: Damping Example	909

Preface

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 guide 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.

Audience

This guide is intended for network administrators who are responsible for configuring the 7450 ESS routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- Multicast — IGMP and PIM-SM
- Routing Reservation Protocol (RIP)
- Open Shortest Path First (OSPF)
- Intermediate System to Intermediate System (IS-IS)
- Border Gateway Protocol (BGP)
- Route policies

List of Technical Publications

The 7450 ESS documentation set is composed of the following guides:

- **7450 ESS Basic System Configuration Guide**
This guide describes basic system configurations and operations.
- **7450 ESS System Management Guide**
This guide describes system security and access configurations as well as event logging and accounting logs.
- **7450 ESS Interface Configuration Guide**
This guide describes card, Media Dependent Adapter (MDA) and port provisioning.
- **7450 ESS Router Configuration Guide**
This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.
- **7450 ESS Routing Protocols Guide**
This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
- **7450 ESS MPLS Guide**
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- **7450 ESS Services Overview Guide**
This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
- **7450 ESS Layer 2 Services and EVPN Guide**
This guide describes Virtual Leased Lines (VLL), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and Ethernet VPN (EVPN).
- **7450 ESS Layer 3 Services Guide**
This guide describes Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.
- **7450 ESS Versatile Service Module Guide**
This guide describes how to configure service parameters for the Versatile Service Module (VSM).
- **7450 ESS OAM and Diagnostics Guide**
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- **7450 ESS Triple Play Guide**
This guide describes Triple Play services and support provided by the 7450 ESS and presents examples to configure and implement various protocols and services.

- 7450 ESS Quality of Service Guide
This guide describes how to configure Quality of Service (QoS) policy management.
- Multi-Service Integrated Service Adapter Guide
This guide describes services provided by integrated service adapters such as Application Assurance, ad insertion (ADI) and Network Address Translation (NAT).

Technical Support

If you purchased a service agreement for your 7450 ESS router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, follow this link to contact an Alcatel-Lucent support representative and to access product manuals and documentation updates:

<http://support.alcatel-lucent.com>

Getting Started

In This Chapter

This chapter provides process flow information to configure IP routing protocols.

Alcatel-Lucent 7450 ESS-Series Router Configuration Process

[Table 1](#) lists the tasks necessary to configure RIP, OSPF, and IS-IS, BGP, and multicast protocols, and route policies. 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 1: Configuration Process

Area	Task	Chapter
Protocol configuration	Configure routing protocols: <ul style="list-style-type: none">• RIP• OSPF• IS-IS	RIP on page 247 OSPF on page 307 IS-IS on page 475
Policy configuration	<ul style="list-style-type: none">• Configure route policies	Route Policies on page 877
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 983

Note: In SR OS 12.0.R4 any function that displays an IPv6 address or prefix changes to reflect rules described in RFC 5952, *A Recommendation for IPv6 Address Text Representation*. Specifically, hexadecimal letters in IPv6 addresses are now represented in lowercase, and the correct compression of all leading zeros is displayed. This changes visible display output compared to previous SR OS releases. Previous SR OS behavior can cause issues with operator

scripts that use standard IPv6 address expressions and with libraries that have standard IPv6 parsing as per RFC 5952 rules. See the section on IPv6 Addresses in the Router Configuration Guide for more information.

In This Chapter

This chapter provides information about IPv6, Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM).

Topics in this chapter include:

- [Introduction to Multicast on page 24](#)
 - [Multicast Models on page 25](#)
 - [Multicast in IP-VPN Networks on page 26](#)
- [Multicast Features on page 27](#)
 - [Internet Group Management Protocol on page 27](#)
 - [Protocol Independent Multicast \(PIM\) on page 30](#)
 - [Multicast Source Discovery Protocol \(MSDP\) on page 42](#)
 - [Dynamic Multicast Signaling over P2MP in GRT Instance on page 45](#)
 - [Multicast Extensions to MBGP on page 47](#)
 - [IPv6 Multicast on page 48](#)
 - [Multicast Debugging Tools on page 54](#)
 - [Multicast Connection Admission Control \(MCAC\) on page 50](#)
 - [Multicast Debugging Tools on page 54](#)

Introduction to Multicast

IP multicast provides an effective method of many-to-many communication. Delivering unicast datagrams is fairly simple. Normally, IP packets are sent from a single source to a single recipient. The source inserts the address of the target host in the IP header destination field of an IP datagram, intermediate routers (if present) simply forward the datagram towards the target in accordance with their respective routing tables.

Sometimes distribution needs individual IP packets be delivered to multiple destinations (like audio or video streaming broadcasts). Multicast is a method of distributing datagrams sourced from one (or possibly more) host(s) to a set of receivers that may be distributed over different (sub) networks. This makes delivery of multicast datagrams significantly more complex.

Multicast sources can send a single copy of data using a single address for the entire group of recipients. The routers between the source and recipients route the data using the group address route. Multicast packets are delivered to a multicast group. A multicast group specifies a set of recipients who are interested in a particular data stream and is represented by an IP address from a specified range. Data addressed to the IP address is forwarded to the members of the group. A source host sends data to a multicast group by specifying the multicast group address in the datagram's destination IP address. A source does not have to register in order to send data to a group nor do they need to be a member of the group.

Routers and Layer 3 switches use the Internet Group Management Protocol (IGMP) to manage membership for a multicast session. When a host wants to receive one or more multicast sessions it will send a join message for each multicast group it wants to join. When a host wants to leave a multicast group, it will send a leave message.

To extend multicast to the Internet, the multicast backbone (Mbone) is used. The Mbone is layered on top of portions of the Internet. These portions, or islands, are interconnected using tunnels. The tunnels allow multicast traffic to pass between the multicast-capable portions of the Internet. As more and more routers in the Internet are multicast-capable (and scalable) the unicast and multicast routing table will converge.

The original Mbone was based on Distance Vector Multicast Routing Protocol (DVMRP) and was very limited. The Mbone is, however, converging around the following protocol set:

- IGMP
- Protocol Independent Multicast (Sparse Mode) (PIM-SM)
- Border Gateway Protocol with multi-protocol extensions (MBGP)
- Multicast Source Discovery Protocol (MSDP)

Multicast Models

Alcatel-Lucent routers support two models to provide multicast:

- [Any-Source Multicast \(ASM\) on page 25](#)
 - [Source Specific Multicast \(SSM\) on page 25](#)
 - [Multicast in IP-VPN Networks on page 26](#)
-

Any-Source Multicast (ASM)

Any-Source Multicast (ASM) is the IP multicast service model defined in RFC 1112, *Host extensions for IP Multicasting*. An IP datagram is transmitted to a host group, a set of zero or more end-hosts identified by a single IP destination address (224.0.0.0 through 239.255.255.255 for IPv4). End-hosts can join and leave the group any time and there is no restriction on their location or number. This model supports multicast groups with arbitrarily many senders. Any end-host can transmit to a host group even if it is not a member of that group.

To combat the vast complexity and scaling issues that ASM represents, the IETF is developing a service model called Source Specific Multicast (SSM).

Source Specific Multicast (SSM)

The Source Specific Multicast (SSM) service model defines a channel identified by an (S,G) pair, where S is a source address and G is an SSM destination address. In contrast to the ASM model, SSM only provides network-layer support for one-to-many delivery.

The SSM service model attempts to alleviate the following deployment problems that ASM has presented:

- Address allocation — SSM defines channels on a per-source basis. For example, the channel (S1,G) is distinct from the channel (S2,G), where S1 and S2 are source addresses, and G is an SSM destination address. This averts the problem of global allocation of SSM destination addresses and makes each source independently responsible for resolving address collisions for the various channels it creates.
- Access control — SSM provides an efficient solution to the access control problem. When a receiver subscribes to an (S,G) channel, it receives data sent only by the source S. In contrast, any host can transmit to an ASM host group. At the same time, when a sender picks a channel (S,G) to transmit on, it is automatically ensured that no other sender will be transmitting on the same channel (except in the case of malicious acts such as address spoofing). This makes it harder to spam an SSM channel than an ASM multicast group.

- Handling of well-known sources — SSM requires only source-based forwarding trees. This eliminates the need for a shared tree infrastructure. In terms of the IGMP, PIM-SM, MSDP, MBGP protocol suite, this implies that neither the RP-based shared tree infrastructure of PIM-SM nor the MSDP protocol is required. Thus, the complexity of the multicast routing infrastructure for SSM is low, making it viable for immediate deployment. Note that MBGP is still required for distribution of multicast reachability information.
 - Anticipating that point-to-multipoint applications such as Internet TV will be significant in the future, the SSM model is better suited for such applications.
-

Multicast in IP-VPN Networks

Multicast can be deployed as part of IP-VPN networks. For details on multicast support in IP-VPNs see SROS Services Guide.

Multicast Features

This section describes the multicast requirements when an Alcatel-Lucent router is deployed as part of the user's core network.

The required protocol set is as follows:

- Internet Group Management Protocol ([Internet Group Management Protocol on page 27](#))
 - Source Specific Multicast Groups ([SSM on page 28](#))
 - Protocol Independent Multicast (Sparse Mode) ([PIM-SM on page 30](#))
 - Multicast Extensions to MBGP ([Multicast Extensions to MBGP on page 47](#))
-

Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is used by IPv4 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

Multicast group memberships include at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router can assume one of two roles, querier or non-querier. There is normally only one querier per physical network.

A querier issues two types of queries, a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are issued when a router receives a leave message from the node it perceives as the last group member remaining on that network segment.

Hosts wanting to receive a multicast session issue a multicast group membership report. These reports must be sent to all multicast enabled routers.

IGMP Versions and Interoperability Requirements

If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported on their subnet and operate in that version.

Version 1 — Specified in RFC-1112, *Host extensions for IP Multicasting*, was the first widely deployed version and the first version to become an Internet standard.

Version 2 — Specified in RFC-2236, *Internet Group Management Protocol*, added support for “low leave latency”, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

Version 3 — Specified in RFC-3376, *Internet Group Management Protocol*, adds support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast (See Source Specific Multicast (SSM)), or from all but specific source addresses, sent to a particular multicast address.

IGMPv3 must keep state per group per attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network.

IGMP Version Transition

Alcatel-Lucent’s routers are capable of interoperating with routers and hosts running IGMPv1, IGMPv2, and/or IGMPv3. RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3)/ Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction* explores some of the interoperability issues and how they affect the various routing protocols.

IGMP version 3 specifies that if at any point a router receives an older version query message on an interface that it must immediately switch into a compatibility mode with that earlier version. Since none of the previous versions of IGMP are source aware, should this occur and the interface switch to Version 1 or 2 compatibility mode, any previously learned group memberships with specific sources (learned via the IGMPv3 specific INCLUDE or EXCLUDE mechanisms) MUST be converted to non-source specific group memberships. The routing protocol will then treat this as if there is no EXCLUDE definition present.

Source-Specific Multicast Groups

IGMPv3 permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic comes from a particular source. If a receiver does this, and no other receiver on

the LAN requires all the traffic for the group, then the designated router (DR) can omit performing a (*,G) join to set up the shared tree, and instead issue a source-specific (S,G) join only.

The range of multicast addresses from 232.0.0.0 to 232.255.255.255 is currently set aside for source-specific multicast in IPv4. For groups in this range, receivers should only issue source-specific IGMPv3 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

An Alcatel-Lucent router PIM router must silently ignore a received (*,G) PIM join message where G is a multicast group address from the multicast address group range that has been explicitly configured for SSM. This occurrence should generate an event. If configured, the IGMPv2 request can be translated into IGMPv3. The router allows for the conversion of an IGMPv2 (*,G) request into a IGMPv3 (S,G) request based on manual entries. A maximum of 32 SSM ranges is supported.

IGMPv3 also permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic does not come from a specific source or sources. In this case, the DR will perform a (*,G) join as normal, but can combine this with a prune for each of the sources the receiver does not wish to receive.

Query Messages

The IGMP query source address is configurable at two hierarchical levels. It can be configured globally at each router instance IGMP level and can be configured at individual at the group-interface level. The group-interface level overrides the src-ip address configured at the router instance level.

By default, subscribers with IGMP policies send IGMP queries with an all zero SRC IP address (0.0.0.0). However, some systems only accept and process IGMP query messages with non-zero SRC IP addresses. This feature allows the BNG to inter-operate with such systems.

Protocol Independent Multicast (PIM)

PIM-SM leverages the unicast routing protocols that are used to create the unicast routing table, OSPF, IS-IS, BGP, and static routes. Because PIM uses this unicast routing information to perform the multicast forwarding function it is effectively IP protocol independent. Unlike DVMRP, PIM does not send multicast routing tables updates to its neighbors.

PIM-SM uses the unicast routing table to perform the Reverse Path Forwarding (RPF) check function instead of building up a completely independent multicast routing table.

PIM-SM only forwards data to network segments with active receivers that have explicitly requested the multicast group. PIM-SM in the ASM model initially uses a shared tree to distribute information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or switch over to an optimized source distribution tree. As multicast traffic starts to flow down the shared tree, routers along the path determine if there is a better path to the source. If a more direct path exists, then the router closest to the receiver sends a join message toward the source and then reroutes the traffic along this path.

As stated above, PIM-SM relies on an underlying topology-gathering protocol to populate a routing table with routes. This routing table is called the Multicast Routing Information Base (MRIB). The routes in this table can be taken directly from the unicast routing table, or it can be different and provided by a separate routing protocol such as MBGP. Regardless of how it is created, the primary role of the MRIB in the PIM-SM protocol is to provide the next hop router along a multicast-capable path to each destination subnet. The MRIB is used to determine the next hop neighbor to whom any PIM join/prune message is sent. Data flows along the reverse path of the join messages. Thus, in contrast to the unicast RIB that specifies the next hop that a data packet would take to get to some subnet, the MRIB gives reverse-path information, and indicates the path that a multicast data packet would take from its origin subnet to the router that has the MRIB.

PIM-SM Functions

PIM-SM functions in three phases:

- [Phase One on page 31](#)
- [Phase Two on page 31](#)
- [Phase Three on page 32](#)

Phase One

In this phase, a multicast receiver expresses its interest in receiving traffic destined for a multicast group. Typically it does this using IGMP or MLD, but other mechanisms might also serve this purpose. One of the receiver's local routers is elected as the DR for that subnet. When the expression of interest is received, the DR sends a PIM join message towards the RP for that multicast group. This join message is known as a (*,G) join because it joins group G for all sources to that group. The (*,G) join travels hop-by-hop towards the RP for the group, and in each router it passes through the multicast tree state for group G is instantiated. Eventually the (*,G) join either reaches the RP or reaches a router that already has (*,G) join state for that group. When many receivers join the group, their join messages converge on the RP and form a distribution tree for group G that is rooted at the RP. This is known as the RP tree and is also known as the shared tree because it is shared by all sources sending to that group. Join messages are resent periodically as long as the receiver remains in the group. When all receivers on a leaf-network leave the group, the DR will send a PIM (*,G) prune message towards the RP for that multicast group. However if the prune message is not sent for any reason, the state will eventually time out.

A multicast data sender starts sending data destined for a multicast group. The sender's local router (the DR) takes those data packets, unicast-encapsulates them, and sends them directly to the RP. The RP receives these encapsulated data packets, removes the encapsulation, and forwards them onto the shared tree. The packets then follow the (*,G) multicast tree state in the routers on the RP tree, being replicated wherever the RP tree branches, and eventually reaching all the receivers for that multicast group. The process of encapsulating data packets to the RP is called registering, and the encapsulation packets are known as PIM register packets.

At the end of phase one, multicast traffic is flowing encapsulated to the RP, and then natively over the RP tree to the multicast receivers.

Phase Two

In this phase, register-encapsulation of data packets is performed. However, register-encapsulation of data packets is unsuitable for the following reasons:

- Encapsulation and de-encapsulation can be resource intensive operations for a router to perform depending on whether or not the router has appropriate hardware for the tasks.
- Traveling to the RP and then back down the shared tree can cause the packets to travel a relatively long distance to reach receivers that are close to the sender. For some applications, increased latency is unwanted.

Although register-encapsulation can continue indefinitely, for these reasons, the RP will normally switch to native forwarding. To do this, when the RP receives a register-encapsulated data packet from source S on group G, it will normally initiate an (S,G) source-specific join towards S. This join message travels hop-by-hop towards S, instantiating (S,G) multicast tree state in the routers along the path. (S,G) multicast tree state is used only to forward packets for group G if those

packets come from source S. Eventually the join message reaches S's subnet or a router that already has (S,G) multicast tree state, and then packets from S start to flow following the (S,G) tree state towards the RP. These data packets can also reach routers with (*,G) state along the path towards the RP - if so, they can short-cut onto the RP tree at this point.

While the RP is in the process of joining the source-specific tree for S, the data packets will continue being encapsulated to the RP. When packets from S also start to arrive natively at the RP, the RP will be receiving two copies of each of these packets. At this point, the RP starts to discard the encapsulated copy of these packets and it sends a register-stop message back to S's DR to prevent the DR unnecessarily encapsulating the packets. At the end of phase 2, traffic will be flowing natively from S along a source-specific tree to the RP and from there along the shared tree to the receivers. Where the two trees intersect, traffic can transfer from the shared RP tree to the shorter source tree.

Note that a sender can start sending before or after a receiver joins the group, and thus, phase two may occur before the shared tree to the receiver is built.

Phase Three

In this phase, the RP joins back towards the source using the shortest path tree. Although having the RP join back towards the source removes the encapsulation overhead, it does not completely optimize the forwarding paths. For many receivers the route via the RP can involve a significant detour when compared with the shortest path from the source to the receiver.

To obtain lower latencies, a router on the receiver's LAN, typically the DR, may optionally initiate a transfer from the shared tree to a source-specific shortest-path tree (SPT). To do this, it issues an (S,G) Join towards S. This instantiates state in the routers along the path to S. Eventually this join either reaches S's subnet or reaches a router that already has (S,G) state. When this happens, data packets from S start to flow following the (S,G) state until they reach the receiver.

At this point the receiver (or a router upstream of the receiver) will be receiving two copies of the data - one from the SPT and one from the RPT. When the first traffic starts to arrive from the SPT, the DR or upstream router starts to drop the packets for G from S that arrive via the RP tree. In addition, it sends an (S,G) prune message towards the RP. The prune message travels hop-by-hop instantiating state along the path towards the RP indicating that traffic from S for G should NOT be forwarded in this direction. The prune message is propagated until it reaches the RP or a router that still needs the traffic from S for other receivers.

By now, the receiver will be receiving traffic from S along the shortest-path tree between the receiver and S. In addition, the RP is receiving the traffic from S, but this traffic is no longer reaching the receiver along the RP tree. As far as the receiver is concerned, this is the final distribution tree.

Encapsulating Data Packets in the Register Tunnel

Conceptually, the register tunnel is an interface with a smaller MTU than the underlying IP interface towards the RP. IP fragmentation on packets forwarded on the register tunnel is performed based upon this smaller MTU. The encapsulating DR can perform path-MTU discovery to the RP to determine the effective MTU of the tunnel. This smaller MTU takes both the outer IP header and the PIM register header overhead into consideration.

PIM Bootstrap Router Mechanism

For proper operation, every PIM-SM router within a PIM domain must be able to map a particular global-scope multicast group address to the same RP. If this is not possible, then black holes can appear (this is where some receivers in the domain cannot receive some groups). A domain in this context is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary.

The bootstrap router (BSR) mechanism provides a way in which viable group-to-RP mappings can be created and distributed to all the PIM-SM routers in a domain. Each candidate BSR originates bootstrap messages (BSMs). Every BSM contains a BSR priority field. Routers within the domain flood the BSMs throughout the domain. A candidate BSR that hears about a higher-priority candidate BSR suppresses its sending of further BSMs for a period of time. The single remaining candidate BSR becomes the elected BSR and its BSMs inform the other routers in the domain that it is the elected BSR.

It is adaptive, meaning that if an RP becomes unreachable, it will be detected and the mapping tables will be modified so the unreachable RP is no longer used and the new tables will be rapidly distributed throughout the domain.

PIM-SM Routing Policies

Multicast traffic can be restricted from certain source addresses by creating routing policies. Join messages can be filtered using import filters. PIM join policies can be used to reduce denial of service attacks and subsequent PIM state explosion in the router and to remove unwanted multicast streams at the edge of the network before it is carried across the core. Route policies are created in the **config>router>policy-options** context. Join and register route policy match criteria for PIM-SM can specify the following:

- Router interface or interfaces specified by name or IP address.
- Neighbor address (the source address in the IP header of the join and prune message).
- Multicast group address embedded in the join and prune message.

- Multicast source address embedded in the join and prune message.

Join policies can be used to filter PIM join messages so no *,G or S,G state will be created on the router.

Table 2: Join Filter Policy Match Conditions

Match Condition	Matches the:
Interface	RTR interface by name
Neighbor	The neighbors source address in the IP header
Group Address	Multicast Group address in the join/prune message
Source Address	Source address in the join/prune message

PIM register message are sent by the first hop designated router that has a direct connection to the source. This serves a dual purpose:

- Notifies the RP that a source has active data for the group
- Delivers the multicast stream in register encapsulation to the RP and its potential receivers.
- If no one has joined the group at the RP, the RP will ignore the registers.

In an environment where the sources to particular multicast groups are always known, it is possible to apply register filters at the RP to prevent any unwanted sources from transmitting multicast stream. You can apply these filters at the edge so that register data does not travel unnecessarily over the network towards the RP.

Table 3: Register Filter Policy Match Conditions

Match Condition	Matches the:
Interface	RTR interface by name
Group Address	Multicast Group address in the join/prune message
Source Address	Source address in the join/prune message

Reverse Path Forwarding Checks

Multicast implements a reverse path forwarding check (RPF). RPF checks the path that multicast packets take between their sources and the destinations to prevent loops. Multicast requires that an incoming interface is the outgoing interface used by unicast routing to reach the source of the multicast packet. RPF forwards a multicast packet only if it is received on an interface that is used by the router to route to the source.

If the forwarding paths are modified due to routing topology changes then any dynamic filters that may have been applied must be re-evaluated. If filters are removed then the associated alarms are also cleared.

Anycast RP for PIM-SM

The implementation of Anycast RP for PIM-SM environments enable fast convergence when a PIM rendezvous point (RP) router fails by allowing receivers and sources to rendezvous at the closest RP. It allows an arbitrary number of RPs per group in a single shared-tree protocol Independent Multicast-Sparse Mode (PIM-SM) domain. This is, in particular, important for triple play configurations that opt to distribute multicast traffic using PIM-SM, not SSM. In this case, RP convergence must be fast enough to avoid the loss of multicast streams which could cause loss of TV delivery to the end customer.


Anycast RP for PIM-SM environments is supported in the base routing/PIM-SM instance of the service router. In the 7750 SR product lines, this feature is supported in Layer 3-VPRN instances that are configured with PIM.

Implementation

The Anycast RP for PIM-SM implementation is defined in *draft-ietf-pim-anycast-rp-03*, *Anycast-RP using PIM*, and is similar to that described in RFC 3446, *Anycast RP Mechanism Using PIM and MSDP*, and extends the register mechanism in PIM so Anycast RP functionality can be retained without using Multicast Source Discovery Protocol (MSDP) (see [Multicast in Virtual Private Networks on page 45](#)).

The mechanism works as follows:

- An IP address is chosen to use as the RP address. This address is statically configured, or distributed using a dynamic protocol, to all PIM routers throughout the domain.
- A set of routers in the domain are chosen to act as RPs for this RP address. These routers are called the Anycast-RP set.
- Each router in the Anycast-RP set is configured with a loopback interface using the RP address.
- Each router in the Anycast-RP set also needs a separate IP address to be used for communication between the RPs.
- The RP address, or a prefix that covers the RP address, is injected into the unicast routing system inside of the domain.
- Each router in the Anycast-RP set is configured with the addresses of all other routers in the Anycast-RP set. This must be consistently configured in all RPs in the set.



OSSG271

Figure 1: Anycast RP for PIM-SM Implementation Example

Assume the scenario in [Figure 1](#) is completely connected where R1A, R1B, and R2 are receivers for a group, and S1 and S2 send to that group. Assume RP1, RP2, and RP3 are all assigned the same IP address which is used as the Anycast-RP address (for example, the IP address is RPA).

Note, the address used for the RP address in the domain (the Anycast-RP address) must be different than the addresses used by the Anycast-RP routers to communicate with each other.

The following procedure is used when S1 starts sourcing traffic:

- S1 sends a multicast packet.
- The DR directly attached to S1 will form a PIM register message to send to the Anycast-RP address (RPA). The unicast routing system will deliver the PIM register message to the nearest RP, in this case RP1A.
- RP1 will receive the PIM register message, de-encapsulate it, send the packet down the shared-tree to get the packet to receivers R1A and R1B.
- RP1 is configured with RP2 and RP3's IP address. Since the register message did not come from one of the RPs in the anycast-RP set, RP1 assumes the packet came from a DR. If the register message is not addressed to the Anycast-RP address, an error has occurred and it should be rate-limited logged.
- RP1 will then send a copy of the register message from S1's DR to both RP2 and RP3. RP1 will use its own IP address as the source address for the PIM register message.
- RP1 may join back to the source-tree by triggering a (S1,G) Join message toward S1. However, RP1 must create (S1,G) state.
- RP2 receives the register message from RP1, de-encapsulates it, and also sends the packet down the shared-tree to get the packet to receiver R2.
- RP2 sends a register-stop message back to the RP1. RP2 may wait to send the register-stop message if it decides to join the source-tree. RP2 should wait until it has received data from the source on the source-tree before sending the register-stop message. If RP2

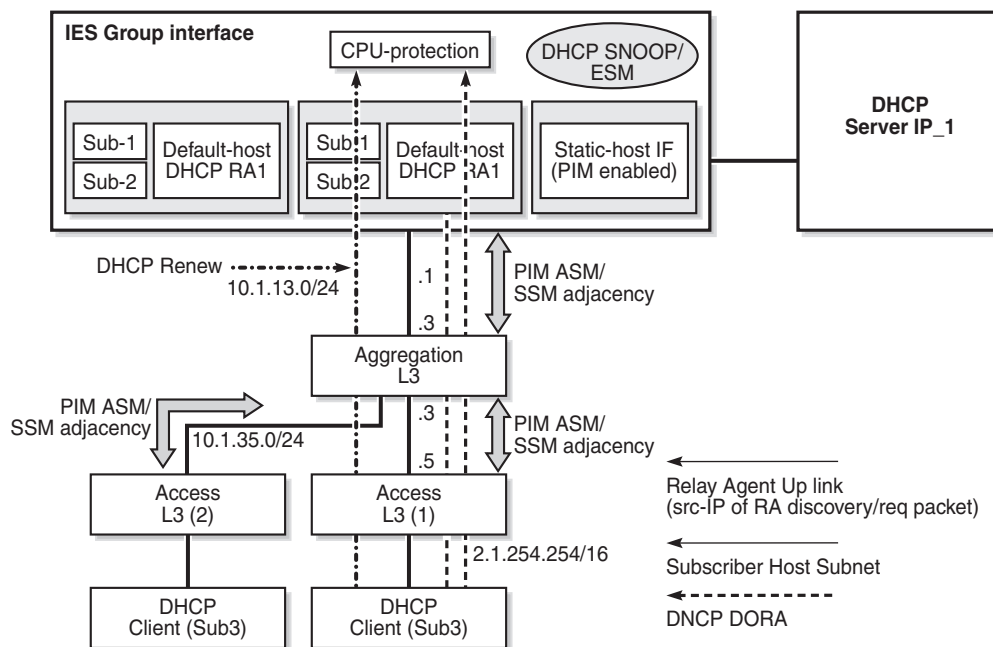
decides to wait, the register-stop message will be sent when the next register is received. If RP2 decides not to wait, the register-stop message is sent now.

- RP2 may join back to the source-tree by triggering a (S1,G) Join message toward S1. However, RP2 must create (S1,G) state.
- RP3 receives the register message from RP1, de-encapsulates it, but since there are no receivers joined for the group, it can discard the packet.
- RP3 sends a register-stop message back to the RP1.
- RP3 creates (S1,G) state so when a receiver joins after S1 starts sending, RP3 can join quickly to the source-tree for S1.
- RP1 processes the register-stop message from each of RP2 and RP3. RP1 may cache on a per-RP/per-(S,G) basis the receipt of register-stop message messages from the RPs in the anycast-RP set. This option is performed to increase the reliability of register message delivery to each RP. When this option is used, subsequent register messages received by RP1 are sent only to the RPs in the Anycast-RP set which have not previously sent register-stop message messages for the (S,G) entry.
- RP1 sends a register-stop message back to the DR the next time a register message is received from the DR and (when the option in the last bullet is in use) if all RPs in the Anycast-RP set have returned register-stop messages for a particular (S,G) route.

The procedure for S2 sending follows the same as above but it is RP3 which sends a copy of the register originated by S2's DR to RP1 and RP2. Therefore, this example shows how sources anywhere in the domain, associated with different RPs, can reach all receivers, also associated with different RPs, in the same domain.

PIM Interface on IES Subscriber Group Interfaces

PIM on a subscriber group interface allows for SAP-level replication over an ESM Group interface by establishing PIM adjacency to a downstream router. The following picture depicts the model:



24824

Figure 2: PIM Interface on IES Subscriber Group Interface

On an IES subscriber-interface, an Ethernet SAP is configured (LAG or physical port). On the SAP, a static-host is configured for connectivity to downstream Layer 3 aggregation devices (including PIM adjacency) while multiple default-hosts can be configured for subscriber traffic. Single SAP with a single static-host per group interface is supported to establish PIM adjacency on a given subscriber group interface. Both IPv4 PIM ASM and SSM are supported.

Feature caveats:

- Only IPv4 PIM is supported with a single static host used to form a PIM interface under a group interface. Using multiple hosts or non-static hosts is not supported. Configuring IPv6-related parameters in **configure>router>pim>interface group-ift** is not blocked, but takes no effect.
- **config>router>pim>apply-to** configuration does not apply to PIM interfaces on IES subscriber group interfaces.
- PIM on group interfaces is not supported in VPRN context.
- Extranet is not supported.
- Locally attached receivers are not supported (no IGMP/MLD and PIM mix in OIF list).

- Default anti-spoofing must be configured (IP+MAC).
 - A subscriber profile with pim-policy enabled cannot combine with the following policies (**config>subscr-mgmt>sub-prof**):
 - **[no] host-tracking** — Apply a host tracking policy
 - **[no] igmp-policy** — Apply an IGMP policy
 - **[no] mld-policy** — Apply an MLD policy
 - **[no] nat-policy** — Apply a NAT policy
 - **[no] sub-mcac-policy** — Apply a subscriber MCAC policy (MCAC policy can be used when configured in PIM interface context)
 - The feature is supported on IOM3-XP or newer line cards. When enabling the feature on older hardware, joins may be accepted and an outgoing interface may be created for the group, but traffic will not be sent out on egress because no OIF is created in forwarding.
-

Multicast Only Fast Reroute (MoFRR)

With large scale multicast deployments, a link or nodal failure impacts multiple subscribers or a complete region/segment of receivers. This failure interrupts the receiver client experience. Besides the impact on user experience, though multicast client applications may buffer streams for short period of time, the loss of stream data may trigger unicast request for the missing stream data to the source in certain middleware implementations. Those requests can overload the network resources, if a traffic loss persists for a prolonged period.

To minimize service interruption to end-users and protect the network from sudden surge of unicast requests, SROS implements a fast failover scheme for native IP networks. SROS MoFRR implementation is based on <http://tools.ietf.org/html/draft-karan-mofrr-02> and relies on:

- Sending a JOIN to a primary and a single standby upstream nodes over disjointed paths.
- Fast failover to a standby stream upon detection of a failure.

The functionality relies on failure detection on the primary path to switch to forwarding the traffic from the standby path. The traffic failure can happen with or without physical links or nodes going down. Various mechanisms for link/node failure detections are supported; however, to achieve best performance and resilience, it is recommended to enable MoFRR on every node in the network and use hop-by-hop BFD for fast link failure or data plane failure detection on each upstream link. Without BFD, the PIM adjacency loss or route change could be used to detect traffic failure. [Figure 3](#) and [Figure 4](#) depict MoFRR behavior.




Figure 3: MoFRR Steady State No Failure




Figure 4: MoFRR Switch to Standby Stream on a Link Failure

The MoFRR functionality on SROS routers supports the following:

- IPv4 link/node failure protection in global routing instance.
- Rosen PIM SSM with MDT SAFI
- Active and a single standby stream JOINS L3 over disjoint ECMP paths
- Active and a single standby stream JOINS over ISIS/OSPF Loop Free Alternative paths.

- When enabled, MoFRR is enabled on all regular PIM interfaces supporting MoFRR for all multicast streams. Tunnel interfaces are ignored.

Multicast Source Discovery Protocol (MSDP)

MSDP-speaking routers in a PIM-SM (RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*) domain have MSDP peering relationship with MSDP peers in another domain. The peering relationship is made up of a TCP connection in which control information is exchanged. Each domain has one or more connections to this virtual topology.

When a PIM-SM RP learns about a new multicast source within its own domain from a standard PIM register mechanism, it encapsulates the first data packet in an MSDP source-active message and sends it to all MSDP peers.

The source-active message is flooded (after an RPF check) by each peer to its MSDP peers until the source-active message reaches every MSDP router in the interconnected networks. If the receiving MSDP peer is an RP, and the RP has a (*.G) entry (receiver) for the group, the RP creates state for the source and joins to the shortest path tree for the source. The encapsulated data is de-encapsulated and forwarded down the shared tree of that RP. When the packet is received by the last hop router of the receiver, the last hop router also may join the shortest path tree to the source.

The MSDP speaker periodically sends source-active messages that include all sources.

Anycast RP for MSDP

MSDP is a mechanism that allows rendezvous points to share information about active sources. When RPs in remote domains hear about the active sources, they can pass on that information to the local receivers and multicast data can be forwarded between the domains. MSDP allows each domain to maintain an independent RP that does not rely on other domains but enables RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

Using PIM-SM, multicast sources and receivers register with their local RP by the closest multicast router. The RP maintains information about the sources and receivers for any particular group. RPs in other domains do not have any knowledge about sources located in other domains.

MSDP is required to provide inter-domain multicast services using Any Source Multicast (ASM). Anycast RP for MSDP enables fast convergence when should an MSDP/PIM PR router fail by allowing receivers and sources to rendezvous at the closest RP.

MSDP Procedure

When an RP in a PIM-SM domain first learns of a new sender, for example, by PIM register messages, it constructs a source-active (SA) message and sends it to its MSDP peers. The SA message contains the following fields:

- Source address of the data source
- Group address the data source sends to
- IP address of the RP

Note that an RP that is not a designated router on a shared network do not originate SAs for directly-connected sources on that shared network. It only originates in response to receiving register messages from the designated router.

Each MSDP peer receives and forwards the message away from the RP address in a peer-RPF flooding fashion. The notion of peer-RPF flooding is with respect to forwarding SA messages. The Multicast RPF Routing Information Base (MRIB) is examined to determine which peer towards the originating RP of the SA message is selected. Such a peer is called an RPF peer.

If the MSDP peer receives the SA from a non-RPF peer towards the originating RP, it will drop the message. Otherwise, it forwards the message to all its MSDP peers (except the one from which it received the SA message).

When an MSDP peer which is also an RP for its own domain receives a new SA message, it determines if there are any group members within the domain interested in any group described by an (S,G) entry within the SA message. That is, the RP checks for a (*,G) entry with a non-empty outgoing interface list. This implies that some system in the domain is interested in the group. In this case, the RP triggers an (S,G) join event toward the data source as if a join/prune message was received addressed to the RP. This sets up a branch of the source-tree to this domain. Subsequent data packets arrive at the RP by this tree branch and are forwarded down the shared-tree inside the domain. If leaf routers choose to join the source-tree they have the option to do so according to existing PIM-SM conventions. If an RP in a domain receives a PIM join message for a new group G, the RP must trigger an (S,G) join event for each active (S,G) for that group in its SA cache.

This procedure is called flood-and-join because if any RP is not interested in the group, the SA message can be ignored, otherwise, they join a distribution tree.

MSDP Peering Scenarios

Draft-ietf-mboned-msdp-deploy-nn.txt, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*, describes how protocols work together to provide intra- and inter-domain ASM service.

Inter-domain peering:

- Peering between PIM border routers (single-hop peering)
- Peering between non-border routers (multi-hop peering)

- MSDP peering without BGP
- MSDP peering between mesh groups
- MSDP peering at a multicast exchange

Intra-domain peering:

- Peering between routers configured for both MSDP and MBGP
 - MSDP peer is not BGP peer (meaning, no BGP peer)
-

MSDP Peer Groups

MSDP peer groups are typically created when multiple peers have a set of common operational parameters. Group parameters not specifically configured are inherited from the global level.

MSDP Mesh Groups

MSDP mesh groups are used to reduce source active flooding primarily in intra-domain configurations. When a number of speakers in an MSDP domain are fully meshed they can be configured as a mesh group. The originator of the source active message forwards the message to all members of the mesh group. Because of this, forwarding the SA between non-originating members of the mesh group is not necessary.

MSDP Routing Policies

MSDP routing policies allow for filtering of inbound and/or outbound active source messages. Policies can be configured at different levels:

- Global level — Applies to all peers
- Group level — Applies to all peers in peer-group
- Neighbor level — Applies only to specified peer

The most specific level is used. If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If no policy is applied source active messages are passed.

Match conditions include:

- Neighbor — Matches on a neighbor address is the source address in the IP header of the source active message.

- Route filter — Matches on a multicast group address embedded in the source active message
- Source address filter — Matches on a multicast source address embedded in the source active message

Multicast in Virtual Private Networks

Draft Rosen

RFC2547bis, *BGP/MPLS IP VPNs*, describes a method of providing a VPN service. A VPN provides secure connections to the network, allowing more efficient service to remote users without compromising the security of firewalls. The Rosen draft specifies the protocols and procedures which must be implemented in order for a service provider to provide a unicast VPN. The draft extends that specification by describing the protocols and procedures which a service provider must implement in order to support multicast traffic in a VPN, assuming that PIM [PIMv2] is the multicast routing protocol used within the VPN, and the SP network can provide PIM as well.

IGMP is not supported for receivers or senders directly attached to the PE.

For further information, refer to the Virtual Private Routed Network Service section of the Services Guide.

Dynamic Multicast Signaling over P2MP in GRT Instance

This feature provides a flexible multicast signaling solution to connect native IP multicast source and receivers running PIM multicast protocol via an intermediate MPLS (P2MP LDP LSP) network. The feature allows each native IP multicast flow to be connected via an intermediate P2MP LSP by dynamically mapping each PIM multicast flow to a P2MP LDP LSP.

The feature uses procedures defined in RFC 6826: Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths. On the leaf node of a P2MP LSP, PIM signaling is dynamically mapped to P2MP LDP tree setup. On the root node of P2MP LSP, P2MP LDP signaling is handed back to PIM. Due to dynamic mapping of multicast IP flow to P2MP LSP, provisioning and maintenance overhead is eliminated as multicast distribution services are added and removed from the network. Per (S, G) IP multicast state is also removed from the network where P2MP LSPs are used to transport multicast flows.

Figure 5 illustrates dynamic mLDP signaling for IP multicast in GRT.

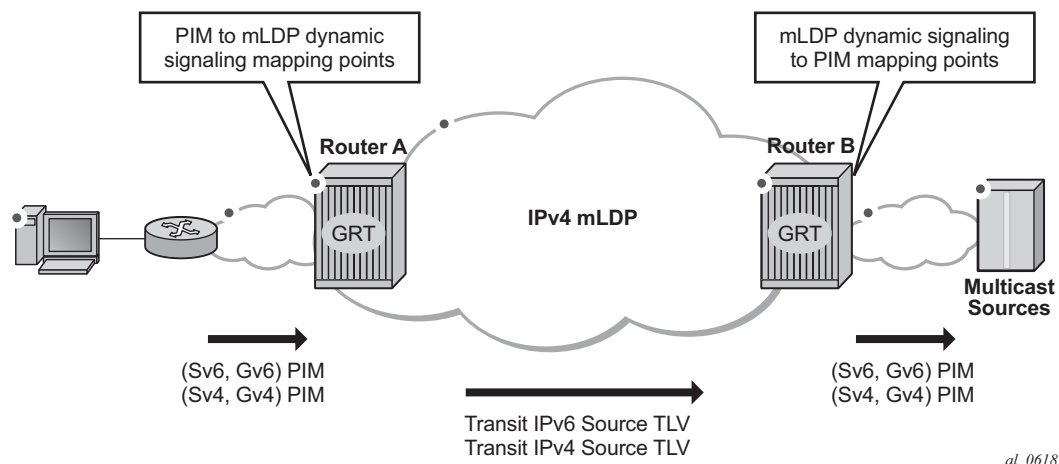


Figure 5: Dynamic mLDP Signaling for IP Multicast in GRT

As illustrated in Figure 5, P2MP LDP LSP signaling is initiated from the router that receives PIM JOIN from a downstream router (Router A). To enable dynamic multicast signaling, p2mp-ldp-tree-join must be configured on PIM outgoing interface of Router A. This enables handover of multicast tree signaling from PIM to P2MP LDP LSP. Being a leaf node of P2MP LDP LSP, Router A selects the upstream-hop as the root node of P2MP LDP FEC based on routing table lookup. If an ECMP path is available for a given route, then the number of trees are equally balanced towards multiple root nodes. The PIM Joins are carried in Transit IPv4 (IPv4 PIM SSM) or IPv6 (IPv6 PIM SSM) mLDP TLVs. On the root node of P2MP LDP LSP (Router B), multicast tree signaling is handed back to PIM and propagated upstream as native-IP PIM JOIN.

The feature is supported with IPv4 and IPv6 PIM SSM and IPv4 mLDP. Directly connected IGMP/MLD receivers are also supported with PIM enabled on outgoing interfaces and SSM mapping configured if required.

If multiple criteria exist to setup a multicast flow, the following priority is given:

1. Multicast (statically provisioned) over P2MP LSP (RSVP-TE or LDP)
2. Dynamic multicast signaling over P2MP LDP
3. PIM native-IP multicast

The following are feature caveats:

- A single instance of P2MP LDP LSP is supported between the root and leaf nodes per multicast flow; there is no stitching of dynamic trees.
- Extranet functionality is not supported.

- The router LSA link ID or the advertising router ID must be a routable IPv4 address (including IPv6 into IPv4 mLDP use cases).
 - IPv6 PIM with dynamic IPv4 mLDP signaling is not supported with e-BGP or i-BGP with IPv6 next-hop.
 - Inter-AS and IGP inter-area scenarios where the originating router is altered at the ASBR and ABR respectively, (hence PIM has no way to create the LDP LSP towards the source), are not supported.
 - The feature requires chassis mode C.
-

Multicast Extensions to MBGP

This section describes the implementation of extensions to MBGP to support multicast. Rather than assuming that all unicast routes are multicast-capable, some routed environments, in some cases, some ISPs do not support or have limited support for multicast throughout their AS.

BGP is capable of supporting two sets of routing information, one set for unicast routing and the other for multicast routing. The unicast and multicast routing sets either partially or fully overlay one another. To achieve this, BGP has added support for IPv4 and mcast-IPv4 address families. Routing policies can be imported or exported.

The multicast routing information can subsequently be used by the Protocol Independent Multicast (PIM) protocol to perform its Reverse Path Forwarding (RPF) lookups for multicast-capable sources. Thus, multicast traffic can only be routed across a multicast topology and not a unicast topology.

MBGP Multicast Topology Support

Recursive Lookup for BGP Next Hops

The next hop for multicast RPF routes learned by MBGP is not always the address of a directly-connected neighbor. For unicast routing, a router resolves the directly-connected next-hop by repeating the IGP routes. For multicast RPF routes, there are different ways to find the real next-hops.

- Scanning to see if a route encompasses the BGP next hop. If one exists, this route is used. If not, the tables are scanned for the best matching route.
- Check to see if the recursed next hop is taken from the protocol routing table with the lowest administrative distance (protocol preference). This means that the operating system algorithm must perform multiple lookups in the order of the lowest admin distance. Note that unlike recursion on the unicast routing table, the longest prefix match rule does not

take effect; protocol preference is considered prior to prefix length. For example, the route 12.0.0.0/14 learned via MBGP will be selected over the route 12.0.0.0/16 learned via BGP.

IPv6 Multicast

IPv6 multicast enables multicast applications over native IPv6 networks. There are two service models: Any Source Multicast (ASM) and Source Specific Multicast (SSM) which includes PIM SSM and MLD (v1 and v2). SSM does not require source discovery and only supports single source for a specific multicast stream. As a result, SSM is easier to operate in a large scale deployment that uses the one-to-many service model.

Multicast Listener Discovery (MLD v1 and v2)

MLD is the IPv6 version of IGMP. The purpose of MLD is to allow each IPv6 router to discover the presence of multicast listeners on its directly attached links, and to discover specifically which multicast groups are of interest to those neighboring nodes.

MLD is a sub-protocol of ICMPv6. MLD message types are a subset of the set of ICMPv6 messages, and MLD messages are identified in IPv6 packets by a preceding Next Header value of 58. All MLD messages are sent with a link-local IPv6 source address, a Hop Limit of 1, and an IPv6 Router Alert option in the Hop-by-Hop Options header.

Similar to IGMPv2, MLDv1 reports only include the multicast group addresses that listeners are interested in, and don't include the source addresses. In order to work with PIM SSM model, a similar SSM translation function is required when MLDv1 is used.

SSM translation allows an IGMPv2 device to join an SSM multicast network through the router that provides such a translation capability. Currently SSM translation can be done at a box level, but this does not allow a per-interface translation to be specified. SSM translation per interface offers the ability to have a same (*,G) mapped to two different (S,G) on two different interfaces to provide flexibility.

MLDv2 is backward compatible with MLDv1 and adds the ability for a node to report interest in listening to packets with a particular multicast group only from specific source addresses or from all sources except for specific source addresses.

PIM SSM

The IPv6 address family for SSM model is supported. This includes the ability to choose which RTM table to use (unicast RTM, multicast RTM, or both). OSPF3, IS-IS and static-route have extensions to support submission of routes into the IPv6 multicast RTM.

IPv6 PIM ASM

IPv6 PIM ASM is supported. All PIM ASM related functions such as bootstrap router, RP, etc., support both IPv4 and IPv6 address-families. IPv6 specific parameters are configured under **configure>router>pim>rp>ipv6**.

Embedded RP

The detailed protocol specification is defined in RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*. This RFC describes a multicast address allocation policy in which the address of the RP is encoded in the IPv6 multicast group address, and specifies a PIM-SM group-to-RP mapping to use the encoding, leveraging, and extending unicast-prefix-based addressing. This mechanism not only provides a simple solution for IPv6 inter-domain ASM but can be used as a simple solution for IPv6 intra-domain ASM with scoped multicast addresses as well. It can also be used as an automatic RP discovery mechanism in those deployment scenarios that would have previously used the Bootstrap Router protocol (BSR).

Multicast Connection Admission Control (MCAC)

Multicast Connection Admission Control (MCAC) allows a router to limit bandwidth used by multicast channels, either on a router or on access links, by controlling the number of channels that are accepted. When a pre-configured limit is reached, the router prevents receivers from joining any new channels not currently established. As result, running the MCAC function might cause some channels to be temporarily unavailable to receivers under overload. However, by rejecting new channel establishment during an overload condition, the degradation of the quality of the existing multicast service offering is avoided.

Operators can configure one or more MCAC policies (`configure>router>mcac`) to specify multicast channel admission rules and then reference a required policy on multicast-enabled IPv4 and IPv6 interfaces or group-interfaces. In addition operators can configure per-interface MCAC behavior.

Multicast CAC is supported on ESM subscriber interfaces as well as multicast interfaces in base router instance and in MVPNs. MCAC is supported for IGMP, IGMP-snooping, MLD, and PIM. When a MCAC policy is applied to a split horizon group, then member SAPs do not permit policy enforcement configurations.

Feature caveats:

- MCAC is not supported with PIM snooping and MLD snooping
-

MCAC Policy Overview

MCAC policy is used to define MCAC rules to be applied on an interface when receivers are trying to join multicast channels. Within each policy, an operator can define:

- Multicast channel:
 - A channel can be defined using multicast group address only or both source and group addresses. Ranges can be used to group multiple multicast channels into a single MCAC channel. When ranges are used, each multicast channel within range will use the same channel BW, class, and priority configuration.
 - Channel BW: a bandwidth value to be used for a channel in MCAC.
 - Channel type (mandatory or optional): mandatory channels have BW pre-reserved on interfaces as soon as they are defined in MCAC policy, while optional channels consume BW on-demand; only when there are active receivers for that channel and the remaining BW allows for channels to be admitted.

→ Channel class: two classes are supported: high and low. For LAG interfaces, the class parameter allows further prioritizing of the mandatory or optional channels. This brings the number of priority levels to four during reshuffles of the joined channels when LAG ports are changing state.

NOTE: Multicast channels not specified in an MCAC policy applicable on a given interface are not subject to MCAC. Treatment of such channels is configurable as either accept or discard.

- Multicast channel bundle:
 - Multicast bundle defines multicast channels as described above. A channel can only be part of one bundle.
 - Maximum bundle BW – the maximum bandwidth the channels forming a given bundle can consume on an interface.
 - MCAC constraints – set of rules governing available BW for multicast channels over LAG as LAG ports are changing state.

MCAC Algorithm

It is important to point out that the MCAC algorithm is based on configured BW values. The configured channel BW based on MCAC policy is CAC-ed against pre-configured maximum bundle BW and pre-configured interface multicast BW limits. A channel must pass all levels of CAC before it is accepted. The statements outline the CAC algorithm for a multicast channel defined in MCAC policy:

A join for a particular multicast channel is accepted if:

1. Mandatory channels:

A sufficient bandwidth exists on the interface according to the policy settings for the interface (Interface-level MCAC) and BW setting for a channel (Bundle-level MCAC). Note, there is always sufficient BW available on the bundle level, because mandatory channels get pre-reserved bandwidth.

2. Optional channels:

A sufficient BW exists on both interface (Interface-level MCAC) and bundle level (Bundle-Level MCAC) based on channel configured BW and currently available BW on both interface and bundle.

When a policy is evaluated over a set of existing channels (adding policy MCAC on LAG), the channels are evaluated and admitted/dropped based on the following priority order: mandatory-high, mandatory-low, optional-high, optional low.

Multicast Connection Admission Control (MCAC)

This method does not guarantee that all bundles are fully allocated while others are not. However it does ensure that all mandatory high channels are allocated before any mandatory lows ones are allocated.

Interface-level MCAC details

Interface-level MCAC constraints are applied to the interface on which the join was received. The channel is allowed when:

- If it is defined as mandatory and the bandwidth for the already accepted mandatory channels plus the bandwidth of this channel is not greater than the configured mandatory bandwidth on this interface.
 - If it is defined optional and the bandwidth for the already accepted optional channels plus the bandwidth of this channel is not greater than the configured amount of unconstrained bandwidth less the configured amount of mandatory bandwidth on this interface.
-

Bundle-Level MCAC details

Bundle-level CAC is applied to the bundle to which the channel belongs that triggered the MCAC algorithm. The channel is allowed when:

- If it is defined as mandatory – always.
 - If it is defined as optional, then the allocated bundle bandwidth cannot exceed the configured bandwidth. The allocated bandwidth equals the bandwidth of all the mandatory channels belonging to that bundle plus the bandwidth of the optional channels already accepted plus the bandwidth of this optional channel.
-

MCAC on Link Aggregation Group Interfaces

When MCAC enabled interfaces reside on a LAG, SROS allows operators to change MCAC behavior when the number of active ports in a LAG changes. Both MCAC policy bundle and MCAC interface allows operators to define multiple MCAC levels per LAG based on the number of active ports in the LAG. For each level, operators can configure corresponding BW limits.

When MCAC LAG constraints are enabled, the level to use is selected automatically based on the configuration and a currently active number of LAG ports. In a case of the available bandwidth reduction (for example, a LAG link failure causes change to a level with smaller BW configured), MCAC attempts first to fit all mandatory channels (in an arbitrary order). If there is no sufficient capacity to carry all mandatory channels in the degraded mode, some channels are dropped and all optional channels are dropped. If after evaluation of mandatory channels, there remains available bandwidth, then all optional channels are re-evaluated (in an arbitrary order). Channel re-evaluation employs the above-described MCAC algorithm applied at the interface and bundle levels that use the constraints for the degraded mode of operation.

Multicast Debugging Tools

This section describes multicast debugging tools requirement for the router family of products.

The debugging tools for multicast consist out of three elements; mtrace, mstat, and mrinfo.

Mtrace

Assessing problems in the distribution of IP multicast traffic can be difficult. The **mtrace** feature utilizes a tracing feature implemented in multicast routers that is accessed via an extension to the IGMP protocol. The **mtrace** feature is used to print the path from the source to a receiver; it does this by passing a trace query hop-by-hop along the reverse path from the receiver to the source. At each hop, information such as the hop address, routing error conditions and packet statistics should be gathered and returned to the requestor.

Data added by each hop includes:

- Query arrival time
- Incoming interface
- Outgoing interface
- Previous hop router address
- Input packet count
- Output packet count
- Total packets for this source/group
- Routing protocol
- TTL threshold
- Forwarding/error code

The information enables the network administrator to determine:

- Where multicast flows stop
- the flow of the multicast stream

When the trace response packet reaches the first hop router (the router that is directly connected to the source's net), that router sends the completed response to the response destination (receiver) address specified in the trace query.

If some multicast router along the path does not implement the multicast traceroute feature or if there is some outage, then no response is returned. To solve this problem, the trace query includes

a maximum hop count field to limit the number of hops traced before the response is returned. This allows a partial path to be traced.

The reports inserted by each router contain not only the address of the hop, but also the TTL required to forward and some flags to indicate routing errors, plus counts of the total number of packets on the incoming and outgoing interfaces and those forwarded for the specified group. Taking differences in these counts for two traces separated in time and comparing the output packet counts from one hop with the input packet counts of the next hop allows the calculation of packet rate and packet loss statistics for each hop to isolate congestion problems.

Finding the Last Hop Router

The trace query must be sent to the multicast router which is the last hop on the path from the source to the receiver. If the receiver is on the local subnet (as determined using the subnet mask), then the default method is to multicast the trace query to all-routers.mcast.net (224.0.0.2) with a TTL of 1. Otherwise, the trace query is multicast to the group address since the last hop router will be a member of that group if the receiver is. Therefore, it is necessary to specify a group that the intended receiver has joined. This multicast is sent with a default TTL of 64, which may not be sufficient for all cases.

When tracing from a multihomed host or router, the default receiver address may not be the desired interface for the path from the source. In that case, the desired interface should be specified explicitly as the receiver.

Directing the Response

By default, mtrace first attempts to trace the full reverse path, unless the number of hops to trace is explicitly set with the hop option. If there is no response within a 3 second timeout interval, a "*" is printed and the probing switches to hop-by-hop mode. Trace queries are issued starting with a maximum hop count of one and increasing by one until the full path is traced or no response is received. At each hop, multiple probes are sent. The first attempt is made with the unicast address of the host running mtrace as the destination for the response. Since the unicast route may be blocked, the remainder of attempts request that the response be multicast to mtrace.mcast.net (224.0.1.32) with the TTL set to 32 more than what's needed to pass the thresholds seen so far along the path to the receiver. For the last attempts the TTL is increased by another 32.

Alternatively, the TTL may be set explicitly with the TTL option.

For each attempt, if no response is received within the timeout, a "*" is printed. After the specified number of attempts have failed, mtrace will try to query the next hop router with a DVMRP_ASK_NEIGHBORS2 request (as used by the mrinfo program) to determine the router type.

The output of `mtrace` is a short listing of the hops in the order they are queried, that is, in the reverse of the order from the source to the receiver. For each hop, a line is printed showing the hop number (counted negatively to indicate that this is the reverse path); the multicast routing protocol; the threshold required to forward data (to the previous hop in the listing as indicated by the up-arrow character); and the cumulative delay for the query to reach that hop (valid only if the clocks are synchronized). The response ends with a line showing the round-trip time which measures the interval from when the query is issued until the response is received, both derived from the local system clock.

Mtrace/mstat packets use special IGMP packets with IGMP type codes of 0x1E and 0x1F.

Mstat

The **mstat** command adds the capability to show the multicast path in a limited graphic display and provide drops, duplicates, TTLs and delays at each node. This information is useful to the network operator because it identifies nodes with high drop & duplicate counts. Duplicate counts are shown as negative drops.

The output of **mstat** provides a limited pictorial view of the path in the forward direction with data flow indicated by arrows pointing downward and the query path indicated by arrows pointing upward. For each hop, both the entry and exit addresses of the router are shown if different, along with the initial ttl required on the packet in order to be forwarded at this hop and the propagation delay across the hop assuming that the routers at both ends have synchronized clocks. The output consists of two columns, one for the overall multicast packet rate that does not contain lost/sent packets and a column for the (S,G)-specific case. The S,G statistics do not contain lost/sent packets.

Mrinfo

mrinfo is a simple mechanism based on the **ask_neighbors igmp** to display the configuration information from the target multicast router. The type of information displayed includes the Multicast of the router, code version, metrics, ttl-thresholds, protocols and status. This information, for instance, can be used by network operators to verify if bi-directional adjacencies exist. Once the specified multicast router responds, the configuration is displayed.

Configuring Multicast Parameters with CLI

This section provides information to configure multicast, IGMP, and PIM.

Topics in this section include:

- [Multicast Configuration Overview on page 58](#)
- [Basic Configuration on page 59](#)
- [Common Configuration Tasks on page 62](#)
- [Disabling IGMP or PIM on page 80](#)

Multicast Configuration Overview

The routers use IGMP to manage membership for a given multicast session. IGMP is not enabled by default. When enabled, at least one interface must be specified in the IGMP context as IGMP is an interface function. Creating an interface enables IGMP. Traffic can only flow away from the router to an IGMP interface and to and from a PIM interface. A router directly connected to a source must have PIM enabled on the interface to that source. The traffic travels in a network from PIM interface to PIM interface and arrives finally on an IGMP enabled interface.

The IGMP CLI context allows you to specify an existing IP interface and modify the interface-specific parameters. Static IGMP group memberships can be configured to test multicast forwarding without a receiver host. When IGMP static group membership is enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP. When a host wants to receive multicast sessions it sends a join message for each multicast group it wants to join. Then, a leave message may be sent for each multicast group it no longer wishes to participate with.

A multicast router keeps a list of multicast group memberships for each attached network, and an interval timer for each membership. Hosts issue a Multicast Group Membership Report when they want to receive a multicast session. The reports are sent to all multicast routers.

PIM is not enabled by default. When PIM is enabled, data is forwarded to network segments with active receivers that have explicitly requested the multicast group. When enabled, at least one interface must be specified in the PIM context as PIM is an interface function. Creating an interface enables PIM.

Basic Configuration

Perform the following basic multicast configuration tasks:

For IGMP:

- Enable IGMP (required)
- Configure IGMP interfaces (required)
- Specify IGMP version on the interface (optional)
- Configure static (S,G)/(*,G) (optional)
- Configure SSM translation (optional)

For PIM:

- Enable PIM (required)
- Add interfaces so the protocol establishes adjacencies with the neighboring routers (required)
- Configure a way to calculate group-to-RP mapping (required) by either:
 - Static group-to-RP mapping
 - Enable Candidate RP/Bootstrap mechanism on some routers.
- Enable unicast routing protocols to learn routes towards the RP/source for reverse path forwarding (required)
- Add SSM ranges (optional)
- Enable Candidate BSR (optional)
- Enable Candidate RP (optional)
- Change hello interval (optional)
- Configure route policies (bootstrap-export, bootstrap-import, import join and register)

For MSDP:

- Enable MSDP (required)
- Configure peer
- Configure local address

For MCAC:

- Configure policy name
- Configure bundle parameters
- Specify default action

The following example displays the enabled IGMP and PIM configurations:

```
A:LAX>config>router>igmp# info
-----
interface "lax-vls"
exit
interface "pl-ix"
exit
-----
A:LAX>config>router>igmp# info detail
-----
interface "lax-vls"
  no import
  version 3
  no shutdown
exit
interface "pl-ix"
  no import
  version 3
  no shutdown
exit
query-interval 125
query-last-member-interval 1
query-response-interval 10
robust-count 2
no shutdown
-----
A:LAX>config>router>igmp# exit
A:LAX>config>router# pim
A:LAX>config>router>pim# info
-----
interface "system"
exit
interface "lax-vls"
exit
interface "lax-sjc"
exit
interface "pl-ix"
exit
rp
  static
    address 2.22.187.237
    group-prefix 224.24.24.24/32
  exit
exit
bsr-candidate
  shutdown
exit
rp-candidate
  shutdown
exit
exit
-----
A:LAX>config>router>pim# info detail
-----
no import join-policy
no import register-policy
interface "system"
```



```

        priority 1
        hello-interval 30
        multicast-senders auto
        no tracking-support
        bsm-check-rtr-alert
        no shutdown
    exit
interface "lax-vls"
    priority 1
    hello-interval 30
    multicast-senders auto
    no tracking-support
    bsm-check-rtr-alert
    no shutdown
exit
interface "lax-sjc"
    priority 1
    hello-interval 30
    multicast-senders auto
    no tracking-support
    bsm-check-rtr-alert
    no shutdown
exit
interface "pl-ix"
    priority 1
    hello-interval 30
    multicast-senders auto
    no tracking-support
    bsm-check-rtr-alert
    no shutdown
exit
apply-to none
rp
    no bootstrap-import
    no bootstrap-export
    static
        address 2.22.187.237
        no override
        group-prefix 224.24.24.24/32
    exit
exit
bsr-candidate
    shutdown
    priority 0
    hash-mask-len 30
    no address
exit
rp-candidate
    shutdown
    no address
    holdtime 150
    priority 192
exit
exit
no shutdown
-----
A:LAX>config>router>pim#

```

Common Configuration Tasks

The following sections describe basic multicast configuration tasks.

- [Configuring IGMP Parameters on page 62](#)
 - [Enabling IGMP on page 62](#)
 - [Configuring an IGMP Interface on page 64](#)
 - [Configuring Static Parameters on page 65](#)
 - [Configuring SSM Translation on page 67](#)
- [Configuring PIM Parameters on page 68](#)
 - [Enabling PIM on page 68](#)
 - [Configuring PIM Interface Parameters on page 69](#)
 - [Importing PIM Join/Register Policies on page 74](#)
- [Configuring Multicast Source Discovery Protocol \(MSDP\) Parameters on page 76](#)
- [Configuring MCAC Parameters on page 77](#)
- [Disabling IGMP or PIM on page 80](#)

Configuring IGMP Parameters

Enabling IGMP

Use the following CLI syntax to enable IGMP.

CLI Syntax: `config>router# igmp`

The following example displays the detailed output when IGMP is enabled.

```
A:LAX>>config>router# info detail
...
#-----
echo "IGMP Configuration"
#-----
      igmp
        query-interval 125
        query-last-member-interval 1
```

```
        query-response-interval 10
        robust-count 2
        no shutdown
    exit
#-----
A:LAX>>config>system#
```

Configuring an IGMP Interface

To configure an IGMP interface:

CLI Syntax:

```
config>router# igmp
      interface ip-int-name
      max-groups value
      import policy-name
      version version
      no shutdown
```

Use the following CLI syntax to configure IGMP interfaces:

Example:

```
config>router#
config>router>igmp# interface "lax-vls"
config>router>igmp>if? no shutdown
config>router>igmp>if# exit
config>router>igmp# interface "p1-ix"
config>router>igmp>if? no shutdown
config>router>igmp>if# exit
config>router>igmp# interface "lax-sjc"
config>router>igmp>if? no shutdown
config>router>igmp>if# exit
```

The following example displays the IGMP configuration:

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      exit
      interface "lax-vls"
      exit
      interface "p1-ix"
      exit
-----
A:LAX>config>router>igmp# exit
```

Configuring Static Parameters

To add an IGMP static multicast source:

CLI Syntax:

```
config>router# igmp
      interface ip-int-name
      no shutdown
      static
      group grp-ip-address
      source ip-address
```

Use the following CLI syntax to configure static group addresses and source addresses for the SSM translate group ranges:

Example:

```
config>router>igmp# interface lax-vls
config>router>igmp>if# static
config>router>igmp>if>static# group 229.255.0.2
config>router>igmp>if>static>group# source 172.22.184.197
config>router>igmp>if>static>group# exit
config>router>igmp>if>static# exit
config>router>igmp>if# exit
```

The following example displays the configuration:

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      exit
      interface "lax-vls"
      static
      group 229.255.0.2
      source 172.22.184.197
      exit
      exit
      exit
      interface "pl-ix"
      exit
-----
A:LAX>config>router>igmp#
```

Configuring IGMP Parameters

To add an IGMP static starg entry:

CLI Syntax:

```
config>router# igmp
      interface ip-int-name
      no shutdown
      static
      group grp-ip-address
      starg
```

Use the following CLI syntax to configure static group addresses and add a static (*,G) entry:

Example:

```
config>router>igmp# interface lax-sjc
config>router>igmp>if# static
config>router>igmp>if>static# group 230.1.1.1
config>router>igmp>if>static>group# starg
config>router>igmp>if>static>group# exit
config>router>igmp>if>static# exit
config>router>igmp>if# exit
config>router>igmp#
```

The following example displays the configuration:

```
A:LAX>config>router>igmp# info
-----
      interface "lax-sjc"
      static
      group 230.1.1.1
      starg
      exit
    exit
  exit
  interface "lax-vls"
  static
  group 229.255.0.2
  source 172.22.184.197
  exit
  exit
  exit
  interface "p1-ix"
  exit
-----
A:LAX>config>router>igmp#
```

Configuring SSM Translation

To configure IGMP parameters:

CLI Syntax: config>router# igmp
 ssm-translate
 grp-range *start end*
 source *ip-address*

The following example displays the command usage to configure IGMP parameters:

Example: config>router# igmp
 config>router>igmp# ssm-translate
 config>router>igmp>ssm# grp-range 229.255.0.1 231.2.2.2
 config>router>igmp>ssm>grp-range# source 10.1.1.1

The following example displays the SSM translation configuration:

```
A:LAX>config>router>igmp# info
-----
      ssm-translate
      grp-range 229.255.0.1 231.2.2.2
      source 10.1.1.1
      exit
    exit
  interface "lax-sjc"
    static
      group 230.1.1.1
      starg
      exit
    exit
  exit
  interface "lax-vls"
    static
      group 229.255.0.2
      source 172.22.184.197
      exit
    exit
  exit
  interface "p1-ix"
  exit
-----
A:LAX>config>router>igmp# exit
```

Configuring PIM Parameters

- [Enabling PIM on page 68](#)
 - [Configuring PIM Interface Parameters on page 69](#)
 - [Importing PIM Join/Register Policies on page 74](#)
-

Enabling PIM

When configuring PIM, make sure to enable PIM on all interfaces for the routing instance, otherwise multicast routing errors can occur.

Use the following CLI syntax to enable PIM.

CLI Syntax:config>router# pim

The following example displays the detailed output when PIM is enabled.

```
A:LAX>>config>router# info detail
...
#-----
echo "PIM Configuration"
#-----
    pim
        no import join-policy
        no import register-policy
        apply-to none
        rp
            no bootstrap-import
            no bootstrap-export
            static
            exit
            bsr-candidate
                shutdown
                priority 0
                hash-mask-len 30
                no address
            exit
            rp-candidate
                shutdown
                no address
                holdtime 150
                priority 192
            exit
        exit
        no shutdown
    exit
#-----
...
A:LAX>>config>system#
```


Configuring PIM Interface Parameters

The following example displays the command usage to configure PIM interface parameters:

```
Example:A:LAX>config>router# pim
A:LAX>config>router>pim# interface "system"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "lax-vls"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "lax-sjc"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# interface "p1-ix"
A:LAX>config>router>pim>if# exit
A:LAX>config>router>pim# rp
A:LAX>config>router>pim>rp# static
A:LAX>config>router>pim>rp>static# address 2.22.187.237
A:LAX>config>router>..>address# group-prefix 224.24.24.24/32
A:LAX>config>router>pim>rp>static>address# exit
A:LAX>config>router>pim>rp>static# exit
A:LAX>config>router>pim>rp# exit
A:LAX>config>router>pim#
```

The following example displays the PIM configuration:

```
A:LAX>config>router>pim# info
-----
        interface "system"
        exit
        interface "lax-vls"
        exit
        interface "lax-sjc"
        exit
        interface "p1-ix"
        exit
        rp
            static
                address 2.22.187.237
                group-prefix 224.24.24.24/32
                exit
                address 10.10.10.10
                exit
            exit
            bsr-candidate
                shutdown
            exit
            rp-candidate
                shutdown
            exit
        exit
-----
A:LAX>config>router>pim#
```

Configuring PIM Parameters

Example: A:SJC>config>router# pim
A:SJC>config>router>pim# interface "system"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-lax"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-nyc"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# interface "sjc-sfo"
A:SJC>config>router>pim>if# exit
A:SJC>config>router>pim# rp
A:SJC>config>router>pim>rp# static
A:SJC>config>router>pim>rp>static# address 2.22.187.237
A:SJC>config>router>pim>rp>static>address# group-prefix
224.24.24.24/32
A:SJC>config>router>pim>rp>static>address# exit
A:SJC>config>router>pim>rp>static# exit
A:SJC>config>router>pim>rp# exit
A:SJC>config>router>pim#

A:SJC>config>router>pim# info

interface "system"
exit
interface "sjc-lax"
exit
interface "sjc-nyc"
exit
interface "sjc-sfo"
exit
rp
static
address 2.22.187.237
group-prefix 224.24.24.24/32
exit
exit
bsr-candidate
shutdown
exit
rp-candidate
shutdown
exit
exit

A:SJC>config>router>pim#

```

Example:A:MV>config>router# pim
A:MV>config>router>pim# interface "system"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# interface "mv-sfo"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# interface "mv-vlc"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# interface "p3-ix"
A:MV>config>router>pim>if# exit
A:MV>config>router>pim# rp
A:MV>config>router>pim>rp# static
A:MV>config>router>pim>rp>static# address 2.22.187.237
A:MV>config>router>pim>rp>static>address# group-prefix
224.24.24.24/32
A:MV>config>router>pim>rp>static>address# exit
A:MV>config>router>pim>rp>static#
A:MV>config>router>pim>rp# exit
A:MV>config>router>pim#

```

```

A:MV>config>router>pim# info
-----
        interface "system"
        exit
        interface "mv-sfo"
        exit
        interface "mv-vlc"
        exit
        interface "p3-ix"
        exit
        rp
            static
                address 2.22.187.237
                group-prefix 224.24.24.24/32
            exit
        exit
        bsr-candidate
            address 2.22.187.236
            no shutdown
        exit
        rp-candidate
            address 2.22.187.236
            no shutdown
        exit
    exit
-----
A:MV>config>router>pim#

```

Configuring PIM Parameters

Example: A:SFO>config>router# pim
A:SFO>config>router>pim# interface "system"
A:SFO>config>router>pim>if# exit
A:SFO>config>router>pim# interface "sfo-sfc"
A:SFO>config>router>pim>if# exit
A:SFO>config>router>pim# interface "sfo-was"
A:SFO>config>router>pim>if# exit
A:SFO>config>router>pim# interface "sfo-mv"
A:SFO>config>router>pim>if# exit
A:SFO>config>router>pim# rp
A:SFO>config>router>pim>rp# static
A:SFO>config>router>pim>rp>static# address 2.22.187.237
A:SFO>config>router>pim>rp>static>address# group-prefix
224.24.24.24/32
A:SFO>config>router>pim>rp>static>address# exit
A:SFO>config>router>pim>rp>static# exit
A:SFO>config>router>pim>rp # exit
A:SFO>config>router>pim#

A:SFO>config>router>pim# info

interface "system"
exit
interface "sfo-sjc"
exit
interface "sfo-was"
exit
interface "sfo-mv"
exit
rp
static
address 2.22.187.237
group-prefix 224.24.24.24/32
exit
exit
bsr-candidate
address 2.22.187.239
no shutdown
exit
rp-candidate
address 2.22.187.239
no shutdown
exit
exit

A:SFO>config>router>pim#

```

Example:A:WAS>config>router# pim
A:WAS>config>router>pim# interface "system"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "was-sfo"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "was-vlc"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# interface "p4-ix"
A:WAS>config>router>pim>if# exit
A:WAS>config>router>pim# rp
A:WAS>config>router>pim>rp# static
A:WAS>config>router>pim>rp>static# address 2.22.187.237
A:WAS>config>router>pim>rp>static>address# group-prefix
224.24.24.24/32
A:WAS>config>router>pim>rp>static>address# exit
A:WAS>config>router>pim>rp>static# exit
A:WAS>config>router>pim>rp# bsr-candidate
A:WAS>config>router>pim>rp>bsr-cand# address 2.22.187.240
A:WAS>config>router>pim>rp>bsr-cand# no shutdown
A:WAS>config>router>pim>rp>bsr-cand# exit
A:WAS>config>router>pim>rp# exit
A:WAS>config>router>pim#

```

```

A:WAS>config>router>pim# info
-----
      interface "system"
      exit
      interface "was-sfo"
      exit
      interface "was-vlc"
      exit
      interface "p4-ix"
      exit
      rp
        static
          address 2.22.187.237
          group-prefix 224.24.24.24/32
          exit
        exit
        bsr-candidate
          address 2.22.187.240
          no shutdown
        exit
        rp-candidate
          address 2.22.187.240
          no shutdown
        exit
      exit
-----
A:WAS>config>router>pim#

```

Importing PIM Join/Register Policies

The import command provides a mechanism to control the (*,G) and (S,G) state that gets created on a router. Import policies are defined in the **config>router>policy-options** context.

Note, in the import policy, if an action is not specified in the entry then the default-action takes precedence. If no entry matches then the default-action also takes precedence. If no default-action is specified, then the default default-action is executed.

Use the following commands to configure PIM parameters:

CLI Syntax:

```
config>router# pim
import {join-policy|register-policy} [policy-name]
[... policy-name]
```

The following example displays the command usage to apply the policy statement which does not allow join messages for group 229.50.50.208/32 and source 192.168.0.0/16 but allows join messages for 192.168.0.0/16, 229.50.50.208 (see [Configuring Route Policy Components on page 912](#)):

Example:

```
config>router# pim
config>router>pim# import join-policy "foo"
config>router>pim# no shutdown
```

The following example displays the PIM configuration:

```
A:LAX>config>router>pim# info
-----
import join-policy "foo"
interface "system"
exit
interface "lax-vls"
exit
interface "lax-sjc"
exit
interface "pl-ix"
exit
rp
static
address 2.22.187.237
group-prefix 224.24.24.24/3
exit
address 10.10.10.10
exit
exit
bsr-candidate
shutdown
```

```
        exit
      rp-candidate
        shutdown
      exit
    exit
  -----
A:LAX>config>router>pim#
```

Configuring Multicast Source Discovery Protocol (MSDP) Parameters

Use the following commands to configure basic MSDP parameters:

CLI Syntax:

```
config>router# msdp
  peer ip-address
    active-source-limit number
    authentication-key [authentication-key|hash-key]
    [hash|hash2]
    default-peer
    export policy-name [policy-name...(up to 5 max)]
    import policy-name [policy-name...(up to 5 max)]
    local-address ip-address
    receive-msdp-msg-rate number intervalseconds [threshold
    threshold]
    no shutdown
no shutdown
```

Use the following CLI syntax to configure MSDP parameters.

Example:

```
config>router>msdp# peer 10.20.1.1
config>router>msdp>peer# local-address 10.20.1.6
config>router>msdp>peer# no shutdown
config>router>msdp>peer# exit
config>router>msdp# no shutdown
config>router>msdp#
```

The following example displays the MSDP configuration:

```
ALA-48>config>router>msdp# info
-----
      peer 10.20.1.1
      local-address 10.20.1.6
      exit
-----
ALA-48>config>router>msdp#
```


Configuring MCAC Parameters

The MCAC policies can be added to a SAP, spoke SDP, mesh SDP, an IGMP interface, and a PIM interface.

The following example displays the command usage to create MCAC policies.

```
Example : config>router# mcac
config>router>mcac# policy "btv_fr"
config>router>mcac>policy# description "foreign TV offering"
config>router>mcac>policy# bundle "FOR" create
config>router>mcac>policy>bundle# bandwidth 30000
config>router>mcac>policy>bundle# channel 224.0.3.1 224.0.3.1 bw 4000
config>router>mcac>policy>bundle# channel 224.0.3.2 224.0.3.2 bw 4000
config>router>mcac>policy>bundle# channel 224.0.4.1 224.0.4.1 bw 3500 class high type
mandatory
config>router>mcac>policy>bundle# channel 224.0.4.2 224.0.4.2 bw 3500 class high
config>router>mcac>policy>bundle# channel 224.0.4.3 224.0.4.3 bw 2800 type mandatory
config>router>mcac>policy>bundle# channel 224.0.4.4 224.0.4.4 bw 2800
config>router>mcac>policy>bundle# mc-constraints
config>router>mcac>policy>bundle>mc-constraints# level 1 bw 20000
config>router>mcac>policy>bundle>mc-constraints# level 2 bw 20000
config>router>mcac>policy>bundle>mc-constraints# level 3 bw 20000
config>router>mcac>policy>bundle>mc-constraints# level 4 bw 20000
config>router>mcac>policy>bundle>mc-constraints# level 5 bw 20000
config>router>mcac>policy>bundle>mc-constraints# level 6 bw 20000
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 1 level 1
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 2 level 3
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 3 level 5
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 1 level 1
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 2 level 3
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 3 level 5
config>router>mcac>policy>bundle>mc-constraints# exit
config>router>mcac>policy>bundle# no shutdown
config>router>mcac>policy>bundle# exit
config>router>mcac>policy# exit
config>router>mcac# policy "btv_vl"
config>router>mcac>policy# description "eastern TV offering"
config>router>mcac>policy# bundle "VRT" create
config>router>mcac>policy>bundle# bandwidth 120000
config>router>mcac>policy>bundle# channel 224.1.2.0 224.1.2.4 bw 4000class high type
mandatory
config>router>mcac>policy>bundle# channel 224.1.2.5 224.1.2.5 bw 20000 type mandatory
config>router>mcac>policy>bundle# channel 224.1.2.10 224.1.2.10 bw 8000 type mandatory
config>router>mcac>policy>bundle# channel 224.2.2.0 224.2.2.4 bw 4000
config>router>mcac>policy>bundle# channel 224.2.2.5 224.2.2.5 bw 10000 class high
config>router>mcac>policy>bundle# channel 224.2.2.6 224.2.2.6 bw 10000 class high
config>router>mcac>policy>bundle# channel 224.2.2.7 224.2.2.7 bw 10000
config>router>mcac>policy>bundle# channel 224.2.2.8 224.2.2.8 bw 10000
config>router>mcac>policy>bundle# mc-constraints
config>router>mcac>policy>bundle>mc-constraints# level 1 bw 60000
config>router>mcac>policy>bundle>mc-constraints# level 2 bw 50000
config>router>mcac>policy>bundle>mc-constraints# level 3 bw 40000
config>router>mcac>policy>bundle>mc-constraints# level 4 bw 30000
config>router>mcac>policy>bundle>mc-constraints# level 5 bw 20000
```

Configuring MCAC Parameters

```
config>router>mcac>policy>bundle>mc-constraints# level 6 bw 10000
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 1 level 1
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 2 level 3
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 1 number-down 3 level 5
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 1 level 1
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 2 level 3
config>router>mcac>policy>bundle>mc-constraints# lag-port-down 2 number-down 3 level 5
config>router>mcac>policy>bundle>mc-constraints# exit
config>router>mcac>policy>bundle# no shutdown
config>router>mcac>policy>bundle# exit
config>router>mcac>policy# exit
```

The following example displays the configuration:

```
*A:ALA-48>config>router>mcac# info
-----
policy "btv_fr"
  description "foreign TV offering"
  bundle "FOR" create
    bandwidth 30000
    channel 224.0.3.1 224.0.3.1 bw 4000
    channel 224.0.3.2 224.0.3.2 bw 4000
    channel 224.0.4.1 224.0.4.1 bw 3500 class high type mandatory
    channel 224.0.4.2 224.0.4.2 bw 3500 class high
    channel 224.0.4.3 224.0.4.3 bw 2800 type mandatory
    channel 224.0.4.4 224.0.4.4 bw 2800
    mc-constraints
      level 1 bw 20000
      level 2 bw 20000
      level 3 bw 20000
      level 4 bw 20000
      level 5 bw 20000
      level 6 bw 20000
      lag-port-down 1 number-down 1 level 1
      lag-port-down 1 number-down 2 level 3
      lag-port-down 1 number-down 3 level 5
      lag-port-down 2 number-down 1 level 1
      lag-port-down 2 number-down 2 level 3
      lag-port-down 2 number-down 3 level 5
    exit
  no shutdown
exit
policy "btv_vl"
  description "eastern TV offering"
  bundle "VRT" create
    bandwidth 120000
    channel 224.1.2.0 224.1.2.4 bw 4000 class high type mandatory
    channel 224.1.2.5 224.1.2.5 bw 20000 type mandatory
    channel 224.1.2.10 224.1.2.10 bw 8000 type mandatory
    channel 224.2.2.0 224.2.2.4 bw 4000
    channel 224.2.2.5 224.2.2.5 bw 10000 class high
    channel 224.2.2.6 224.2.2.6 bw 10000 class high
    channel 224.2.2.7 224.2.2.7 bw 10000
    channel 224.2.2.8 224.2.2.8 bw 10000
    mc-constraints
      level 1 bw 60000
```

```
level 2 bw 50000
level 3 bw 40000
level 4 bw 30000
level 5 bw 20000
level 6 bw 10000
lag-port-down 1 number-down 1 level 1
lag-port-down 1 number-down 2 level 3
lag-port-down 1 number-down 3 level 5
lag-port-down 2 number-down 1 level 1
lag-port-down 2 number-down 2 level 3
lag-port-down 2 number-down 3 level 5
exit
no shutdown
exit
exit
-----
*A:ALA-48>config>router>mcac#
```

Disabling IGMP or PIM

Use the following CLI syntax to disable IGMP and PIM:

CLI Syntax:

```
config>router#  
  igmp  
    shutdown  
  msdp  
    shutdown  
  pim  
    shutdown
```

The following example displays the command usage to disable multicast:

Example:

```
config>router# igmp  
config>router>igmp# shutdown  
config>router>igmp# exit  
config>router#  
config>router>msdp# shutdown  
config>router>msdp# exit  
config>router# pim  
config>router>pim# shutdown  
config>router>pim# exit
```

The following example displays the configuration output:

```
A:LAX>config>router# info
-----
...
#-----
echo "IGMP Configuration"
#-----
    igmp
        shutdown
        ssm-translate
            grp-range 229.255.0.1 231.2.2.2
            source 10.1.1.1
        exit
    exit
    interface "lax-sjc"
        static
            group 230.1.1.1
            starg
        exit
    exit
    interface "lax-vls"
        static
            group 229.255.0.2
            source 172.22.184.197
        exit
    exit
    interface "p1-ix"
    exit
exit
#-----

#-----
echo "MSDP Configuration"
#-----
    msdp
        shutdown
        peer 10.20.1.1
            local-address 10.20.1.6
        exit
        group "test"
            active-source-limit 50000
            receive-msdp-msg-rate 100 interval 300 threshold 5000
            export "LDP-export"
            import "LDP-import"
            local-address 10.10.10.103
            mode mesh-group
            peer 10.10.10.104
        exit
    exit
exit
#-----

echo "PIM Configuration"
#-----
    pim
```

Disabling IGMP or PIM

```
shutdown
import join-policy "foo"
interface "system"
exit
interface "lax-sjc"
exit
interface "lax-vls"
exit
interface "p1-ix"
exit
rp
    static
        address 2.22.187.237
        group-prefix 224.24.24.24/32
        exit
        address 10.10.10.10
        exit
    exit
    bsr-candidate
        shutdown
    exit
    rp-candidate
        shutdown
    exit
exit
exit
#-----
....
-----
A:LAX>config>router#
```

Multicast Command Reference

Command Hierarchies

- [Configuration Commands on page 83](#)
 - [IGMP Commands on page 83](#)
 - [PIM Commands on page 86](#)
 - [MSDP Commands on page 89](#)
 - [Multicast CAC Policy Commands on page 91](#)
 - [Multicast Listener Discovery \(MLD\) Commands on page 92](#)
- [Operational Commands on page 94](#)
- [Show Commands on page 94](#)
 - [Clear Commands on page 95](#)
- [Debug Commands on page 97](#)

Configuration Commands

```

config
  — router
    — [no] ip-fast-reroute
    — mc-maximum-routes number [log-only] [threshold threshold]
    — no mc-maximum-routes
    — multicast-info policy-name
    — no multicast-info

```

IGMP Commands

```

config
  — router
    — [no] igmp
      — [no] group-interface ip-int-name
        — [no] disable-router-alert-check
        — import policy-name
        — no import
        — max-groups [1..16000]
        — no max-groups
        — max-grp-sources [1..32000]
        — no max-grp-sources
        — max-sources [1..1000]
        — no max-sources
        — mcac
          — mc-constraints
            — [no] shutdown
          — policy policy-name
          — no policy
          — unconstrained-bw bandwidth mandatory-bw mandatory-bw

```

```

— no unconstrained-bw
— query-src-ip ip-address
— no query-src-ip
— [no] shutdown
— [no] sub-hosts-only
— [no] subnet-check
— version version
— no version
— grp-if-query-src-ip ip-address
— no grp-if-query-src-ip
— [no] interface ip-int-name
— [no] disable-router-alert-check
— [no] group-interface ip-int-name
— [no] shutdown
— import policy-name
— no import
— max-groups [1..16000]
— no max-groups
— max-grp-sources [1..32000]
— no max-grp-sources
— max-sources [1..1000]
— no max-sources
— mcac
— mc-constraints
— level level-id bw bandwidth
— no level level-id
— number-down number-lag-port-down level level-id
— no number-down number-lag-port-down
— [no] shutdown
— policy mcac-policy-name
— no policy
— unconstrained-bw bandwidth mandatory-bw mandatory-bw
— no unconstrained-bw
— [no] shutdown
— ssm-translate
— [no] grp-range start end
— [no] source ip-address
— static
— [no] group grp-ip-address
— [no] source ip-address
— [no] starg
— [no] subnet-check
— version version
— no version
— query-interval seconds
— no query-interval
— query-last-member-interval seconds
— no query-last-member-interval
— query-response-interval seconds
— no query-response-interval
— robust-count robust-count
— no robust-count
— [no] shutdown
— ssm-translate
— [no] grp-range start end

```


- [no] **source** *ip-address*
- [no] **sub-hosts-only**
- [no] **tunnel-interface** **rsvp-p2mp** *lsp-name*
 - **static**
 - [no] **group** *grp-ip-address*
 - [no] **source** *ip-address*
 - [no] **starg**

PIM Commands

```

config
  — router
    — [no] pim
      — apply-to {ies | non-ies | all | none}
      — [no] enable-mdt-spt
      — import {join-policy | register-policy} policy-name [.. policy-name]
      — no import {join-policy | register-policy}
      — [no] interface ip-int-name
        — assert-period assert-period
        — no assert-period
        — [no] bfd-enable [ipv4 | ipv6]
        — [no] bsm-check-rtr-alert
        — hello-interval hello-interval
        — no hello-interval
        — hello-multiplier deci-units
        — no hello-multiplier
        — [no] improved-assert
        — [no] ipv4-multicast-disable
        — [no] ipv6-multicast-disable
        — max-groups value
        — no max-groups
        — mcac
          — mc-constraints
            — level level bw bandwidth
            — no level level
            — number-down number-lag-port-down level level-id
            — no number-down number-lag-port-down
            — [no] shutdown
          — policy policy-name
          — no policy
          — unconstrained-bw bandwidth mandatory-bw mandatory-bw
          — no unconstrained-bw
        — multicast-senders {auto | always | never}
        — no multicast-senders
        — [no] p2mp-ldp-tree-join [ipv4] [ipv6]
        — priority dr-priority
        — no priority
        — [no] shutdown
        — sticky-dr [priority dr-priority]
        — no sticky-dr
        — three-way-hello [compatibility-mode]
        — no three-way-hello
        — [no] tracking-support
      — [no] ipv4-multicast-disable
      — ipv6-multicast-disable
      — [no] lag-usage-optimization
      — [no] mc-ecmp-balance
      — mc-ecmp-balance-hold minutes
      — no mc-ecmp-balance-hold
      — [no] mc-ecmp-hashing-enabled
      — [no] non-dr-attract-traffic
      — rp
        — [no] anycast rp-ip-address

```

```

— [no] rp-set-peer ip-address
— [no] auto-rp-discovery
— bootstrap-export policy-name [.. policy-name]
— no bootstrap-export
— bootstrap-import policy-name [.. policy-name]
— no bootstrap-import
— bsr-candidate
— address ip-address
— no address
— hash-mask-len hash-mask-length
— no hash-mask-len
— priority bootstrap-priority
— no priority
— [no] shutdown
— ipv6
— [no] anycast rp-ip-address
— [no] rp-set-peer ip-address
— bsr-candidate
— address ip-address
— no address
— hash-mask-len hash-mask-length
— no hash-mask-len
— priority bootstrap-priority
— no priority
— [no] shutdown
— [no] embedded-rp
— [no] group-range ipv6-address/prefix-length
— [no] shutdown
— rp-candidate
— address ip-address
— no address
— [no] group-range {grp-ip-address/mask | grp-ip-address netmask}
— holdtime holdtime
— no holdtime
— priority priority
— no priority
— [no] shutdown
— static
— [no] address ip-address
— [no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}
— [no] override
— rp-candidate
— address ip-address
— no address
— [no] group-range {grp-ip-address/mask | grp-ip-address netmask}
— holdtime holdtime
— no holdtime
— priority priority
— no priority
— [no] shutdown
— static
— [no] address ip-address

```

- [no] **group-prefix** {grp-ip-address/mask | grp-ip-address netmask}
- [no] **override**
- **rpf-rtm** rtm-id | rtm-name
- [no] **rpf6-table** {rtable6-m | rtable6-u | both}
- **rpfv** core
- **rpfv** mvpn
- **rpfv** core mvpn
- no **rpfv** [core] [mvpn]
- [no] **shutdown**
- **spt-switchover-threshold** {grp-ip-address/mask | grp-ip-address netmask} spt-threshold
- no **spt-switchover-threshold** {grp-ip-address/mask | grp-ip-address netmask}
- [no] **ssm-groups**
 - [no] **group-range** {ip-prefix/mask | ip-prefix netmask}
- [no] **tunnel-interface** {rsvp-p2mp lsp-name | ldp-p2mp p2mp-id} [sender ip-address]

MSDP Commands

```

config
— router
— [no] msdp
— [no] active-source-limit number
— [no] data-encapsulation
— export [policy-name...(up to 5 max)]
— no export
— [no] group group-name
— [no] active-source-limit number
— export policy-name [policy-name...(up to 5 max)]
— no export
— import policy-name [policy-name...(up to 5 max)]
— no import
— local-address address
— no local-address
— mode {mesh-group | standard}
— [no] peer peer-address
— [no] active-source-limit number
— authentication-key [authentication-key | hash-key]
  [hash|hash2]
— no authentication-key
— [no] default-peer
— export policy-name [policy-name...(up to 5 max)]
— no export
— import policy-name [policy-name...(up to 5 max)]
— no import
— local-address address
— no local-address
— [no] shutdown
— receive-msdp-msg-rate number interval seconds [threshold number]
— no receive-msdp-msg-rate
— [no] shutdown
— import policy-name [policy-name...(up to 5 max)]
— no import
— local-address address
— no local-address
— [no] peer peer-address
— [no] active-source-limit number
— authentication-key [authentication-key | hash-key] [hash | hash2]
— no authentication-key
— [no] default-peer
— export policy-name [policy-name...(up to 5 max)]
— no export
— import policy-name [policy-name...(up to 5 max)]
— no import
— local-address address
— no local-address
— receive-msdp-msg-rate number interval seconds [threshold number]
— no receive-msdp-msg-rate
— [no] shutdown
— receive-msdp-msg-rate number interval seconds [threshold number]
— no receive-msdp-msg-rate
— rpf6-table {rtable-m | rtable-u | both}

```

- **no rpf6-table**
- **sa-timeout** *seconds*
- **no sa-timeout**
- **[no] shutdown**
- **[no] source** *prefix/mask*
 - **active-source-limit** *number*
 - **no active-source-limit** *number*

Multicast CAC Policy Commands

```

config
  — [no] router
    — mcac
      — [no] policy policy-name
        — [no] bundle bundle-name
          — bandwidth bandwidth
          — no bandwidth
          — channel start-address end-address bw bandwidth [class class]
            [type type] [source ip-prefix/ip-prefix-length]
          — no channel start-address end-address [source ip-prefix/ip-
            prefix-length]
          — description description-string
          — no description
          — mc-constraints
            — lag-port-down lag-id number-down number-lag-
              port-down level level-id
            — no lag-port-down lag-id number-down number-
              lag-port-down
            — level level bw bandwidth
            — no level level
          — [no] shutdown
        — default-action {accept | discard}
        — description description-string
        — no description

```

Multicast Listener Discovery (MLD) Commands

For more information about MLD commands, refer to the SR OS Triple Play Guide.

```

config
  — [no] router
    — [no] mld
      — [no] group-interface ip-int-name
        — [no] disable-router-alert-check
        — import policy-name
        — no import
        — max-groups value
        — no max-groups
        — max-grp-sources [1..32000]
        — no max-grp-sources
        — max-sources [1..1000]
        — no max-sources
        — mcac
          — mc-constraints
            — [no] shutdown
          — policy policy-name
          — no policy
          — unconstrained-bw-bw bandwidth mandatory-bw mandatory-bw
          — no unconstrained-bw
        — query-src-ip ipv6-address
        — no query-src-ip
        — [no] shutdown
        — [no] sub-hosts-only
        — [no] subnet-check
        — version version
        — no version
        — grp-if-query-src-ip ipv6-address
        — no grp-if-query-src-ip
      — [no] interface ip-int-name
        — [no] disable-router-alert-check
        — import policy-name
        — no import
        — max-groups value
        — no max-groups
        — query-interval seconds
        — no query-interval
        — query-last-member-interval seconds
        — no query-last-member-interval
        — query-response-interval seconds
        — no query-response-interval
        — [no] shutdown
        — static
          — [no] group grp-ipv6-address
            — [no] source src-ipv6-address
            — [no] starg
          — version version
          — no version
        — query-interval seconds
        — no query-interval

```


- **query-last-member-interval** *seconds*
- **no query-last-member-interval**
- **query-response-interval** *seconds*
- **no query-response-interval**
- **robust-count** *robust-count*
- **no robust-count**
- **[no] shutdown**
- **ssm-translate**
 - **[no] grp-range** *start end*
 - **[no] source** *src-ipv6-address*

Operational Commands

<GLOBAL>

- **mrinfo** *ip-address* [**router** *router-name* | *service*]
- **mstat** **source** *ip-address* [**group** *grp-ip-address*] [**destination** *dst-ip-address*] [**hop** *hop*] [**router** *router-name* | *service*] [**wait-time** *wait-time*]
- **mtrace** **source** *ip-address* [**group** *grp-ip-address*][**destination** *dst-ip-address*] [**hop** *hop*] [**router** *router-name* | *service*] [**wait-time** *wait-time*]

Show Commands

- show
 - router
 - igmp
 - **group** [*grp-ip-address*]
 - **group** **summary**
 - **hosts** [**group** *grp-address*] [**detail**] [**fwd-service** *service-id*] [**grp-interface** *ip-int-name*]
 - **hosts** [**host** *ip-address*] [**group** *grp-address*] [**detail**]
 - **hosts** **summary**
 - **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-address*] [**detail**]
 - **ssm-translate**
 - **ssm-translate** **interface** *interface-name*
 - **static** [*ip-int-name* | *ip-addr*]
 - **statistics** [*ip-int-name* | *ip-address*]
 - **statistics** **host** [*ip-address*]
 - **status**
 - pim
 - **anycast** [**detail**]
 - **crp** [*ip-address*]
 - **s-pmsi** [*data-mt-interface-name*] [**detail**]
 - **group** [*grp-ip-address*] [**source** *ip-address*] [**type** {**starstarrp**|**starg**|**sg**}] [**detail**] [*family*]
 - **interface** [*ip-int-name* | *mt-int-name* | *ip-address*] [**group** [*grp-ip-address*] **source** *ip-address*] [**type** {**starstarrp** | **starg** | **sg**}] [**detail**] [*family*]
 - **neighbor** [*ip-address* | *ip-int-name* [**address** *ip-address*]] [**detail**] [*family*]
 - **rp** [*ip-address*]
 - **rp-hash** *grp-ip-address*
 - **statistics** [*ip-int-name* | *mt-int-name* | *ip-address*] [*family*]
 - **status** [**detail**] [*family*]
 - show
 - router
 - **mld**
 - **group** [*grp-ipv6-address*]
 - **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-ipv6-address*] [**detail**]
 - **ssm-translate**
 - **static** [*ip-int-name* | *ip-address*]
 - **statistics** [*ip-int-name* | *ipv6-address*]
 - **status**

```

show
  — router
    — msdp
      — group [group-name] [detail]
      — peer [ip-address] [group group-name] [detail]
      — source [ip-address/mask] [type {configured | dynamic | both}] [detail]
      — source-active [group ip-address | local | originator ip-address | peer ip-address |
        source ip-address | {group ip-address source ip-address}] [detail]
      — source-active-rejected [peer-group name] [group ip-address] [source ip-
        address] [originator ip-address] [peer ip-address]
      — statistics [peer ip-address]
      — status
    — mcac
      — policy [policy-name [bundle bundle-name] [protocol protocol-name] [interface
        if-name] [detail]]
      — statistics

show
  — router {router-instance}
    — ldp
      — bindings active
    — mvpn

show
  — router
    — tunnel-table [ip-address [/mask]] [protocol | sdp sdp-id]
    — tunnel-table [summary]

```

Clear Commands

```

clear
  — router
    — igmp
      — database [interface ip-int-name|ip-address] group grp-ip-address [source src-ip-
        address]
      — database grp-interface interface-name [fwd-service service-id]
      — database [interface ip-int-name|ip-address] group grp-ip-address source src-ip-
        address
      — database host [ip-address]
      — database interface ip-int-name|ip-address [group grp-ip-address] [source src-ip-
        address]
      — statistics [interface ip-int-name | ip-address]
      — version [interface ip-int-name | ip-address]
    — mld
      — database [interface ip-int-name|ipv6-address] [group ip-address [source ip-
        address]]
      — statistics [ip-int-name|ipv6-address]
      — version [ip-int-name | ip-address]
    — msdp
      — cache [peer ip-address] [group ip-address] [source ip-address] [originrp ip-
        address]
      — statistics [peer ip-address]
    — pim

```

```
clear
  — service
    — id
      — database [interface ip-int-name | ip-address | mt-int-name] [group grp-ip-address
        [source ip-address]][family]]
      — neighbor [interface ip-int-name | ip-address] [family]
      — s-pmsi [mdSrcAddr] [mdGrpAddr] [vprnSrcAddr vprnGrpAddr]
      — statistics [ {[interface ip-int-name | ip-address | mt-int-name]} {[group grp-ip-
        address [source ip-address]]}[family]]

      — igmp-snooping
        — port-db sap sap-id [group grp-address [source ip-address]]
        — port-db sdp sdp-id:vc-id [group grp-address [source ip-address]]
        — querier
        — statistics [all | sap sap-id | sdp sdp-id:vc-id]

      — pim-snooping
        — database [[sap sap-id | sdp sdp-id:vc-id] [group grp-ip-address]
          [source src-ip-address]]
        — neighbor [ip-address | sap sap-id | sdp sdp-id:vc-id]
        — statistics [sap sap-id | sdp sdp-id:vc-id]
```

Debug Commands

```

debug
  — router
    — igmp
      — [no] group-interface [fwd-service service-id] [ip-int-name]
      — host [ip-address]
      — host [fwd-service service-id] group-interface ip-int-name
      — no host [ip-address]
      — no host [fwd-service service-id] group-interface ip-int-name
      — [no] interface [ip-int-name | ip-address]
      — mcs [ip-int-name]
      — no mcs
      — [no] misc
      — packet [query|v1-report|v2-report|v3-report|v2-leave] host ip-address
      — no packet [query|v1-report|v2-report|v3-report|v2-leave] [ip-int-name|ip-
        address]
      — no packet [query|v1-report|v2-report|v3-report|v2-leave] host ip-address
      — packet [query|v1-report|v2-report|v3-report|v2-leave] [ip-int-name|ip-address]

debug
  — router
    — pim
      — [no] adjacency
      — all [group grp-ip-address] [source ip-address] [detail]
      — no all
      — assert [group grp-ip-address] [source ip-address] [detail]
      — no assert
      — bsr [detail]
      — no bsr
      — data [group grp-ip-address] [source ip-address] [detail]
      — no data
      — db [group grp-ip-address] [source ip-address] [detail]
      — no db
      — interface [ip-int-name | mt-int-name| ip-address] [detail]
      — no interface
      — jp [group grp-ip-address] [source ip-address] [detail]
      — no jp
      — mrrib [group grp-ip-address] [source ip-address] [detail]
      — no mrrib
      — msg [detail]
      — no msg
      — packet [hello | register | register-stop| jp | bsr | assert | crp] [ip-int-name | ip-
        address]
      — no packet
      — register [group grp-ip-address] [source ip-address] [detail]
      — no register
      — rtm [detail]
      — no rtm
      — s-pmsi [{vpnSrcAddr [vpnGrpAddr]}] [mdSrcAddr]] [detail]
      — no s-pmsi

debug
  — router
    — [no] msdp
      — packet [pkt-type] [peer ip-address]
      — no packet

```

Command Hierarchies

- **pim** [*grp-address*]
- **no pim**
- **rtm** [*rp-address*]
- **no rtm**
- **sa-db** [**group** *grpAddr*] [*source srcAddr*] [**rp** *rpAddr*]
- **no sa-db**

Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context

```

config>router>igmp
config>router>igmp>interface
config>router>igmp>interface>group-interface
config>router>igmp>if>mcac>mc-constraints
config>router>pim
config>router>pim>interface
config>router>pim>rp>rp-candidate
config>router>pim>rp>bsr-candidate
config>router>pim>rp>ipv6>rp-candidate
config>router>pim>rp>ipv6>bsr-candidate
config>router>pim>if>mcac>mc-constraints
config>router>msdp
config>router>msdp>peer
config>router>msdp>group
config>router>mcac>policy>bundle
config>router>mld
config>router>mld>group-interface>mcac>mc-constraints
config>router>mld>group-interface
config>router>mld>interface

```

Description The **shutdown** 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 **no shutdown** command.

The **shutdown** 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default

```

no shutdown:  config>router>igmp
               config>router>igmp>interface ip-int-name
               config>router>pim
               config>router>pim>rp>rp-candidate

shutdown:     config>router>pim>rp>bsr-candidate

```

Multicast Commands

ssm-translate

Syntax	ssm-translate
Context	config>router>igmp>interface>shutdown
Description	This command adds or removes ssm-translate group ranges.

source

Syntax	[no] source <i>ip-address</i>
Context	config>router>igmp>interface>shutdown>ssm-translate>grp-range
Description	This command adds or removes source addresses for the SSM translate group range.
Parameters	<i>ip-address</i> — a.b.c.d - unicast source address

grp-range

Syntax	[no] grp-range <i>start end</i>
Context	config>router>igmp>interface>shutdown>ssm-translate
Description	This command adds or removes SSM translate group range entries.
Parameters	<i>start</i> — a.b.c.d - multicast group range start address <i>end</i> — a.b.c.d - multicast group range end address

description

Syntax	description <i>description-string</i> no description
Context	config>router>mcac>policy config>router>mcac>policy>bundle
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of the command removes any description string from the context.

Default No description associated with the configuration context.

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, etc.), the entire string must be enclosed within double quotes.

ip-fast-reroute

Syntax [no] ip-fast-reroute

Context config>router

Description This command configures IP fast reroute.

mc-maximum-routes

Syntax mc-maximum-routes *number* [log-only] [threshold *threshold*]
no mc-maximum-routes

Context config>router

Description This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins will be processed.

The **no** form of the command disables the limit of multicast routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default no mc-maximum-routes

Parameters *number* — Specifies the maximum number of routes to be held in a VRF context.

Values 1 — 2147483647

log-only — Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold *threshold* — The percentage at which a warning log message and SNMP trap should be sent.

Values 0 — 100

Default 1

multicast-info

Syntax **multicast-info-policy** *policy-name*
 no multicast-info-policy

Context configure>router

Description This command configures multicast information policy.

Parameters *policy-name* — Specifies the policy name.

Values 32 chars max

Router IGMP Commands

igmp

Syntax	[no] igmp
Context	config>router
Description	<p>This command enables the Internet Group Management Protocol (IGMP) context. When the context is created, the IGMP protocol is enabled.</p> <p>The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to neighboring multicast routers. An IP multicast router can be a member of one or more multicast groups, in which case it performs both the “multicast router part” of the protocol which collects the membership information needed by its multicast routing protocol, and the “group member part” of the protocol which informs itself and other neighboring multicast routers of its memberships.</p> <p>The no form of the command disables the IGMP instance. To start or suspend execution of IGMP without affecting the configuration, use the no shutdown command.</p>
Default	none

grp-if-query-src-ip

Syntax	grp-if-query-src-ip <i>ip-address</i> no grp-if-query-src-ip
Context	config>router>igmp
Description	<p>This command configures the query source IP address for all group interfaces.</p> <p>The no form of the command removes the IP address.</p>
Default	none

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>igmp
Description	<p>This command enables the context to configure an IGMP interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.</p> <p>The no form of the command deletes the IGMP interface. The shutdown command in the config>router>igmp>interface context can be used to disable an interface without removing the configuration for the interface.</p>

Router IGMP Commands

Default **no interface** — No interfaces are defined.

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, etc.), 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.

disable-router-alert-check

Syntax **[no] disable-router-alert-check**

Context config>router>igmp>if
 config>router>igmp>group-interface

Description This command enables the router alert checking for IGMP messages received on this interface. The **no** form of the command disables the IGMP router alert check option.

group-interface

Syntax **[no] group-interface** *ip-int-name*

Context config>router>igmp>if

Description This command enables IGMP on a group-interface in a VRF context. Activating IGMP under the group-interface is a prerequisite for subscriber replication. The group-interface is also needed so that mcac can be applied and various IGMP parameters defined.

This command can be used in a regular, wholesaler or retailer type of VRF. Note that the retailer VRF does not have the concept of group-interfaces under the subscriber-interface hierarchy. In case that this command is applied to a retailer VRF instance, the optional fwd-service command must be configured. The fwd-service command is referencing the wholesaler VRF in which the traffic is ultimately replicated. Note that redirection in the retailer VRF is supported.

This command enables IGMP on a group-interface in the Global Routing Table (GRT). The group-interface in GRT is defined under the IES service. Activating IGMP under the group-interface is a prerequisite for subscriber replication. The group-interface is also needed so that MCAC can be applied and various IGMP parameters defined.

Default none

Parameters *ip-int-name* — Specifies the name of the group interface.

import

Syntax	import <i>policy-name</i> no import
Context	configure>router>igmp>interface configure>router>igmp>group-interface configure>service>vprn>igmp>interface configure>service>vprn>igmp>group-interface configure>subscr-mgmt>igmp-policy
Description	<p>This command applies the referenced IGMP policy (filter) to a subscriber or a group-interface. An IGMP filter is also known as a black/white list and it is defined under the configure>router>policy-options.</p> <p>When redirection is applied, only the import policy from the subscriber will be in effect. The import policy under the group interface is applicable only for IGMP states received directly on the sap (AN in IGMP proxy mode).</p> <p>The no form of the command removes the policy association from the IGMP instance.</p>
Default	no import — No import policy specified.
Parameters	<i>policy-name</i> — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the config>router>policy-options context.

query-src-ip

Syntax	query-src-ip <i>ip-address</i> no query-src-ip
Context	config>router>igmp>group-interface
Description	<p>This command configures the query source IP address for the group interface. This IP address overrides the source IP address configured at the router level.</p> <p>The no form of the command removes the IP address.</p>
Default	none
Parameters	<i>ip-address</i> — Sets the source IPv4 address for all subscriber's IGMP queries.

sub-hosts-only

Syntax	[no] sub-hosts-only
Context	config>router>igmp>group-interface>mcac config>router>mld>group-interface

Router IGMP Commands

Description	This command enables the handling of IGMP joins received from hosts that are not known in subscriber management or on which no IGMP policy is applied. The no form of the command disables the command.
Default	sub-hosts-only

sub-hosts-only

Syntax	[no] sub-hosts-only
Context	config>router>igmp>group-interface
Description	<p>This command disables processing of IGMP messages outside of the subscriber-host context. No other hosts outside of the subscriber-hosts can create IGMP states.</p> <p>Disabling this command will allow creation of the IGMP states that correspond to the AN that operate in IGMP proxy mode. In this mode the AN will hide source IP addresses of IGMP messages and will source IGMP messages with its own IP address. In this case an IGMP state can be created under the sap context. This IGMP state creation under the sap is controlled via the import policy under the group-interface.</p> <p>IGMP state processing for regular subscriber-hosts is unaffected by this command.</p> <p>The no form of the command disables the command.</p>
Default	sub-hosts-only

max-groups

Syntax	max-groups [1..16000] no max-groups
Context	config>router>igmp>if config>router>igmp>group-interface config>router>pim>if
Description	This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. This command is applicable for IPv4 and IPv6.
Default	0, no limit to the number of groups.
Parameters	<i>value</i> — Specifies the maximum number of groups for this interface. Values 1 — 16000

max-grp-sources

Syntax	max-grp-sources [1..32000] no max-grp-sources
Context	config>router>igmp>interface config>router>igmp>group-interface config>router>mld>group-interface
Description	This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. The no form of the command reverts to the default.
Default	0
Parameters	1 — 32000 — Specifies the maximum number of group source. Values 1 — 32000

max-sources

Syntax	max-sources [1..1000] no max-sources
Context	config>router>igmp>group-interface config>router>mld>group-interface
Description	This command configures the maximum number of group sources for this group-interface.

static

Syntax	static
Context	config>router>igmp>if
Description	This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.
Default	none

group

Syntax	[no] group grp-ip-address
Context	config>router>igmp>if>static
Description	<p>This command enables the context to add a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.</p> <p>When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.</p>
Default	none
Parameters	<i>grp-ip-address</i> — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

source

Syntax	[no] source ip-address
Context	config>router>igmp>if>static>group config>router>igmp>ssm-translate>grp-range
Description	<p>This command specifies a IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the source(s) that the traffic is expected.</p> <p>The source command is mutually exclusive with the specification of individual sources for the same group.</p> <p>The source command in combination with the group is used to create a specific (S,G) static group entry.</p> <p>Use the no form of the command to remove the source from the configuration.</p>
Default	none
Parameters	<i>ip-address</i> — Specifies the IPv4 unicast address.

starg

Syntax	[no] starg
Context	config>router>igmp>if>static>group
Description	<p>This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.</p> <p>Use the no form of the command to remove the starg entry from the configuration.</p>
Default	none

subnet-check

Syntax	[no] subnet-check
Context	config>router>igmp>interface config>router>mld>group-interface config>router>igmp>group-interface>mcac
Description	This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.
Default	enabled

version

Syntax	version <i>version</i> no version
Context	config>router>igmp>if config>router>mld>group-interface config>router>igmp>group-interface>mcac
Description	<p>This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.</p> <p>For IGMPv3, note that a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.</p>
Default	3
Parameters	<i>version</i> — Specifies the IGMP version number.
	Values 1, 2, 3
	Values >= 1000

query-interval

Syntax	query-interval <i>seconds</i> no query-interval
Context	config>router>igmp
Description	This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.
Default	125

seconds — The time frequency, in seconds, that the router transmits general host-query messages.

Values 2 — 1024

query-last-member-interval

Syntax **query-last-member-interval** *seconds*

Context config>router>igmp

Description This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.

Default 1

Parameters *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 — 1024

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>router>igmp

Description This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default 10

Parameters *seconds* — Specifies the the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 1023

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>router>igmp

Description This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.

Default 2

Parameters *robust-count* — Specify the robust count value.

Values 2 — 10

ssm-translate

Syntax **ssm-translate**

Context config>router>igmp

Description This command enables the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the **source** command with **starg** command enabled.

grp-range

Syntax [**no**] **grp-range** *start end*

Context config>router>igmp>ssm-translate

Description This command is used to configure group ranges which are translated to SSM (S,G) entries.

Parameters *start* — An IP address that specifies the start of the group range.
 end — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the *start* value.

source

Syntax [**no**] **source** *ip-address*

Context config>router>igmp>ssm-translate>grp-range

Description This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters *ip-address* — Specifies the IP address that will be sending data.

tunnel-interface

Syntax [**no**] **tunnel-interface** **rsvp-p2mp** *lsp-name*

Context config>router>pim
 config>router>igmp

Router IGMP Commands

Description	<p>This command creates a tunnel interface associated with an RSVP P2MP LSP. IPv4 multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP.</p> <p>At ingress LER, the tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. The user can create one or more tunnel interfaces in PIM and associate each to a different RSVP P2MP LSP. P2mp-ID is required to configure LDP P2MP LSP tunnel interfaces. Sender address for a tunnel itnerface must be specified only on the leaf node.</p> <p>At egress LER, the tunnel interface identifier consists of a couple of string of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER. The LSP name string must not contain “:” (two :s) nor contain a “.” (single “.”) at the end of the LSP name. However, a “.” (single “.”) can appear anywhere in the string except at the end of the name.</p>
Default	none
Parameters	<p>rsvp-p2mp <i>lsp-name</i> — Specifies the LSP. The LSP name can be up to 32 characters long and must be unique.</p> <p><i>psmp-id</i> — Identifier used for signaling mLDP P2MP LSP.</p> <p>Values 1-4294967296</p>

static

Syntax	static
Context	config>router>igmp>tunnel-interface
Description	<p>This command provides the context to configure static multicast receiver hosts on a tunnel interface associated with an RSVP P2MP LSP.</p> <p>When enabled, data is forwarded to an interface without receiving membership reports from host members.</p>
Default	none

group

Syntax	[no] group <i>grp-ip-address</i>
Context	config>router>igmp>tunnel-interface>static
Description	<p>This command enables the context to add a static multicast group either as a (*,G) or one or more (S,G) records.</p> <p>The user can assign static multicast group joins to a tunnel interface associated with an RSVP P2MP LSP. Note that a given <*,G> or <S,G> can only be associated with a single tunnel interface.</p>

A multicast packet which is received on an interface and which succeeds the RPF check for the source address will be replicated and forwarded to all OIFs which correspond to the branches of the P2MP LSP. The packet is sent on each OIF with the label stack indicated in the NHLFE of this OIF. The packets will also be replicated and forwarded natively on all OIFs which have received IGMP or PIM joins for this <S,G>.

The multicast packet can be received over a PIM or IGMP interface which can be an IES interface, a spoke SDP terminated IES interface, or a network interface.

Default none

Parameters *grp-ip-address* — Specifies a multicast group address that receives data on a tunnel interface. The IP address must be unique for each static group.

source

Syntax [no] **source** *ip-address*

Context config>router>igmp>tunnel-interface>static>group

Description This command specifies a IPv4 unicast address of a multicast source. The source command is mutually exclusive with the specification of individual sources for the same group. The source command in combination with the group is used to create a specific (S,G) group entry in a static group join on a tunnel interface associated with a P2MP RSVP LSP.

The **no** form of the command removes the source from the configuration.

Default none

Parameters *ip-address* — Specifies the IPv4 unicast address.

starg

Syntax [no] **starg**

Context config>router>igmp>tunnel-interface>static>group

Description This command adds a static (*,G) group entry in a static group join on a tunnel interface associated with a P2MP RSVP LSP.

This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of the command removes the starg entry from the configuration.

Default none

Router PIM Commands

pim

Syntax	[no] pim
Context	config>router
Description	<p>This command configures a Protocol Independent Multicast (PIM) instance.</p> <p>PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router OS supports PIM sparse mode (PIM-SM).</p>
Default	not enabled

interface

Parameters	[no] interface <i>ip-int-name</i>
Context	config>router>pim
Description	<p>This command creates a PIM interface.</p> <p>Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for config router interface, config service ies interface, and config service ies subscriber-interface group-interface. Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it may be confusing.</p> <p>The no form of the command removes the IP interface and all the associated configurations.</p>
Default	No interfaces or names are defined within PIM.
Parameters	<p><i>ip-int-name</i> — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for config router interface, config service ies interface, and config service ies subscriber-interface group-interface commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p>Values 1 — 32 alphanumeric characters.</p> <p>If the <i>ip-int-name</i> already exists, the context is changed to maintain that IP interface. If <i>ip-int-name</i> does not exist, the interface is created and the context is changed to that interface for further command processing.</p>

apply-to

Syntax	apply-to {ies non-ies all none}
Context	config>router>pim
Description	<p>This command creates a PIM interface with default parameters.</p> <p>If a manually created or modified interface is deleted, the interface will be recreated when (re)processing the apply-to command and if PIM is not required on a specific interface a shutdown should be executed.</p> <p>The apply-to command is first saved in the PIM configuration structure. Then, all subsequent commands either create new structures or modify the defaults as created by the apply-to command.</p>
Default	none (keyword)
Parameters	<p>ies — Creates all IES interfaces in PIM.</p> <p>non-ies — Non-IES interfaces are created in PIM.</p> <p>all — All IES and non-IES interfaces are created in PIM.</p> <p>none — Removes all interfaces that are not manually created or modified. It also removes explicit no interface commands if present.</p>

assert-period

Syntax	assert-period <i>assert-period</i> no assert-period
Context	config>router>pim>if
Description	<p>This command configures the period for periodic refreshes of PIM Assert messages on an interface.</p> <p>The no form of the command removes the assert-period from the configuration.</p>
Default	no assert-period
Parameters	<i>assert-period</i> — Specifies the period for periodic refreshes of PIM Assert messages on an interface.
Values	1 — 300 seconds

bfd-enable

Parameters	[no] bfd-enable [ipv4 ipv6]
Context	config>router>pim>interface
Description	<p>This command enables the use of IPv4 or IPv6 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>

Router PIM Commands

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

Default no bfd-enable

enable-mdt-spt

Syntax [no] enable-mdt-spt

Context config>router>pim

Description This command is used to enable SPT switchover for default MDT. On enable, PIM instance resets all MDTs and reinitiate setup.

The **no** form of the command disables SPT switchover for default MDT. On disable, PIM instance resets all MDTs and reinitiate setup.

Default no enable-mdt-spt

import

Syntax import {join-policy | register-policy} [*policy-name* [.. *policy-name*]]
no import {join-policy | register-policy}

Context config>router>pim

Description This command specifies the import route policy to be used. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, BGP routes are accepted by default. Up to five import policy names can be specified.

The **no** form of the command removes the policy association from the instance.

Default no import join-policy
no import register-policy

Parameters **join-policy** — Use this command to filter PIM join messages which prevents unwanted multicast streams from traversing the network.

register-policy — This keyword filters register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

policy-name — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

ipv4-multicast-disable

Syntax	[no] ipv4-multicast-disable
Context	configure>router>pim configure>router>pim>interface
Description	This command administratively disables/enables PIM operation for IPv4. Note that IPv4 multicast must be enabled to enable mLDLP in-band signaling for IPv4 PIM joins; see p2mp-ldp-tree-join .
Default	no ipv4-multicast-disable

lag-usage-optimization

Syntax	[no] lag-usage-optimization
Context	configure>router>pim
Description	This command specifies whether the router should optimize usage of the LAG such that traffic for a given multicast stream destined to an IP interface using the LAG is sent only to the forwarding complex that owns the LAG link on which it will actually be forwarded. Changing the value causes the PIM protocol to be restarted. If this optimization is disabled, the traffic will be sent to all the forwarding complexes that own at least one link in the LAG. Note that changes made for 9G multicast hashing causes Layer 4 multicast traffic to not hashed. This is independent whether lag-usage-optimization is enabled or disabled.

mc-ecmp-balance

Syntax	[no] mc-ecmp-balance
Context	configure>router>pim
Description	This command enables multicast balancing of traffic over ECMP links. When enabled, each multicast stream that needs to be forwarded over an ECMP link will be re-evaluated for the total multicast bandwidth utilization. Re-evaluation occurs on the ECMP interface in question. The no form of the command disables the multicast balancing.

mc-ecmp-balance-hold

Syntax	mc-ecmp-balance-hold <i>minutes</i> no mc-ecmp-balance-hold
Context	configure>router>pim

Router PIM Commands

Description	This command configures the hold time for multicast balancing over ECMP links.
Parameters	<i>minutes</i> — Specifies the hold time, in minutes, that applies after an interface has been added to the ECMP link.

mc-ecmp-hashing-enabled

Syntax	[no] mc-ecmp-hashing-enabled
Context	configure>router>pim
Description	<p>This command enables hash-based multicast balancing of traffic over ECMP links and causes PIM joins to be distributed over the multiple ECMP paths based on a hash of S and G (and possibly next-hop IP). When a link in the ECMP set is removed, the multicast streams that were using that link are re-distributed over the remaining ECMP links using the same hash algorithm. When a link is added to the ECMP set new joins may be allocated to the new link based on the hash algorithm, but existing multicast streams using the other ECMP links stay on those links until they are pruned.</p> <p>Hash-based multicast balancing is supported for both IPv4 and IPv6.</p> <p>This command is mutually exclusive with the <code>mc-ecmp-balance</code> command in the same context.</p> <p>The no form of the command disables the hash-based multicast balancing of traffic over ECMP links.</p>
Default	no mc-ecmp-hashing-enabled

ipv6-multicast-disable

Syntax	ipv6-multicast-disable
Context	configure>router>pim configure>router>pim>interface
Description	<p>This command administratively disables/enables PIM operation for IPv6.</p> <p>Note that IPv6 multicast must be enabled to enable mLDP in-band signaling for IPv6 PIM joins; see p2mp-ldp-tree-join.</p>
Default	ipv6-multicast-disable

bsm-check-rtr-alert

Syntax	[no] bsm-check-rtr-alert
Context	config>router>pim>interface
Description	This command enables the checking of the router alert option in the bootstrap messages received on this interface.
Default	no bsm-check-rtr-alert

hello-interval

Syntax	hello-interval <i>hello-interval</i> no hello-interval
Context	config>router>pim>interface
Description	This command configures the frequency at which PIM Hello messages are transmitted on this interface. The no form of this command reverts to the default value of the hello-interval.
Default	30
Parameters	<i>hello-interval</i> — Specifies the hello interval in seconds. A 0 (zero) value disables the sending of hello messages (the PIM neighbor will never timeout the adjacency).
Values	0 — 255 seconds

hello-multiplier

Syntax	hello-multiplier <i>deci-units</i> no hello-multiplier
Context	config>router>pim>interface
Description	This command configures the multiplier to determine the holdtime for a PIM neighbor on this interface. The hello-multiplier in conjunction with the hello-interval determines the holdtime for a PIM neighbor.
Parameters	<i>deci-units</i> — Specify the value, specified in multiples of 0.1, for the formula used to calculate the hello-holdtime based on the hello-multiplier: $(\text{hello-interval} * \text{hello-multiplier}) / 10$ This allows the PIMv2 default timeout of 3.5 seconds to be supported.
Values	20 — 100
Default	35

improved-assert

Syntax	[no] improved-assert
Context	config>router>pim>interface
Description	The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes. The assert process is started when data is received on an outgoing interface meaning that duplicate traffic is forwarded to the LAN until the forwarder is negotiated among the routers. When the improved-assert command is enabled, the PIM assert process is done entirely in the control plane. The advantages are that it eliminates duplicate traffic forwarding to the LAN. It also improves performance since it removes the required interaction between the control and data planes.

NOTE: improved-assert is still fully interoperable with the draft-ietf-pim-sm-v2-new-xx, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Revised*, and RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM)*, implementations. However, there may be conformance tests that may fail if the tests expect control-data plane interaction in determining the assert winner. Disabling the **improved-assert** command when performing conformance tests is recommended.

Default enabled

multicast-senders

Syntax **multicast-senders {auto | always | never}**
no multicast-senders

Context config>router>pim>interface

Description This command configures how traffic from directly-attached multicast sources should be treated on broadcast interfaces. It can also be used to treat all traffic received on an interface as traffic coming from a directly-attached multicast source. This is particularly useful if a multicast source is connected to a point-to-point or unnumbered interface.

Default auto

Parameters **auto** — Specifies that, on broadcast interfaces, the forwarding plane performs subnet-match check on multicast packets received on the interface to determine if the packet is from a directly-attached source. On unnumbered/point-to-point interfaces, all traffic is implicitly treated as coming from a remote source.

always — Treats all traffic received on the interface as coming from a directly-attached multicast source.

never — Specifies that, on broadcast interfaces, traffic from directly-attached multicast sources will not be forwarded. Note that traffic from a remote source will still be forwarded if there is a multicast state for it. On unnumbered/point-to-point interfaces, it means that all traffic received on that interface must not be forwarded.

p2mp-ldp-tree-join

Syntax **[no] p2mp-ldp-tree-join [ipv4] [ipv6]**

Context config>router>pim>interface

Description This command configures the option to join the P2MP LDP tree towards the multicast source. If **p2mp-ldp-tree-join** is enabled, a PIM multicast join received on an interface is processed to join the P2MP LDP LSP, using the in-band signaled P2MP tree for the same multicast flow. LDP P2MP tree is set up towards the multicast source. The route to the multicast node source is looked up from the RTM. The next-hop address for the route to source is set as the root of LDP P2MP tree.

The **no** form of the command disables joining the P2MP LDP tree for IPv4 or IPv6 or for both (if both or none is specified).

Parameters	ipv4 — Enables dynamic mLDP in-band signaling for IPv4 PIM joins. IPv4 multicast must be enabled; see ipv4-multicast-disable . For backward compatibility p2mp-ldp-tree-join is equivalent to p2mp-ldp-tree-join ipv4 .
	ipv6 — Enables dynamic mLDP in-band signaling for IPv6 PIM joins. IPv6 multicast must be enabled; see ipv6-multicast-disable).
	no p2mp-ldp-tree-join

priority

Syntax	priority <i>dr-priority</i> no priority
Context	config>router>pim>interface
Description	This command sets the priority value to elect the designated router (DR). The DR election priority is a 32-bit unsigned number and the numerically larger priority is always preferred. The no form of the command restores the default values.
Default	1
Parameters	<i>priority</i> — Specifies the priority to become the designated router. The higher the value, the higher the priority. Values 1 — 4294967295

priority

Syntax	priority <i>bootstrap-priority</i> no priority
Context	config>router>pim>rp>bsr-candidate
Description	This command configures the bootstrap priority of the router. The RP is sometimes called the bootstrap router. The priority determines if the router is eligible to be a bootstrap router. In the case of a tie, the router with the highest IP address is elected to be the bootstrap router.
Default	0
Parameters	<i>bootstrap-priority</i> — Specifies the priority to become the bootstrap router. The higher the value, the higher the priority. A 0 value the router is not eligible to be the bootstrap router. A value of 1 means router is the least likely to become the designated router. Values 0 — 255

priority

Syntax	priority <i>priority</i> no priority
Context	config>router>pim>rp>rp-candidate config>router>pim>rp>ipv6>rp-candidate
Description	This command configures the Candidate-RP priority for becoming a rendezvous point (RP). This value is used to elect RP for a group range.
Default	192
Parameters	<i>priority</i> — Specifies the priority to become a rendezvous point (RP). A value of 0 is considered as the highest priority. Values 0 — 255

sticky-dr

Syntax	sticky-dr [priority <i>dr-priority</i>] no sticky-dr
Context	config>router>pim>interface
Description	<p>This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the designated router (DR) will be modified to the value configured in <i>dr-priority</i>. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.</p> <p>By enabling sticky-dr on this interface, it will continue to act as the DR for the LAN even after the old DR comes back up.</p> <p>The no form of the command disables sticky-dr operation on this interface.</p>
Default	disabled
Parameters	priority <i>dr-priority</i> — Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled. Values 1 — 4294967295

three-way-hello

Syntax	three-way-hello [compatibility-mode] no three-way-hello
Context	config>router>pim>interface

Description This command configures the compatibility mode to enable three-way hello. By default value is disabled on all interface which specifies that the standard two- way hello is supported. When enabled, the three way hello is supported.

Default no three-way-hello

tracking-support

Syntax [no] tracking-support

Context config>router>pim>interface

Description This command sets the the T bit in the LAN Prune Delay option of the Hello Message. This indicates the router's capability to enable join message suppression. This capability allows for upstream routers to explicitly track join membership.

Default no tracking-support

rp

Syntax rp

Context config>router>pim

Description This command enables the context to configure rendezvous point (RP) parameters. The address of the root of the group's shared multicast distribution tree is known as its RP. Packets received from a source upstream and join messages from downstream routers rendezvous at this router.

If this command is not enabled, then the router can never become the RP.

ipv6

Syntax ipv6

Context config>router>pim>rp

Description This command enables the context to configure IPv6 parameters.

anycast

Syntax [no] anycast *rp-ip-address*

Context config>router>pim>rp
config>router>pim>rp>ipv6

Router PIM Commands

Description	<p>This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.</p> <p>The no form of the command removes the anycast instance from the configuration.</p>
Default	none
Parameters	<p><i>rp-ip-address</i> — Configure the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level.</p> <p>Values Any valid loopback address configured on the node.</p>

auto-rp-discovery

Syntax	[no] auto-rp-discovery
Context	config>router>pim>rp
Description	<p>This command enables Auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP-candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn about availability of RP nodes present in the network.</p> <p>The no form of the command disables auto RP.</p>
Default	no auto-rp-discovery

rp-set-peer

Syntax	[no] rp-set-peer ip-address
Context	config>router>pim>rp>anycast config>router>pim>rp>ipv6>anycast
Description	<p>This command configures a peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.</p> <p>This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this rp-set.</p> <p>Although there is no set maximum number of addresses that can be configured in an rp-set, up to 15 IP addresses is recommended.</p> <p>The no form of the command removes an entry from the list.</p>
Default	None
Parameters	<p><i>ip-address</i> — Specifies a peer in the anycast rp-set.</p> <p>Values Any valid ip-address within the scope outlined above.</p>

bsr-candidate

Syntax	bsr-candidate
Context	config>router>pim>rp config>router>pim>rp>ipv6
Description	This command enables the context to configure Candidate Bootstrap (BSR) parameters.

rp-candidate

Syntax	rp-candidate
Context	config>router>pim>rp config>router>pim>rp>ipv6
Description	<p>This command enables the context to configure the Candidate RP parameters.</p> <p>Routers use a set of available rendezvous points distributed in Bootstrap messages to get the proper group-to-RP mapping. A set of routers within a domain are also configured as candidate RPs (C-RPs); typically these will be the same routers that are configured as candidate BSRs.</p> <p>Every multicast group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) is the root of this shared tree.</p>
Default	shutdown

static

Syntax	static
Context	config>router>pim>rp config>router>pim>rp>ipv6
Description	<p>This command enables the context to configure static Rendezvous Point (RP) addresses for a multicast group range.</p> <p>Entries can be created or destroyed. If no IP addresses are configured in the config>router>pim>rp>static>address context, then the multicast group to RP mapping is derived from the RP-set messages received from the Bootstrap Router.</p>

address

Syntax	address <i>ip-address</i>
Context	config>router>pim>rp>bsr-candidate config>router>pim>rp>ipv6>bsr-cand

Router PIM Commands

Description	This command is used to configure the candidate BSR IP address. This address is for Bootstrap router election.
Default	none
Parameters	<i>ip-address</i> — The <i>ip-address</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
Values	1.0.0.0 – 223.255.255.255

address

Syntax	[no] address <i>ip-address</i>
Context	config>router>pim>rp>rp-candidate config>router>pim>rp>ipv6>bsr-cand
Description	This command configures the local RP address. This address is sent in the RP candidate advertisements to the bootstrap router.
Default	none
Parameters	<i>ip-address</i> — The <i>ip-address</i> .
Values	1.0.0.0 – 223.255.255.255

address

Syntax	address <i>ip-address</i> no address
Context	config>router>pim>rp>static config>router>pim>rp>ipv6>static
Description	This command indicates the Rendezvous Point (RP) address that should be used by the router for the range of multicast groups configured by the range command.
Default	none
Parameters	<i>ip-address</i> — The static IP address of the RP. The <i>ip-addr</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
Values	1.0.0.0 – 223.255.255.255

embedded-rp

Syntax	[no] embedded-rp
Context	config>router>pim>rp>ipv6
Description	<p>This command enables the context to configure embedded RP parameters.</p> <p>Embedded RP is required to support IPv6 inter-domain multicast because there is no MSDP equivalent in IPv6.</p> <p>The detailed protocol specification is defined in RFC 3956, <i>Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address</i>. This RFC describes a multicast address allocation policy in which the address of the RP is encoded in the IPv6 multicast group address, and specifies a PIM-SM group-to-RP mapping to use the encoding, leveraging, and extending unicast-prefix-based addressing. This mechanism not only provides a simple solution for IPv6 inter-domain ASM but can be used as a simple solution for IPv6 intra-domain ASM with scoped multicast addresses as well. It can also be used as an automatic RP discovery mechanism in those deployment scenarios that would have previously used the Bootstrap Router protocol (BSR).</p> <p>The no form of the command disables embedded RP.</p>

group-range

Syntax	[no] group-range <i>ipv6-address/prefix-length</i>		
Context	config>router>pim>ipv6>rp>embedded-rp		
Description	<p>This command defines which multicast groups can embed RP address information besides FF70::/12. Embedded RP information is only used when the multicast group is in FF70::/12 or the configured group range.</p>		
Parameters	<p><i>ipv6-address/prefix-length</i> — Specifies the group range for embedded RP.</p> <table> <tr> <td>Values</td><td> ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D prefix-length: 16 — 128 </td></tr> </table>	Values	ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D prefix-length: 16 — 128
Values	ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D prefix-length: 16 — 128		

group-range

Syntax	[no] group-range {<i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i>}
Context	config>router>pim>rp>rp-candidate config>router>pim>rp>static>rp>ipv6>rp-candidate
Description	This command configures the address ranges of the multicast groups for which this router can be an RP.
Default	none

Router PIM Commands

Parameters	<i>grp-ip-address</i> — The multicast group IP address expressed in dotted decimal notation. Values 224.0.0.0 — 239.255.255.255 <i>mask</i> — The mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decimal notation (255.255.0.0). Values 4 — 32 <i>netmask</i> — The subnet mask in dotted decimal notation. Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)
-------------------	--

group-range

Syntax	[no] group-range { <i>ip-prefix/mask</i> <i>ip-prefix netmask</i> }
Context	config>router>pim>ssm-groups
Description	This command configures the address ranges of the multicast groups for this router.
Default	none
Parameters	<i>ip-prefix/mask</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area. Values <i>ipv4-prefix:</i> a.b.c.d <i>ipv4-prefix-le:</i> 0 — 32 <i>ipv6-prefix:</i> x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D <i>ipv6-prefix-le:</i> 0 — 128 Values 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal) <i>netmask</i> — The subnet mask in dotted decimal notation. Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

holdtime

Syntax	holdtime <i>holdtime</i> no holdtime
Context	config>router>pim>rp>rp-candidate config>router>pim>rp>ipv6>rp-candidate
Description	This command configures the length of time, in seconds, that neighbors should consider the sending router to be operationally up. A local RP cannot be configured on a logical router.

Parameters *holdtime* — Specifies the hold time, in seconds.

Values 5 — 255

group-prefix

Syntax **[no] group-prefix** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context config>router>pim>rp>static>address
config>router>pim>rp>ipv6>static>address

Description This command specifies the range of multicast group addresses which should be used by the router as the Rendezvous Point (RP). The config>router>pim>rp>static>address a.b.c.d implicitly defaults to deny all for all multicast groups (224.0.0.0/4). A group-prefix must be specified for that static address. This command does not apply to the whole group range.

The **no** form of the command removes the group-prefix from the configuration.

Default none

Parameters *grp-ip-address* — The multicast group IP address expressed in dotted decimal notation.

Values 224.0.0.0 — 239.255.255.255

mask — The mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decimal notation (255.255.0.0).

Values 4 — 32

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

override

Syntax **[no] override**

Context config>router>pim>rp>static>address
config>router>pim>rp>ipv6>static>address

Description This command changes the precedence of static RP over dynamically learned Rendezvous Point (RP). When enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.

Default no override

non-dr-attract-traffic

Syntax	[no] non-dr-attract-traffic
Context	config>router>pim
Description	<p>This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designater router.</p> <p>An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. Note that while using this flag the router may not send the stream down to the DSLAM while not DR.</p> <p>When enabled, the designated router state is ignored. When disabled, no non-dr-attract-traffic, the designated router value is honored.</p>
Default	no non-dr-attract-traffic

rpf-rtm

Syntax	[no] rpf-rtm <i>rtm-id</i> <i>rtm-name</i>
Context	config>router>pim
Description	<p>This command associates the specified RTM instance with the PIM protocol. This RTM will then be used to generate the RPF table for multicast.</p> <p>The no form of this command removes the association with the specified RTM instance and will cause PIM to use the unicast RTM.</p>
Default	No default
Parameters	<p><i>rtm-id</i> — RTM Instance ID that is to be associated with the new IS-IS topology.</p> <p>Values integer: 3 — 32</p> <p><i>rtm-name</i> — string name given to the RTM instance.</p>

rpf6-table

Syntax	rpf6-table {<i>rtable6-m</i> <i>rtable6-u</i> both} no rpf6-table
Context	config>router>pim config>router>msdp

Description	<p>This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route.</p> <p>By default, only the unicast route table is looked up to calculate RPF interface towards the source/ rendezvous point. However the operator can specify the following:</p> <ul style="list-style-type: none"> a) Use unicast route table only b) Use multicast route table only or c) Use both the route tables.
Parameters	<p>rtable6-m — Specifies that only the multicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by static routes, ISIS and OSPF.</p> <p>rtable6-u — Specifies only that the unicast route table will be used by the multicast protocol (PIM) for IPv6 RPF checks. This route table will contain routes submitted by all the unicast routing protocols.</p> <p>both — Will always lookup first in the multicast route table and if there is a route, it will use it. If PIM does not find a route in the first lookup, it will try to find it in the unicast route table. Rtable6-m is checked before rtable6-u.</p>
Default	rtable-u

rpfv

Syntax	<p>rpfv core rpfv mvpn rpfv core mvpn no rpfv [core] [mvpn]</p>
Context	config>router>pim
Description	This command enables RPF Vector processing for Inter-AS Rosen MVPN Option-B and Option-C. The rpfv must be enabled on every node for Inter-AS Option B/C MVPN support.
Parameters	<p>mvpn — Enables mvpn RPF vector processing for Inter-AS Option B/C MVPN based on RFC 5496 and RFC6513. If a core RPF vector is received, it will be dropped before a message is processed.</p> <p>core — Enables core RPF vector (no RD) processing for Inter-AS Option B/C MVPN, which allows SROS interoperability as P-router with third-party vendors that do not encode RD in the RPF vector for Inter-AS MVPN.</p> <p>core mvpn — Enables core RPF vector (no RD) processing for Inter-AS Option B/C MVPN, which allows SROS interoperability as P-router with third-party vendors that do not encode RD in the RPF vector for Inter-AS MVPN.</p> <p>The no version of this command disables RPF Vector processing. If RPF vector is received in a PIM join message, the vector will be removed before local processing of PIM message starts.</p>
Default	no rpfv

sa-timeout

Syntax	sa-timeout <i>seconds</i> no sa-timeout
Context	config>router>msdp
Description	This command configures the value for the SA entries in the cache. If these entries are not refreshed within the timeout value then they are removed from the cache. Normally the entries are refreshed at least once a minute. But under high load with many of MSDP peers the refresh cycle could be incomplete. A higher timeout value (more than 90) could be useful to prevent unstabilities in the MSDP cache.
Default	90
Parameters	<i>seconds</i> — Specifies the time, in seconds, to wait for a response from the peer before declaring the peer unavailable. Values 90 — 600

spt-switchover-threshold

Syntax	spt-switchover-threshold { <i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i> } <i>spt-threshold</i> no spt-switchover-threshold { <i>grp-ip-address/mask</i> <i>grp-ip-address netmask</i> }
Context	config>router>pim
Description	<p>This command configures shortest path (SPT) tree switchover thresholds for group prefixes.</p> <p>PIM-SM routers with directly connected routers receive multicast traffic initially on a shared tree rooted at the Rendezvous Point (RP). Once the traffic arrives on the shared tree and the source of the traffic is known, a switchover to the SPT tree rooted at the source is attempted.</p> <p>For a group that falls in the range of a prefix configured in the table, the corresponding threshold value determines when the router should switch over from the shared tree to the source specific tree. The switchover is attempted only if the traffic rate on the shared tree for the group exceeds the configured threshold.</p> <p>In the absence of any matching prefix in the table, the default behavior is to switchover when the first packet is seen. In the presence of multiple prefixes matching a given group, the most specific entry is used.</p>
Parameters	<p><i>grp-ip-address</i> — The multicast group IP address expressed in dotted decimal notation.</p> <p>Values 224.0.0.0 — 239.255.255.255</p> <p><i>spt-threshold</i> — Specifies the configured threshold in kilobits per second (kbps) for a group prefix. A switchover is attempted only if the traffic rate on the shared tree for the group exceeds this configured threshold.</p> <p>Values 1 — 4294967294 infinity</p>

mask — The mask associated with the IP prefix expressed as a mask length or in dotted decimal notation; for example /16 for a sixteen-bit mask. The mask can also be entered in dotted decimal notation (255.255.0.0).

Values 4 — 32

infinity — When the **infinity** keyword is specified, no switchover will occur at any time, regardless of the traffic level is detected. The threshold, in kilobits per second (KBPS), value is 4294967295.

ssm-groups

Syntax [no] **ssm-groups**

Context config>router>pim

Description This command enables the context to enable an ssm-group configuration instance.

bootstrap-export

Syntax **bootstrap-export** *policy-name* [*..policy-name*]

Context config>router>pim>rp

Description Use this command to apply export policies to control the flow of bootstrap messages from the RP, and apply them to the PIM configuration. Up to 5 policy names can be specified.

Default no bootstrap-export

Parameters *policy-name* — Specify the export policy name up to 32 characters in length.

bootstrap-import

Syntax **bootstrap-import** *policy-name* [*..policy-name*]

Context config>router>pim>rp

Description Use this command to apply import policies to control the flow of bootstrap messages to the RP, and apply them to the PIM configuration. Up to 5 policy names can be specified.

Default no bootstrap-import

Parameters *policy-name* — Specify the import policy name up to 32 characters in length.

hash-mask-len

Syntax	hash-mask-len <i>hash-mask-length</i> no hash-mask-len
Context	config>router>pim>rp>bsr-candidate config>router>pim>rp>ipv6>bsr-candidate
Description	This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.
Parameters	<i>hash-mask-length</i> — The hash mask length. Values 0 — 32

Router Multicast Source Discovery Protocol (MSDP) Commands

msdp

Syntax	[no] msdp
Context	config>router
Description	<p>This command enables a Multicast Source Discovery Protocol (MSDP) instance. When an MSDP instance is created, the protocol is enabled. To start or suspend execution of the MSDP protocol without affecting the configuration, use the [no] shutdown command.</p> <p>The no form of the command deletes the MSDP protocol instance removing all associated configuration parameters.</p>
Default	no msdp
Interactions:	<p>In order for the MSDP protocol to function at least one peer must be configured.</p> <p>When MSDP is configured and started an appropriate event message should be generated.</p> <p>When the no form of the command is executed all sessions must be terminated and an appropriate event message should be generated.</p> <p>When all peering sessions are terminated an event message per peer is not required.</p>

active-source-limit

Syntax	active-source-limit <i>number</i> no active-source-limit
Context	config>router>msdp config>router>msdp>group config>router>msdp>group>peer
Description	<p>This option controls the maximum number of active source messages that will be accepted by Multicast Source Discovery Protocol (MSDP). This effectively controls the number of active sources that can be stored on the system.</p> <p>The no form of this command reverts the number of source message limit to default operation</p>
Default	No limit is placed on the number of source active records
Parameters	<i>number</i> — This parameter defines how many active sources can be maintained by MSDP.
Values	0 — 1000000

receive-msdp-msg-rate

Syntax	receive-msg-rate <i>number interval seconds</i> [threshold <i>number</i>] no receive-msg-rate
Context	config>router>msdp config>router>msdp>peer config>router>msdp>group config>router>msdp>source
Description	<p>This command limits the number of Multicast Source Discovery Protocol (MSDP) messages that are read from the TCP session. It is possible that an MSDP/ RP router may receive a large number of MSDP protocol message packets in a particular source active message.</p> <p>The no form of this command reverts this active-source limit to default operation</p>
Default	No limit is placed on the number of MSDP and source active limit messages will be accepted.
Parameters	<p><i>number</i> — Defines the number of MSDP messages (including source active messages) that are read from the TCP session per the number of seconds.</p> <p>Values 10 — 10000</p> <p>Default 0</p> <p><i>interval seconds</i> — This defines the time that together with the <i>number</i> parameter defines the number of MSDP messages (including source active messages) that are read from the TCP session within the configured number of seconds.</p> <p>Values 1 — 600</p> <p>Default 0</p> <p><i>threshold</i> — This number reflects the number of MSDP messages can be processed before the MSDP message rate limiting function described above is activated; this is of use in particular during at system startup and initialization.</p> <p>Values 1 — 1000000</p> <p>Default 0</p>
Interactions:	Once the number of MSDP packets (including source active messages) defined in the threshold have been processed the rate of all other MSDP packets is rate limited by no longer accepting messages from the TCP session until the time (seconds) has elapsed.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [<i>hash</i> <i>hash2</i>] no authentication-key
Context	config>router>msdp>peer config>router>msdp>group>peer

Description	This command configures a Message Digest 5 (MD5) authentication key to be used with a specific Multicast Source Discovery Protocol (MSDP) peering session. The authentication key must be configured per peer as such no global or group configuration is possible.
Default	Authentication-key. All MSDP messages are accepted and the MD5 signature option authentication key is disabled.
Parameters	<p><i>authentication-key</i> — 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> — 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>hash — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.</p>

data-encapsulation

Syntax	[no] data-encapsulation
Context	config>router>msdp
Description	This command configures a rendezvous point (RP) using Multicast Source Discovery Protocol (MSDP) to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.
Default	data-encapsulation

default-peer

Syntax	default-peer no default-peer
Context	config>router>msdp>peer config>router>msdp>group>peer
Description	Using the default peer mechanism a peer can be selected as the default Multicast Source Discovery Protocol (MSDP) peer, as a result all source-active messages from the peer will be accepted without the usual peer-reverse-path-forwarding (RPF) check.

The MSDP peer-RPF check is different from the normal multicast RPF checks. The peer-RPF check is used to stop source-active messages from looping. A router validates source-active messages originated from other routers in a deterministic fashion.

A set of rules is applied in order to validate received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected. The rules applied to source-active messages originating at Router S received at Router R from Router N are as follows:

- If Router N and router S are one and the same, then the message is originated by a direct peer-RPF neighbor and will be accepted.
- If Router N is a configured peer, or a member of the Router R mesh group then its source-active messages are accepted.
- If Router N is the Border Gateway Protocol (BGP) next hop of the active multicast RPF route toward Router S then Router N is the peer-RPF neighbor and its source-active messages are accepted.
- If Router N is an external BGP peer of Router R and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router N's AS number, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
- If Router N uses the same next hop as the next hop to Router S, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
- If Router N fits none of the above rules, then Router N is not a peer-RPF neighbor, and its source-active messages are rejected.

Default No default peer is established and all active source messages must be RPF checked.

export

Syntax **export** *policy-name* [*policy-name*...(up to 5 max)]
no export

Context config>router>msdp
config>router>msdp>peer
config>router>msdp>group
config>router>msdp>group>peer

Description This command specifies the policies to export source active state from the source active list into Multicast Source Discovery Protocol (MSDP).

Default No export policies are applied and all SA entries are announced.

Interactions: 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.

If you configure an export policy at the global level, each individual peer inherits the global policy. If you configure an export policy at the group level, each individual peer in a group inherits the group's policy. If you configure an export policy at the peer level then policy only applies to the peer where it is configured.

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

group

Syntax	[no] group <i>group-name</i>
Context	config>router>msdp
Description	<p>This command enables access to the context to create or modify a Multicast Source Discovery Protocol (MSDP) group. To configure multiple MSDP groups, include multiple group statements.</p> <p>By default, the group's options are inherited from the global MSDP options. To override these global options, group-specific options within the group statement can be configured.</p> <p>In order for a group to be of use at least one peer must be configured.</p>
Default	no group
Parameters	<i>group-name</i> — Specifies a unique name for the MSDP group.
Interactions:	<p>If the group name provided is already configured then this command only provides the context to configure the options pertaining to this group.</p> <p>If the group name provided is not already configured, then the group name must be created and the context to configure the parameters pertaining to the group should be provided. In this case the \$ prompt to indicate that a new entity (group) is being created should be used.</p>

import

Syntax	import <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no import
Context	config>router>msdp config>router>msdp>peer config>router>msdp>group config>router>msdp>group>peer
Description	This command specifies the policies to import source active state from Multicast Source Discovery Protocol (MSDP) into source active list.
Default	No import policies are applied and all source active messages are allowed.
Interactions:	<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 import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.</p> <p>The no form of the command removes all policies from the configuration.</p> <p>If you configure an import policy at the global level, each individual peer inherits the global policy.</p> <p>If you configure an import policy at the group level, each individual peer in a group inherits the group's policy.</p> <p>If you configure an import policy at the peer level then policy only applies to the peer where it is configured.</p>

local-address

Syntax	local-address <i>address</i> no local-address
Context	config>router>msdp config>router>msdp>peer config>router>msdp>group config>router>msdp>group>peer
Description	<p>This command configures the local end of a Multicast Source Discovery Protocol (MSDP) session. In order for MSDP to function at least one peer must be configured. When configuring a peer, you must include this local-address command to configure the local end of the MSDP session. This address must be present on the node and is used to validate incoming connections to the peer and to establish connections to the remote peer.</p> <p>The no local address format of this command removes the local-address from the configuration.</p>
Default	No local address is configured.
Parameters	<i>address</i> — Specifies an existing address on the node.
Interactions:	<p>If the user enters this command then the address provided is validated and will be used as the local address for MSDP peers from that point. If a subsequent local-address command is entered it will replace the existing configuration and existing session(s) will be terminated.</p> <p>Similarly when the no form of this command is entered the existing local-address will be removed from the configuration and the existing session(s) will be terminated.</p> <p>Whenever a session is terminated all information pertaining to and learned from that peer and will be removed.</p> <p>Whenever a new peering session is created or a peering session is lost an event message should be generated.</p>

mode

Syntax	mode { mesh-group standard }
Context	config>router>msdp>group
Description	<p>This command configures groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers.</p> <p>Multicast Source Discovery Protocol (MSDP) peers can be configured grouped in a full-mesh topology that prevents excessive flooding of source-active messages to neighboring peers.</p>
Default	standard (non-meshed)
Parameters	<p>mesh-group — Specifies that source-active message received from a mesh group member are always accepted but are not flooded to other members of the same mesh group. These source-active messages are only flooded to non-mesh group peers or members of other mesh groups.</p> <p>standard — Specifies a non-meshed mode.</p>

Interactions: In a meshed configuration all members of the group must have a peer connection with every other mesh group member. If this rule is not adhered to then unpredictable results may occur.

peer

Syntax [no] **peer** *peer-address*

Context config>router>msdp
config>router>msdp>group

Description This command configures peer parameters. Multicast Source Discovery Protocol (MSDP) must have at least one peer configured. A peer is defined by configuring a local-address that can be used by this node to set up a peering session and the address of a remote MSDP router. It is the address of this remote peer that is configured in this command and it identifies the remote MSDP router address.

After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. It may be required to have multiple peering sessions in which case multiple peer statements should be included in the configurations.

By default the options applied to a peer are inherited from the global or group-level. To override these inherited options, include peer-specific options within the peer statement.

At least one peer must be configured for MSDP to function.

Default none

Parameters *peer-address* — The address configured in this statement must identify the remote MSDP router that the peering session must be established with.

Interactions: If the peer address provided is already a configured peer then this command only provides the context to configure the parameters pertaining to this peer.

If the peer address provided is not already a configured peer, then the peer instance must be created and the context to configure the parameters pertaining to this peer should be provided. In this case the \$ prompt to indicate that a new entity (peer) is being created should be used.

The peer address provided will be validated and assuming it is valid it will be used as the remote address for an MSDP peering session.. When the no form of this command is entered the existing peering address will be removed from the configuration and the existing session will be terminated. Whenever a session is terminated all source active information pertaining to and learned from that peer and will be removed. Whenever a new peering session is created or a peering session is lost an event message should be generated.

source

Syntax [no] **source** *ip-prefix/mask*

Context config>router>msdp

Description This command limits the number of active source messages the router accepts from sources in the specified address range.

The **no** form of this message removes the source active rate limiter for this source address range.

Default None. The source active **msdp** messages are not rate limited based on the source address range.

Interactions: If the prefix and mask provided is already a configured then this command only provides the context to configure the parameters pertaining to this active source-message filter.

If the prefix and mask provided is not already a configured, then the source node instance must be created and the context to configure the parameters pertaining to this node should be provided. In this case the \$ prompt to indicate that a new entity (source) is being created should be used.

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 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal)

Multicast CAC Policy Configuration Commands

mcac

Parameters	mcac
Context	config>router config>router>pim>if config>router>mld>group-interface
Description	This command enables the context to configure multicast CAC parameters.
Default	none

policy

Syntax	policy <i>mcac-policy-name</i> no policy <i>mcac-policy-name</i>
Context	configure>router>igmp>interface>mcac configure>service>vpn>igmp>interface >mcac
Description	<p>This command references the global channel bandwidth definition policy that is used for (H)mcac and HQoS Adjust.</p> <p>HQoS Adjustment is supported only with redirection enabled. In other words, the policy from the redirected interface is used for HQoS Adjustment.</p> <p>Hierarchical mcac (Hmcac) is supported only with redirection enabled. In Hmcac, the subscriber is checked first against its bandwidth limits followed by the check on the redirected interface against the bandwidth limits defined under the redirected interface. In the Hmcac case the channel definition policy must be referenced under the redirected interface level.</p>
Parameters	<i>mcac-policy-name</i> — Specifies the name of the global mcac channel definition policy defined under the hierarchy configure>router>mcac>policy.
Default	No policy is referenced.

bundle

Parameters	[no] bundle <i>bundle-name</i>
Context	config>router>mcac>policy
Description	This command creates the context that enables the grouping of MCAC group addresses into bundles.

Multicast CAC Policy Configuration Commands

When a number of multicast groups or BTV channels are grouped into a single bundle, then policing, if a join for a particular MC-group (BTV channel), can depend on whether:

1. There is enough physical bandwidth on the egress interface.
2. The given channel is a mandatory or optional channel.
 - If optional, is there sufficient bandwidth according to the policy settings for the relevant interface.
 - If optional, is there sufficient bandwidth within the bundle.

The **no** form of the command removes the named bundle from the configuration.

Default none

Parameters *bundle-name* — Specifies the multicast bundle 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, etc.), the entire string must be enclosed within double quotes.

bw *bandwidth* — Defines the bandwidth available to this bundle when unconstrained.

bandwidth

Syntax **bandwidth** *bandwidth*
no bandwidth

Context config>router>mcac>policy>bundle

Description This command configures the MCAC policy bundle maximum bandwidth.

Parameters *bandwidth* — Specifies the MCAC policy bandwidth.

channel

Syntax **channel** *start-address end-address* **bw** *bandwidth* [**class** {**high** | **low**}] [**type** {**mandatory** | **optional**}] [**source** *source-prefix*]
no channel *start-address end-address* [**source** *source-prefix*]

Context config>router>mcac>policy>bundle

Description This command creates a multicast channel within the bundle where it is configured. A join for a particular multicast channel can be accepted if:

- 1) Mandatory channels:

A sufficient bandwidth exists on the interface according to the policy settings for the interface.
Note, there is always sufficient BW available on the bundle level, as mandatory channels get BW pre-reserved.

- 2) Optional channels:

A sufficient BW exists on both interface and bundle level.

A channel definition can be either IPv4 (*start-address*, *end-address*, *source-address* are IPv4 addresses) or IPv6. A single bundle can have either IPv4 or IPv6 or IPv6 and IPv4 channel definitions. A single policy can mix any of those bundles.

Overlapping channels are not allowed. Two channels overlap if they contain same groups and the same source address prefix (or both do not specify source address prefix). Two channels with same groups and different source prefixes (including one of the channels having no source configured or one of the channels having more specific prefix than the other) do not overlap and are treated as separate channels.

When joining a group from multiple sources, MCAC accounts for that only once when no source address is specified or a prefix for channel covers both sources. Channel BW should be adjusted accordingly or source-aware channel definition should be used if that is not desired.

If a bundle is removed, the channels associated are also removed and every multicast group that was previously policed (because it was in the bundle that contained the policy) becomes free of constraints.

When a new bundle is added to a MCAC policy, the bundle's established groups on a given interfaces are accounted by the policy. Even if this action results in exceeding the bundle's constrain, no active multicast groups are removed. When a leave message is received for an existing optional channel, then the multicast stream is pruned and subsequent new joins may be denied in accordance with the policy. It is possible that momentarily there may be insufficient bandwidth, even for mandatory channels, in this bundle.

Default No channels are specified as part of a bundle on default.

Parameters *start-address end-address* — Specifies the beginning and ending multicast IP addresses that identifies a multicast stream (BTV channel).

Values This must be a valid IPv4 or IPv6 multicast group address

source *source-prefix* — Specifies the source of the multicast IP stream. This must be a valid IPv4 or IPv6 multicast source address prefix.

Values address-prefix/prefix-length

address-prefix is valid IPv4/IPv6 multicast source IP address prefix (local scope excluded)

prefix-length [0..32] for IPv4 [0..128] for IPv6

bw *bandwidth* — Specifies the bandwidth required by this channel in kbps. If this bandwidth is configured for a mandatory channel then this bandwidth is reserved by subtracting the amount from the total available bandwidth for all potential egress interfaces and the bundle.

If this bandwidth is configured as an optional channel then this bandwidth must be available for both the bundle and the egress interface requesting the channel to be added. Once the channel has been added the available bandwidth for the bundle and the interface must be reduced by the configured bandwidth of channel.

Values 10 — 20000 kbps

class {high | low} — Provides deeper classification of channels used in the algorithm when LAG ports change state.

Default low

type {mandatory | optional} — Specifies the channel to be either mandatory or optional.

mandatory — When the **mandatory** keyword is specified, then the bandwidth is reserved by subtracting it from the total available for all the potential egress interfaces and the bundle.

Multicast CAC Policy Configuration Commands

optional — When the **optional** keyword is specified then the bandwidth must be available on both the bundle and the egress interface that requests the channel to be added. Once the channel has been added the available bandwidth for the bundle and the interface must be reduced by the configured bandwidth of channel.

Default optional

mc-constraints

Parameters mc-constraints

Context config>router>mcac>policy>bundle
config>router>igmp>group-interface>mcac
config>router>mld>group-interface

Description This command enables the context to configure the level and its associated bandwidth for a bundle or a logical interface.

Default none

policy

Syntax **policy** *policy-name*
no policy

Context configure>router>igmp>interface>mcac
configure>router>igmp>group-interface>mcac
configure>service>vprn>igmp>interface>mcac
config>router>mld>group-interface
configure>service>vprn>igmp>group-interface>mcac

Description This command references the global channel bandwidth definition policy that is used for (H)mcac and HQoS Adjust.

Within the scope of HQoS Adjustment, the channel definition policy under the group-interface is used if redirection is disabled. In such case HQoS Adjustment can be applied to IPoE subscribers in per-sap replication mode.

In case that redirection is enabled, the channel bandwidth definition policy applied under the Layer 3 redirected interface is in effect.

Hierarchical mcac (Hmcac) is supported on two levels simultaneously:

subscriber level and redirected interface in case that redirection is enabled

subscriber level and group-interface level in case that redirection is disabled.

In Hmcac, the subscriber is first checked against its bandwidth limits followed by the check on the redirected interface (or group-interface) against the bandwidth limits there.

In the case that the redirection is enabled but the policy is referenced ONLY under the group-interface, no admission control will be executed (Hmcac or Mcac).

Default No policy is referenced.

Parameters *policy-name* — Specifies the name of the global mcac channel definition policy defined under the hierarchy **configure>router>mcac>policy**.

lag-port-down

Parameters **lag-port-down** *lag-id* **number-down** *number-lag-port-down* **level** *level-id*
no lag-port-down *lag-id* **number-down** *number-lag-port-down*

Context config>router>mcac>policy>bundle>mc-constraints

Description This command configures the bandwidth available both at the interface and bundle level when a specific number of ports in a LAG group fail.

Default none

Parameters *lag-id* — When the number of ports available in the LAG link is reduced by the number of ports configured in this context then the *level-id* specified here must be applied.

number-down *number-lag-port-down* — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

Values 1 — 64 (for 64-link LAG)
 1 — 32 (for other LAGs)

level *level-id* — Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.

number-down

Parameters **number-down** *number-lag-port-down* **level** *level-id*
no number-down *number-lag-port-down*

Context config>router>pim>if>mcac>mc-constraints

Description This command configures the number of ports down along with level for multicast cac policy on this interface.

Default none

Parameters **number-down** *number-lag-port-down* — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

Values 1 — 64 (for 64-link LAG)
 1 — 32 (for other LAGs)

Multicast CAC Policy Configuration Commands

level *level-id* — Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.

level

Parameters **level** *level* **bw** *bandwidth*
no level *level*

Context config>router>mcac>policy>bundle>mc-constraints

Description This command configures the amount of bandwidth available within a given bundle for MC traffic for a specified level. The amount of allowable BW for the specified level is expressed in kbps and this can be defined for up to eight different levels.

The **no** form of the command removes the level from the configuration.

Default none (If no bandwidth is defined for a given level then no limit is applied.)

Parameters *level* — Specifies the bandwidth for a given level. Level 1 has the highest priority. Level 8 has the lowest priority.

Values 1 — 8

bw *bandwidth* — Specifies the bandwidth, in kbps, for the level.

Values 1 — 2147483647 kbps

Default 1

number-down

Syntax **number-down** *number-lag-port-down* **level** *level-id*
no number-down *number-lag-port-down*

Context config>router>igmp>mcac>mc-constraints

Description This command configures the number of ports down along with level for the MCAC policy.

Parameters *number-lag-port-down* — Specifies the number of ports down along with level for the MCAC policy.

Values 1 — 64

level *level-id* — Specifies the bandwidth for a given level. Level 1 has the highest priority. Level 8 has the lowest priority.

Values 1 — 8

unconstrained-bw

Syntax	unconstrained-bw <i>bandwidth</i> mandatory-bw <i>mandatory-bw</i> no unconstrained-bw
Context	configure>router>igmp>interface>mcac configure>router>igmp>group-interface>mcac configure>service>vprn>igmp>interface >mcac config>router>mld>group-interface>mcac configure>service>vprn>igmp>group-interface >mcac configure>subscr-mgmt>sub-mcac-policy
Description	<p>This command enables Mcac (or Hmcac) function on the corresponding level (subscriber, group-interface or redirected interface). When Mcac (or Hmcac) is enabled and a channel definition policy is referenced, admission control is performed. The allocated bandwidth for optional channels should not exceed the unconstrained-bw minus the mandatory-bw. The mandatory channels have to stay below the specified value for the mandatory-bw.</p> <p>In Hmcac, the subscriber is checked first against its bandwidth limits followed by the check on the redirected interface or the group-interface against the bandwidth limits defined there.</p> <p>In case that redirection is enabled and Hmcac enabled, the channel definition policy must be referenced under the redirected interface level. If it is referenced under the group-interface level, it will be ignored.</p> <p>Subscriber Mcac (only subscriber is checked for available resources) is supported only with direct subscriber replication (no redirection). In this case the channel definition policy must be referenced under the group-interface.</p> <p>In the case that the redirection is enabled but the policy is referenced ONLY under the group-interface, no admission control will be executed (Hmcac or Mcac).</p>
Default	none
Parameters	<p><i>bandwidth</i> — Specifies the unconstrained bandwidth in kbps for the MCAC policy.</p> <p>Values 0 — 2147483647</p> <p><i>mandatory-bw</i> mandatory-bw — Specifies the mandatory bandwidth in kbps for the MCAC policy.</p> <p>Values 0 — 2147483647</p>

default-action

Parameters	default-action {accept discard}
Context	config>router>mcac>policy
Description	<p>This command specifies the action to be applied to multicast streams (channels) when the streams do not match any of the multicast addresses defined in the MCAC policy.</p> <p>When multiple default-action commands are entered, the last command will overwrite the previous command.</p>
Default	discard (all multicast stream not defined in a MCAC policy will be discarded)

Multicast CAC Policy Configuration Commands

Parameters **accept** — Specifies multicast streams (channels) not defined in the MCAC policy will be accepted.
 discard — Specifies multicast streams (channels) not defined in the MCAC policy will be dropped.

shutdown

Parameters **[no] shutdown**

Context **config>router>mcac>policy>bundle**

Description This 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 **no shutdown** command.

When an entity is shutdown, the operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shutdown before they may be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

When a shutdown is performed then all constraints placed on either a bundle or an interface are removed and multicast can potentially take up the full bandwidth of the interface. Furthermore, when a **no shutdown** command is executed then policing of the policy must be in a gradual fashion. No active multicast groups may be removed. When a leave message is received for an optional channel then the multicast stream should be pruned and subsequent new joins can be denied in accordance with the policy. This may mean that for a period of time insufficient bandwidth is available even for mandatory channels.

MLD Commands

mld

Syntax	[no] mld
Context	config>router
Description	This command enables the context to configure Multicast Listener Discovery (MLD) parameters. The no form of the command disables MLD.
Default	no mld

group-interface

Syntax	[no] group-interface <i>ip-int-name</i>
Context	config>router>mld
Description	This command creates and enables the context to configure MLD group interface parameters.

grp-if-query-src-ip

Syntax	grp-if-query-src-ip <i>ipv6-address</i> no grp-if-query-src-ip
Context	config>router>mld>group-interface
Description	This command configures the query source IPv6 address for all group interfaces. The no form of the command removes the IP address.
Default	none

query-src-ip

Syntax	query-src-ip <i>ipv6-address</i> no query-src-ip
Context	config>router>mld>group-interface
Description	This command configures the query source IPv6 address for the group interface. This IP address overrides the source IP address configured at the router level. The no form of the command removes the IPv6 address.

MLD Commands

Default none

Parameters *ipv6-address* — Sets the source IPv4 address for all subscriber's IGMP queries.

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>mld

Description This command enables the context to configure an Multicast Listener Discovery (MLD) interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.

The **no** form of the command deletes the MLD interface. The **shutdown** command in the **config>router>mld>interface** context can be used to disable an interface without removing the configuration for the interface.

Default **no interface** — No interfaces are defined.

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, etc.), 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.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>router>mld>if

Description This command enables router alert checking for MLD messages received on this interface. The no form of the command disables the router alert checking.

Default none

import

Syntax	import <i>policy-name</i> no import
Context	config>router>mld>if
Description	<p>This command specifies the import route policy to be used for determining which membership reports are accepted by the router. Route policies are configured in the config>router>policy-options context.</p> <p>When an import policy is not specified, all the MLD reports are accepted.</p> <p>The no form of the command removes the policy association from the MLD instance.</p>
Default	no import — No import policy specified.
Parameters	<i>policy-name</i> — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the config>router>policy-options context.

max-groups

Syntax	max-groups <i>value</i> no max-groups
Context	config>router>mld>if
Description	<p>This command specifies the maximum number of groups for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.</p>
Default	0, no limit to the number of groups.
Parameters	<p><i>value</i> — Specifies the maximum number of groups for this interface.</p> <p>Values 1 — 16000</p>

query-interval

Syntax	query-interval <i>seconds</i> no query-interval
Context	config>router>mld config>router>mld>if
Description	This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.
Default	125
Parameters	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages.
	Values 2 — 1024

query-last-member-interval

Syntax	query-last-member-interval <i>seconds</i>
Context	config>router>mld config>router>mld>if
Description	This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.
Default	1
Parameters	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent.
	Values 1 — 1024

query-response-interval

Syntax	query-response-interval <i>seconds</i>
Context	config>router>mld config>router>mld>if
Description	This command specifies how long the querier router waits to receive a response to a host-query message from a host.
Default	10
Parameters	<i>seconds</i> — Specifies the the length of time to wait to receive a response to the host-query message from the host.
	Values 1 — 1023

static

Syntax	static
Context	config>router>mld>if
Description	This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.
Default	none

group

Syntax	[no] group <i>ipv6-address</i>
Context	config>router>mld>if>static
Description	<p>This command enables the context to add a static multicast group either as a (*,G) or one or more (S,G) records. Use MLD static group memberships to test multicast forwarding without a receiver host. When MLD static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.</p> <p>When static MLD group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static MLD group entries do not generate join messages toward the RP.</p> <p>The no form of the command removes the IPv6 address from the configuration.</p>
Default	none
Parameters	<i>ipv6-address</i> — Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

source

Syntax	[no] source <i>ipv6-address</i>
Context	config>router>mld>if>static>group config>router>mld>ssm-translate>grp-range
Description	<p>This command specifies an IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the source(s) that the traffic is expected.</p> <p>The source command is mutually exclusive with the specification of individual sources for the same group.</p> <p>The source command, in combination with the group, is used to create a specific (S,G) static group entry.</p> <p>The no form of the command removes the source from the configuration.</p>
Default	none
Parameters	<i>ip-address</i> — Specifies the IPv6 unicast address.

starg

Syntax	[no] starg
Context	config>router>mld>if>static>group
Description	This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified. Use the no form of the command to remove the starg entry from the configuration.
Default	none

subnet-check

Syntax	[no] subnet-check
Context	config>router>mld>interface
Description	This command enables subnet checking for MLD messages received on this interface. All MLD packets with a source address that is not in the local subnet are dropped.
Default	enabled

version

Syntax	version <i>version</i> no version
Context	config>router>mld>if
Description	This command specifies the MLD version. If routers run different versions, they will negotiate the lowest common version of MLD that is supported by hosts on their subnet and operate in that version. For MLD to function correctly, all routers on a LAN should be configured to run the same version of MLD on that LAN.
Default	1
Parameters	<i>version</i> — Specifies the MLD version number. Values 1, 2

robust-count

Syntax	robust-count <i>robust-count</i> no robust-count
Context	config>router>mld
Description	This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.
Default	2
Parameters	<i>robust-count</i> — Specify the robust count value. Values 2 — 10

ssm-translate

Syntax	ssm-translate
Context	config>router>mld
Description	This command enables the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the source command with starg command enabled.

grp-range

Syntax	[no] grp-range <i>start end</i>
Context	config>router>mld>ssm-translate
Description	This command is used to configure group ranges which are translated to SSM (S,G) entries.
Parameters	<i>start</i> — An IP address that specifies the start of the group range. <i>end</i> — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the <i>start</i> value.

source

Syntax [no] **source** *ip-address*

Context config>router>mld>ssm-translate>grp-range

Description This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters *ip-address* — Specifies the IP address that will be sending data.

Operational Commands

mrinfo

Syntax `mrinfo ip-address [router router-name|service]`

Context <GLOBAL>

Description This command is used to display relevant multicast information from the target multicast router. Information displayed includes adjacency information, protocol, metrics, thresholds, and flags from the target multicast router. This information can be used by network operators to determine whether bi-directional adjacencies exist.

Parameters *ip-address* — Specify the IP address of the multicast capable target router should be entered.

router router-name — Specify the router instance that this command applies to.

Default management Base

service — Specify the service instance that this command applies to.

Values 1 — 2147483647

Mrinfo Output Fields — The following table describes the output fields:

Label	Description
General flags	
version	Indicates software version on queried router.
prune	Indicates that router understands pruning.
genid	Indicates that router sends generation IDs.
mtrace	Indicates that the router handles mtrace requests.
Neighbors flags	
1	Metric
0	Threshold (multicast time-to-live)
pim	PIM enabled on interface.
down	Operational status of interface.
disabled	Administrative status of interface.
leaf	No downstream neighbors on interface.
querier	Interface is IGMP querier.
tunnel	Neighbor reached via tunnel.

Operational Commands

```
A:dut-f# mrimfo 10.1.1.2

10.1.1.2 [version 3.0,prune,genid,mtrace]:
10.1.1.2 -> 10.1.1.1 [1/0/pim]
16.1.1.1 -> 0.0.0.0 [1/0/pim/down/disabled]
17.1.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
200.200.200.3 -> 200.200.200.5 [1/0/tunnel/pim]...
```

mstat

Syntax **mstat source** *ip-address* **group** *grp-ip-address* [**destination** *dst-ip-address*] [**hop** *hop*] [**router** *router-name*]**[service]** [**wait-time** *wait-time*]

Context <GLOBAL>

Description This command traces a multicast path from a source to a receiver and displays multicast packet rate and loss information. The **mstat** command adds the capability to show the multicast path in a limited graphic display and provide drops, duplicates, TTLs, and delays at each node. This information is useful to network operators because it identifies nodes with high drop and duplicate counts. Duplicate counts are shown as negative drops.

Parameters **source** *ip-address* — Specify the IP address of the multicast-capable source. This is a unicast address of the beginning of the path to be traced.

group *group-ip-address* — Specify the multicast address that will be used.

destination *dst-ip-address* — Specify the IP address of the unicast destination. If this parameter is omitted, the IP address of the system where the command is entered is used. The destination parameter can also be used to specify a local interface address as the destination address to send the trace query.

Default The default address for the destination address is the incoming IETF format for that (S,G)

hop *hop* — Specify the maximum number of hops that will be traced from the receiver back toward the source.

Values 1 — 255

Default 32 hops (infinity for the DVMRP routing protocol).

router *router-name* — Specify the router instance that this command applies to.

service — Specify the service instance that this command applies to.

Values 1 — 2147483647

wait-time *wait-time* — Specify the number of seconds to wait for the response.

Values 1 — 60

Default 10

Mstat Output Fields — The following table describes the output fields:

Label	Description
hop	Number of hops from the source to the listed router.
router name	Name of the router for this hop or “?” when not reverse DNS translated.
address	Address of the router for this hop.
protocol	Protocol used.
ttl	Forward TTL threshold. TTL that a packet is required to have before it will be forwarded over the outgoing interface.
forwarding code	Forwarding information/error code for this hop.

For each interface between 2 nodes a line is printed, following the same layout as other routers with an implementation derived from mroute. Note the following:

- The forwarding information/error code is only displayed when different from “No Error”.
- “?” means the there is no reverse DNS translation.
- There is no “Overall Mcast Pkt Rate” available in the PE for the VPRN case.

Operational Commands

Source	Response Dest	Overall	Packet Statistics For Traffic From		
10.10.16.9	10.20.1.6	Mcast Pkt	10.10.16.9 To 224.5.6.7		
	rtt 29 ms	Rate	Lost/Sent	= Pct	Rate
v	/	-----	-----		
10.10.16.3					
10.10.2.3	?				
	^	ttl 2	1 pps	0/0	= -- 0 pps
v					
10.10.2.1					
10.10.1.1	?				
	^	ttl 3	0 pps	0/0	= -- 0 pps
v					
10.10.1.2					
10.10.4.2	?	Reached RP/Core			
	^	ttl 4	0 pps	0/0	= -- 0 pps
v					
10.10.4.4					
10.10.6.4	?				
	^	ttl 5	0 pps	0/0	= -- 0 pps
v					
10.10.6.5					
10.10.10.5	?				
	\	ttl 6	0 pps	0/0	= -- 0 pps
v	\				
10.10.10.6	10.20.1.6				
Receiver	Query Source				

mtrace

Syntax	mtrace source <i>ip-address</i> group <i>grp-ip-address</i> [destination <i>dst-ip-address</i>] [hop <i>hop</i>] [router <i>router-name service</i>] [wait-time <i>wait-time</i>]
Context	<GLOBAL>
Description	This command traces the multicast path from a source to a receiver by passing a trace query hop-by-hop along the reverse path from the receiver to the source. At each hop, information such as the hop address, routing error conditions, and packet statistics are gathered and returned to the requestor. A network administrator can determine where multicast flows stop and verify the flow of the multicast stream.
Parameters	source <i>ip-address</i> — Specify the IP address of the multicast-capable source. This is a unicast address of the beginning of the path to be traced. group <i>group-ip-address</i> — Specify the multicast address that will be used. destination <i>dst-ip-address</i> — Specify the IP address of the unicast destination. If this parameter is omitted, the IP address of the system where the command is entered is used. The destination parameter can also be used to specify a local interface address as the destination address to send the trace query. Default The default address for the desination address is the incoming IETF format for that (S,G)

hop *hop* — Specify the maximum number of hops that will be traced from the receiver back toward the source.

Values 1 — 255

Default 32 hops (infinity for the DVMRP routing protocol).

router *router-name* — Specify the router instance that this command applies to.

service — Specify the service instance that this command applies to.

Values 1 — 2147483647

wait-time *wait-time* — Specify the number of seconds to wait for the response.

Values 1 — 60

Default 10

Mtrace Output Fields — The following table describes the output fields:

Label	Description
hop	Number of hops from the source to the listed router.
router name	Name of the router for this hop. If a DNS name query is not successful a “?” displays.
address	Address of the router for this hop.
protocol	Protocol used.
tth	Forward TTL threshold. TTL that a packet is required to have before it will be forwarded over the outgoing interface.
forwarding code	Forwarding information/error code for this hop.

```
A:Dut-F# mtrace source 10.10.16.9 group 224.5.6.7
```

```
Mtrace from 10.10.16.9 via group 224.5.6.7
Querying full reverse path...
```

```
0 ? (10.10.10.6)
-1 ? (10.10.10.5) PIM thresh^ 1 No Error
-2 ? (10.10.6.4) PIM thresh^ 1 No Error
-3 ? (10.10.4.2) PIM thresh^ 1 Reached RP/Core
-4 ? (10.10.1.1) PIM thresh^ 1 No Error
-5 ? (10.10.2.3) PIM thresh^ 1 No Error
-6 ? (10.10.16.9)
```

```
Round trip time 29 ms; total ttl of 5 required.
```

Show Commands

IGMP Commands

group

Syntax **group** [*grp-ip-address*]
group summary

Context show>router>igmp

Description This command displays the multicast group and (S,G) addresses. If no *grp-ip-address* parameters are specified, then all IGMP group, (*,G) and (S,G) addresses are displayed.

Parameters *grp-ip-address* — Displays specific multicast group addresses.

Output **IGMP Group Output** — The following table describes the output fields for IGMP group information.

Label	Description
IGMP Groups	Displays the IP multicast sources corresponding to the IP multicast groups that are statically configured.
Fwd List	Displays the list of interfaces in the forward list.
Blk List	Displays the list of interfaces in the bulk list.

Sample Output

```
*B:Dut-C# show router igmp group
=====
IGMP Interface Groups
=====
IGMP Host Groups
=====
(*,225.0.0.1)
  Fwd List   : 112.112.1.2           Up Time : 0d 00:00:21
(11.11.0.1,225.0.0.1)
  Fwd List   : 112.112.1.1           Up Time : 0d 00:00:30
  Blk List   : 112.112.1.2           Up Time : 0d 00:00:21
(11.11.0.2,225.0.0.1)
  Fwd List   : 112.112.1.1           Up Time : 0d 00:00:30
(*,225.0.0.2)
  Fwd List   : 112.112.1.2           Up Time : 0d 00:00:21
(11.11.0.1,225.0.0.2)
  Blk List   : 112.112.1.2           Up Time : 0d 00:00:21
-----
(*,G)/(S,G) Entries : 5
=====
```

IGMP Commands

```
*B:Dut-C#

*B:Dut-C# show router igmp group summary
=====
IGMP Interface Groups
=====
IGMP Host Groups Summary          Nbr Fwd Hosts          Nbr Blk Hosts
=====
(*,225.0.0.1)                     1                      0
(11.11.0.1,225.0.0.1)             1                      1
(11.11.0.2,225.0.0.1)             1                      0
(*,225.0.0.2)                     1                      0
(11.11.0.1,225.0.0.2)             0                      1
-----
(*,G)/(S,G) Entries : 5
=====
*B:Dut-C#

A:NYC# show router igmp group 224.24.24.24
=====
IGMP Groups
=====
(*,224.24.24.24)                  Up Time : 0d 05:23:23
    Fwd List : nyc-vlc
-----
(*,G)/(S,G) Entries : 1
=====
A:NYC#
```

hosts

Syntax **hosts** [group *grp-address*] [detail] [fwd-service *service-id*] [grp-interface *ip-int-name*]
hosts [host *ip-address*] [group *grp-address*] [detail]
hosts summary

Context show>router>igmp

Description This command shows IGMP hosts information.

Sample Output

```
*B:Dut-C# show router igmp hosts
=====
IGMP Hosts
=====
Host          Oper   Oper   Fwd   GrpItf          Num   Subscriber
              State  Version Svc                    Groups
-----
112.112.1.1   Up     3      1     gi_1_1          1     sub_1
112.112.1.2   Up     3      1     gi_1_1          2     sub_1
112.112.1.3   Up     3      1     gi_1_2          0     sub_2
-----
Hosts : 3
=====
```

*B:Dut-C#

*B:Dut-C# show router igmp hosts detail

=====

IGMP Host 112.112.1.1

=====

Oper Status	: Up	MacAddress	: 00:00:00:00:00:01
Oper version	: 3	Subscriber	: sub_1
Num Groups	: 1	GrpItf	: gi_1_1
Max Grps Till Now:	2	IGMP-Policy	: poll
PPPoE SessionId	: 1	Next query time:	0d 00:02:03
FwdSvcId	: 1		

IGMP Group

Group Address	: 225.0.0.1	Up Time	: 0d 00:00:24
Expires	: Not running	Mode	: Include
V1 Host Timer	: Not running	Type	: Dynamic
V2 Host Timer	: Not running	Compat Mode:	IGMP Version 3
Redir.vRtrId	: N/A	Redir.Intf	: N/A

Source Address	Expires	Type	Fwd/Blk
11.11.0.1	0d 00:03:56	Dynamic	Fwd
11.11.0.2	0d 00:03:56	Dynamic	Fwd

=====

IGMP Host 112.112.1.2

=====

Oper Status	: Up	MacAddress	: 00:00:00:00:00:01
Oper version	: 3	Subscriber	: sub_1
Num Groups	: 2	GrpItf	: gi_1_1
Max Grps Till Now:	2	IGMP-Policy	: poll
PPPoE SessionId	: 2	Next query time:	0d 00:02:03
FwdSvcId	: 1		

IGMP Group

Group Address	: 225.0.0.1	Up Time	: 0d 00:00:16
Expires	: 0d 00:04:05	Mode	: Exclude
V1 Host Timer	: Not running	Type	: Dynamic
V2 Host Timer	: Not running	Compat Mode:	IGMP Version 3
Redir.vRtrId	: N/A	Redir.Intf	: N/A

Source Address	Expires	Type	Fwd/Blk
11.11.0.1	0d 00:00:00	Dynamic	Blk

IGMP Group

Group Address	: 225.0.0.2	Up Time	: 0d 00:00:16
Expires	: 0d 00:04:04	Mode	: Exclude
V1 Host Timer	: Not running	Type	: Dynamic
V2 Host Timer	: Not running	Compat Mode:	IGMP Version 3
Redir.vRtrId	: N/A	Redir.Intf	: N/A

Source Address	Expires	Type	Fwd/Blk
11.11.0.1	0d 00:00:00	Dynamic	Blk

IGMP Commands

```
=====
IGMP Host 112.112.1.3
=====
Oper Status      : Up           MacAddress       : 00:00:00:00:00:02
Oper version     : 3            Subscriber      : sub_2
Num Groups       : 0            GrpItf         : gi_1_2
Max Grps Till Now: 1            IGMP-Policy    : poll
PPPoE SessionId  : 1            Next query time: 0d 00:00:48
FwdSvcId         : 1
-----
Hosts : 3
=====
*B:Dut-C#

*B:Dut-C# show router igmp statistics host 112.112.1.1
=====
IGMP Host Statistics 112.112.1.1
=====
Message Type      Received      Transmitted
-----
Queries           0              580
Report V1         0              0
Report V2         0              0
Report V3         5              0
Leaves            0              0
-----
General Host Statistics
-----
Bad Length        : 0
Bad Checksum      : 0
Unknown Type      : 0
Bad Receive If    : 0
Rx Non Local      : 0
Rx Wrong Version  : 0
Policy Drops      : 0
No Router Alert   : 0
Rx Bad Encodings  : 0
Local Scope Pkts  : 0
Resvd Scope Pkts : 0
MCAC Policy Drops : 0
-----
Source Group Statistics
-----
(S,G)             : 0
(*,G)             : 0
=====
*B:Dut-C# show subscriber-mgmt igmp-policy
```

ssm-translate

Syntax **ssm-translate**
ssm-translate interface *interface-name*

Context show>router>igmp

Description This command displays IGMP SSM translate configuration information.

Output **GMP Interface Output** — The following table provides IGMP field descriptions.

Label	Description
Group Range	Displays the address ranges of the multicast groups for which this router can be an RP.
Source	Displays the unicast address that sends data on an interface.
SSM Translate Entries	Displays the total number of SSM translate entries.

Sample Output

```
=====
IGMP SSM Translate Entries
=====
Group Range                Source                Interface
-----
<234.1.1.1 - 234.1.1.2>    100.1.1.1            -
<232.1.1.1 - 232.1.1.5>    100.1.1.2            ies-abc
-----
```

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-address*] [**detail**]

Context show>router>igmp

Description This command displays IGMP interface information.

Parameters

- ip-int-name* — Only displays the information associated with the specified IP interface name.
- ip-address* — Only displays the information associated with the specified IP address.
- group** *grp-address* — Only displays IP multicast group address for which this entry contains information.
- detail** — Displays detailed IP interface information along with the source group information learned on that interface.

Output **IGMP Interface Output** — The following table provides IGMP field descriptions.

Label	Description
Interface	Specifies the interfaces that participate in the IGMP protocol.
Adm Admin Status	Displays the administrative state for the IGMP protocol on this interface.
Oper Oper Status	Displays the current operational state of IGMP protocol on the interface.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Querier Up Time	Displays the time since the querier was last elected as querier.
Querier Expiry Timer	Displays the time remaining before the querier ages out. If the querier is the local interface address, the value will be zero.
Cfg/Opr Version Admin/Oper version	Cfg — The configured version of IGMP running on this interface. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN. Opr — The operational version of IGMP running on this interface. If the cfg value is 3 but all of the routers in the local subnet of this interface use IGMP version v1 or v2, the operational version will be v1 or v2.
Num Groups	The number of multicast groups which have been learned by the router on the interface.
Policy	Specifies the policy that is to be applied on the interface.
Group Address	Specifies the IP multicast group address for which this entry contains information.
Up Time	Specifies the time since this source group entry got created.
Last Reporter	Specifies the IP address of the source of the last membership report received for this IP Multicast group address on this interface. If no membership report has been received, this object has the value 0.0.0.0.
Mode	The mode is based on the type of membership report(s) received on the interface for the group. In the 'include' mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In 'exclude' mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter.

Label	Description (Continued)
V1 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
V2 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 Membership Report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 Leave messages for this group that it receives on this interface.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, it will be set to "dynamic". For statically configured groups, the value will be set to 'static'.
Compat Mode	Used in order for routers to be compatible with older version routers. IGMPv3 hosts MUST operate in version 1 and version 2 compatibility modes. IGMPv3 hosts MUST keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the Host Compatibility Mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of General Queries heard on that interface as well as the Older Version Querier Present timers for the interface.

Sample Output

```
*A:ALA-BA# show router 100 interface
=====
Interface Table (Service: 100)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
IP-Address
-----
IGMP_to_CE          Up        Up           VPRN      1/1/7
11.1.1.1/24         n/a
system              Up        Up           VPRN      loopback
10.20.1.2/32        n/a
-----
Interfaces : 2
=====
*A:ALA-BA#

*A:ALA-BA# show router 100 interface IGMP_to_CE
=====
Interface Table (Service: 100)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
```

IGMP Commands

```

      IP-Address                                     PfxState
-----
IGMP_to_CE                Up      Up      VPRN    1/1/7
      11.1.1.1/24                                     n/a
-----
Interfaces : 1
=====
*A:ALA-BA#

*A:ALA-BA# show router 100 igmp interface
=====
IGMP Interfaces
=====
Interface                Adm  Oper  Querier      Cfg/Opr Num    Policy
                        Version Groups
-----
IGMP_to_CE                Up   Up    11.1.1.1     1/1      3      igmppol
-----
Interfaces : 1
=====
*A:ALA-BA#

*A:ALA-BA# show router 100 igmp interface IGMP_to_CE
=====
IGMP Interface IGMP_to_CE
=====
Interface                Adm  Oper  Querier      Cfg/Opr Num    Policy
                        Version Groups
-----
IGMP_to_CE                Up   Up    11.1.1.1     1/1      3      igmppol
-----
Interfaces : 1
=====
*A:ALA-BA#

*A:ALA-BA# show router 100 igmp interface 11.1.1.1
=====
IGMP Interface 11.1.1.1
=====
Interface                Adm  Oper  Querier      Cfg/Opr Num    Policy
                        Version Groups
-----
IGMP_to_CE                Up   Up    11.1.1.1     1/1      3      igmppol
-----
Interfaces : 1
=====
*A:ALA-BA#
```



```

*A:ALA-BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1
=====
IGMP Interface IGMP_to_CE
=====
Interface                Adm  Oper Querier                Cfg/Opr Num    Policy
                          Up   Up   11.1.1.1                1/1    3      igmppol
-----
IGMP Group
-----
Group Address : 227.1.1.1          Up Time       : 0d 00:03:52
Interface     : IGMP_to_CE        Expires       : never
Last Reporter : 0.0.0.0           Mode          : exclude
V1 Host Timer : Not running       Type          : static
V2 Host Timer : Not running       Compat Mode   : IGMP Version 3
-----
Interfaces : 1
=====
*A:ALA-BA#

*A:ALA-BA# show router 100 igmp interface IGMP_to_CE group 227.1.1.1 detail
=====
IGMP Interface IGMP_to_CE
=====
Interface      : IGMP_to_CE
Admin Status   : Up
Querier        : 11.1.1.1
Querier Expiry Time: N/A
Admin/Oper version : 1/1
Policy         : igmppol
Max Groups Allowed : 16000
MCAC Policy Name :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Oper Status    : Up
Querier Up Time : 0d 00:04:01
Time for next query: 0d 00:13:42
Num Groups     : 3
Subnet Check   : Disabled
Max Groups Till Now: 3
MCAC Const Adm St : Enable
MCAC Max Mand BW : no limit
MCAC Avail Mand BW : unlimited
MCAC Avail Opnl BW : unlimited
-----
IGMP Group
-----
Group Address : 227.1.1.1          Up Time       : 0d 00:04:02
Interface     : IGMP_to_CE        Expires       : never
Last Reporter : 0.0.0.0           Mode          : exclude
V1 Host Timer : Not running       Type          : static
V2 Host Timer : Not running       Compat Mode   : IGMP Version 3
-----
Interfaces : 1
=====
*A:ALA-BA#

```

static

- Syntax** **static** [*ip-int-name* | *ip-addr*]
- Context** show>router>igmp
- Description** This command displays static IGMP, (*,G) and (S,G) information.
- Parameters** *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-addr — Only displays the information associated with the specified IP address.
- Output** **Static IGMP Output** — The following table provides static IGMP field descriptions.

Label	Description
Source	Displays entries which represents a source address from which receivers are interested/not interested in receiving multicast traffic.
Group	Displays the IP multicast group address for which this entry contains information.
Interface	Displays the interface name.

Sample Output

```
*A:ALA-BA# show router 100 igmp static
=====
IGMP Static Group Source
=====
Source          Group          Interface
-----
11.11.11.11     226.136.22.3   IGMP_to_CE
*               227.1.1.1      IGMP_to_CE
22.22.22.22     239.255.255.255 IGMP_to_CE
-----
Static (*,G)/(S,G) Entries : 3
=====
*A:ALA-BA#
```

statistics

- Syntax** **statistics** [*ip-int-name* | *ip-address*]
statistics host [*ip-address*]
- Context** show>router>igmp
- Description** This command displays IGMP statistics information.
- Parameters** *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-address — Only displays the information associated with the specified IP address.

Output **IGMP Statistics Output** — The following table provides statistical IGMP field descriptions.

Label	Description
IGMP Interface Statistics	The section listing the IGMP statistics for a particular interface.
Message Type	<p>Queries — The number of IGMP general queries transmitted or received on this interface.</p> <p>Report — The total number of IGMP V1, V2, or V3 reports transmitted or received on this interface.</p> <p>Leaves — The total number of IGMP leaves transmitted on this interface.</p>
Received	Displays the total number of IGMP packets received on this interface.
Transmitted	Column that displays the total number of IGMP packets transmitted from this interface.
General Interface Statistics	The section listing the general IGMP statistics.
Bad Length	Displays the total number of IGMP packets with bad length received on this interface.
Bad Checksum	Displays the total number of IGMP packets with bad checksum received on this interface.
Unknown Type	Displays the total number of IGMP packets with unknown type received on this interface.
Bad Receive If	Displays the total number of IGMP packets incorrectly received on this interface.
Rx Non Local	Displays the total number of IGMP packets received from a non-local sender.
Rx Wrong Version	Displays the total number of IGMP packets with wrong versions received on this interface.
Policy Drops	Displays the total number of times IGMP protocol instance matched the host IP address or group/source addresses specified in the import policy.
No Router Alert	Displays the total number of IGMPv3 packets received on this interface which did not have the router alert flag set.

Sample Output

```
*A:ALA-BA# show router 100 igmp statistics
=====
IGMP Interface Statistics
=====
Message Type           Received      Transmitted
-----
Queries                0             5
Report V1              0             0
Report V2              0             0
Report V3              0             0
Leaves                 0             0
-----
General Interface Statistics
-----
Bad Length             : 0
Bad Checksum           : 0
Unknown Type          : 0
Bad Receive If         : 0
Rx Non Local           : 0
Rx Wrong Version       : 0
Policy Drops           : 0
No Router Alert        : 0
Rx Bad Encodings       : 0
Rx Pkt Drops           : 0
-----
Source Group Statistics
-----
(S,G)                  : 2
(*,G)                  : 1
=====
*A:ALA-BA#
```

```
*B:Dut-C# show router igmp statistics host
=====
IGMP Host Statistics
=====
Message Type           Received      Transmitted
-----
Queries                0            1739
Report V1              0             0
Report V2              0             0
Report V3              10             0
Leaves                 0             0
-----
General Host Statistics
-----
Bad Length             : 0
Bad Checksum           : 0
Unknown Type          : 0
Bad Receive If         : 0
Rx Non Local           : 0
Rx Wrong Version       : 0
Policy Drops           : 0
No Router Alert        : 0
Rx Bad Encodings       : 0
Local Scope Pkts       : 0
```

```

Resvd Scope Pkts : 0
MCAC Policy Drops : 0
=====
*B:Dut-C#

```

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.
If IGMP is not enabled, the following message appears:

```

A:NYC# show router igmp status
MINOR: CLI IGMP is not configured.
A:NYC#

```

Output **IGMP Status Output** — The following table provides IGMP status field descriptions.

Label	Description
Admin State	Displays the administrative status of IGMP.
Oper State	Displays the current operating state of this IGMP protocol instance on this router.
Query Interval	The frequency at which IGMP query packets are transmitted.
Last Member Query Interval	The maximum response time inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages.
Query Response Interval	The maximum query response time advertised in IGMPv2 queries.
Robust Count	Displays the number of times the router will retry a query.

Sample Output

```

*A:ALA-BA# show router 100 igmp status
=====
IGMP Status
=====
Admin State           : Up
Oper State            : Up
Query Interval        : 1024
Last Member Query Interval : 1024
Query Response Interval : 1023
Robust Count          : 10
=====
*A:ALA-BA#

```

Show Router PIM Commands

anycast

- Syntax** **anycast [detail]**
- Context** show>router>pim
- Description** This command displays PIM anycast rp-set information.
- Parameters** **detail** — Displays detailed information.
- Output** **PIM anycast Output —** The following table provides PIM anycast field descriptions

Label	Description
Anycast Address	Displays the candidate anycast address.
Anycast RP Peer	Displays the candidate anycast RP peer address.

Sample Output

```
A:dut-d# show router pim anycast
=====
PIM Anycast RP Entries
=====
Anycast RP           Anycast RP Peer
-----
100.100.100.1        102.1.1.1
                     103.1.1.1
                     104.1.1.1
-----
PIM Anycast RP Entries : 3
=====
```

crp

- Syntax** **crp [ip-address]**
- Context** show>router>pim
- Description** Display PIM candidate RP (CRP) information received at the elected Bootstrap router (BSR).
- Parameters** *ip-address* — The candidate RP IP address.

Output **PIM CRP Output** — The following table provides PIM CRP field descriptions.

Label	Description
RP Address	Displays the Candidate RP address.
Group Address	Displays the range of multicast group addresses for which the CRP is the Candidate RP.
Priority	Displays the Candidate RP's priority for becoming a rendezvous point (RP). This value is used to elect RP for a group range. A value of 0 is considered as the highest priority.
Holdtime	Displays the hold time of the candidate RP. It is used by the Bootstrap router to time out the RP entries if it does not listen to another CRP advertisement within the holdtime period.
Expiry	The minimum time remaining before the CRP will be declared down. If the local router is not the BSR, this value is 0.
Candidate RPs	Displays the number of CRP entries.

Sample Output

```
A:WAS# show router pim crp
=====
PIM Candidate RPs
=====
RP Address      Group Address    Priority    Holdtime    Expiry Time
-----
2.22.187.236    224.0.0.0/4      192        150         0d 00:02:19
2.22.187.239    224.0.0.0/4      192        150         0d 00:02:19
2.22.187.240    224.0.0.0/4      192        150         0d 00:02:09
-----
Candidate RPs : 3
=====
A:WAS#

A:WAS# show router pim crp 2.22.187.236
=====
PIM Candidate RPs
=====
RP Address      Group Address    Priority    Holdtime    Expiry Time
-----
2.22.187.236    224.0.0.0/4      192        150         0d 00:01:43
-----
Candidate RPs : 1
=====
A:WAS#
```

s-psmi

- Syntax

s-psmi [mdSrcAddr [mdGrpAddr]] [detail]
- Context

show>router>pim
- Description

Displays the list of selective provider multicast service interfaces that are currently active.
- Parameters

mdSrcAddr — Specifies the source address of the multicast sender.

mdGrpAddr — Specifies the group address of the multicast sender.

detail — Displays detailed output.
- Output

PIM data MDT Output — The following table provides PIM data MDT descriptions.

Label	Description
MD Grp Address	Displays the IP multicast group address for which this entry contains information.
MD Src Address	Displays the source address of the multicast sender. It will be 0 if the type is configured as starg . It will be the address of the Rendezvous Point (RP) if the type is configured as starRP .
MT Index	Displays the index number.
Num VP SGs	Displays the VPN number.

Sample Output PIM Selective Provider Tunnel

```
*B:node-6# show router 100 pim s-psmi
=====
PIM Selective provider tunnels
=====
MD Src Address      MD Grp Address      MT Index      Num VPN SGs
-----
200.200.200.7       230.0.89.72         24603         1
200.200.200.7       230.0.89.73         24604         1
200.200.200.7       230.0.89.74         24605         1
200.200.200.7       230.0.89.75         24606         1
200.200.200.7       230.0.89.76         24607         1
200.200.200.7       230.0.89.77         24608         1
200.200.200.7       230.0.89.78         24609         1
200.200.200.7       230.0.89.79         24610         1
200.200.200.7       230.0.89.80         24611         1
200.200.200.7       230.0.89.81         24612         1
200.200.200.7       230.0.89.82         24613         1
200.200.200.7       230.0.89.83         24614         1
200.200.200.7       230.0.89.84         24615         1
200.200.200.7       230.0.89.85         24616         1
200.200.200.7       230.0.89.86         24617         1
200.200.200.7       230.0.89.87         24618         1
...
=====
*B:node-6#
```


Sample Output PIM Selective Provider Tunnel Detail

```

*B:node-6# show router 100 pim s-pmsi detail
=====
PIM Selective provider tunnels
=====
Md Source Address   : 200.200.200.7      Md Group Address   : 230.0.89.72
Number of VPN SGs   : 1                  Uptime             : 0d 00:00:18
MT IfIndex           : 24603              Egress Fwding Rate : 163.2 kbps

VPN Group Address   : 228.1.0.0          VPN Source Address : 11.2.102.1
State               : RX Joined
Expiry Timer        : 0d 00:02:41
=====
PIM Selective provider tunnels
=====
Md Source Address   : 200.200.200.7      Md Group Address   : 230.0.89.73
Number of VPN SGs   : 1                  Uptime             : 0d 00:00:18
MT IfIndex           : 24604              Egress Fwding Rate : 163.2 kbps

VPN Group Address   : 228.1.0.1          VPN Source Address : 11.2.102.1
State               : RX Joined
Expiry Timer        : 0d 00:02:41
=====
PIM Selective provider tunnels
=====
Md Source Address   : 200.200.200.7      Md Group Address   : 230.0.89.74
Number of VPN SGs   : 1                  Uptime             : 0d 00:00:20
MT IfIndex           : 24605              Egress Fwding Rate : 165.7 kbps

VPN Group Address   : 228.1.0.2          VPN Source Address : 11.2.102.1
State               : RX Joined
Expiry Timer        : 0d 00:02:39
=====
PIM Selective provider tunnels
=====
Md Source Address   : 200.200.200.7      Md Group Address   : 230.0.89.75
Number of VPN SGs   : 1                  Uptime             : 0d 00:00:20
MT IfIndex           : 24606              Egress Fwding Rate : 165.7 kbps

VPN Group Address   : 228.1.0.3          VPN Source Address : 11.2.102.1
State               : RX Joined
Expiry Timer        : 0d 00:02:39
=====
*B:node-6#

```

Sample Output RX Tracking for RSVP S-PMSI Tunnel

```

*A:Dut-C# show router 1 pim s-pmsi
=====
PIM RSVP Spmsi tunnels
=====
P2mp ID   Tunnel ID   Ext Tunnel Adrs      SPMSI Index   Num VPN   State
                               SGs
-----

```

Show Router PIM Commands

```
0          0          10.20.1.4          1030144          1          RX Tracking
0          0          10.20.1.4          1030144          1          RX Tracking
=====
PIM RSVP Spmsi Interfaces : 2
=====
*A:Dut-C# show router 21 pim s-pmsi

=====
PIM LDP Spmsi tunnels
=====
Lsp ID      Root Addr          SPMSI Index    Num VPN    State
              SGs
-----
0          10.20.1.4          1030144        1          RX Tracking
0          10.20.1.4          1030144        1          RX Tracking
=====
PIM LDP Spmsi Interfaces : 2
=====
*A:Dut-C#
```

Sample Output RX Tracking for RSVP S-PMSI Tunnel Detail

```
*A:Dut-C# show router 1 pim s-pmsi detail

=====
PIM RSVP Spmsi tunnels
=====
P2MP ID      : 0          Tunnel ID      : 0
Ext Tunnel Addr : 10.20.1.4    Spmsi IfIndex : 1030144
Number of VPN SGs : 1          Uptime        : 0d 00:02:48

VPN Group Address : 225.100.0.0
VPN Source Address : 10.1.101.2
State            : RX Tracking    Mdt Threshold : 0

=====
PIM RSVP Spmsi tunnels
=====
P2MP ID      : 0          Tunnel ID      : 0
Ext Tunnel Addr : 10.20.1.4    Spmsi IfIndex : 1030144
Number of VPN SGs : 1          Uptime        : 0d 00:02:47

VPN Group Address : ff0e:225:100::
VPN Source Address : 2001:10:1:101::2
State            : RX Tracking    Mdt Threshold : 0

=====
PIM RSVP Spmsi Interfaces : 2
=====
*A:Dut-C# show router 21 pim s-pmsi detail

=====
PIM LDP Spmsi tunnels
=====
LSP ID      : 0
Root Addr   : 10.20.1.4          Spmsi IfIndex : 1030144
Number of VPN SGs : 1          Uptime        : 0d 00:03:35
```

```

VPN Group Address : 225.100.0.0
VPN Source Address : 10.1.101.2
State              : RX Tracking      Mdt Threshold      : 0

```

```

=====
PIM LDP Spmsi tunnels
=====

```

```

LSP ID           : 0
Root Addr        : 10.20.1.4          Spmsi IfIndex      : 1030144
Number of VPN SGs : 1                  Uptime             : 0d 00:03:34

```

```

VPN Group Address : ff0e:225:100::
VPN Source Address : 2001:10:1:101::2
State              : RX Tracking      Mdt Threshold      : 0

```

```

=====
PIM LDP Spmsi Interfaces : 2
=====

```

```

*A:Dut-C#

```

Sample Output TX Tracking for RSVP S-PMSI Tunnel Detail

```

*A:Dut-C# show router 1 pim s-pmsi detail

```

```

=====
PIM RSVP Spmsi tunnels
=====

```

```

P2MP ID          : 1                  Tunnel ID          : 61442
Ext Tunnel Addr   : 10.20.1.4          Spmsi IfIndex      : 74230
Number of VPN SGs : 1                  Uptime             : 0d 00:05:11

```

```

VPN Group Address : 225.100.0.0
VPN Source Address : 10.1.101.2
State              : TX Join Pending    Mdt Threshold      : 1
Join Timer         : N/A                 Holddown Timer     : 0d 00:00:47
Receiver Count     : 4

```

```

=====
PIM RSVP Spmsi tunnels
=====

```

```

P2MP ID          : 1                  Tunnel ID          : 61443
Ext Tunnel Addr   : 10.20.1.4          Spmsi IfIndex      : 74231
Number of VPN SGs : 1                  Uptime             : 0d 00:05:10

```

```

VPN Group Address : ff0e:225:100::
VPN Source Address : 2001:10:1:101::2
State              : TX Join Pending    Mdt Threshold      : 1
Join Timer         : N/A                 Holddown Timer     : 0d 00:00:50
Receiver Count     : 4

```

```

=====
PIM RSVP Spmsi Interfaces : 2
=====

```

```

*A:Dut-D# show router 21 pim s-pmsi detail

```

```

=====
PIM LDP Spmsi tunnels
=====

```

```

LSP ID           : 8194

```

Show Router PIM Commands

```
Root Addr       : 10.20.1.4           Spmsi IfIndex   : 74228
Number of VPN SGs : 1                 Uptime         : 0d 00:05:56

VPN Group Address : 225.100.0.0
VPN Source Address : 10.1.101.2
State             : TX Join Pending    Mdt Threshold   : 1
Join Timer        : N/A               Holddown Timer   : 0d 00:00:02
Receiver Count    : 4
```

```
=====
PIM LDP Spmsi tunnels
=====
LSP ID          : 8195
Root Addr       : 10.20.1.4           Spmsi IfIndex   : 74229
Number of VPN SGs : 1                 Uptime         : 0d 00:05:55

VPN Group Address : ff0e:225:100::
VPN Source Address : 2001:10:1:101::2
State            : TX Join Pending    Mdt Threshold   : 1
Join Timer       : N/A               Holddown Timer   : 0d 00:00:05
Receiver Count    : 4
=====
PIM LDP Spmsi Interfaces : 2
=====
*A:Dut-D#
```

group

Syntax **group** *grp-ip-address* [**source** *ip-address* [**type** {**starstarrp** | **starg** | **sg**}] [**detail**] [*family*]

Context show>router>pim

Description This command displays PIM source group database information.

Parameters *grp-ip-address* — Specifies the IP multicast group address for which this entry contains information.
source ip-address — Specifies the source address for which this entry contains information.
type starstarrp — Specifies that only (*, *, rp) entries be displayed.
type starg — Specifies that only (*,G) entries be displayed.
type sg — specifies that only (S,G) entries be displayed.
detail — Displays detailed group information.
family — Displays either IPv4 or IPv6 information.

Output **PIM Group Output** — The following table provides PIM Group field descriptions.

Label	Description
Group Address	Displays the IP multicast group address for which this entry contains information.

Label	Description (Continued)
Source Address	Displays the source address of the multicast sender. It will be 0 if the type is configured as starg. It will be the address of the Rendezvous Point (RP) if the type is configured as starRP.
RP Address	Displays the RP address.
Type	Specifies the type of entry, (*,*, rp)/(*,G) or (S,G).
Spt Bit	Specifies whether to forward on (*,*, rp)/(*,G) or on (S,G) state. It is updated when the (S,G) data comes on the RPF interface towards the source.
Incoming Intf	Displays the interface on which the traffic comes in. It can be the RPF interface to the RP (if starg) or the source (if sg).
Num Oifs	Displays the number of interfaces in the inherited outgoing interface list. An inherited list inherits the state from other types.
Flags	Displays the different lists that this interface belongs to.
Keepalive Timer Exp	The keepalive timer is applicable only for (S,G) entries. The (S,G) keepalive timer is updated by data being forwarded using this (S,G) Forwarding state. It is used to keep (S,G) state alive in the absence of explicit (S,G) joins.
MRIB Next Hop	Displays the next hop address towards the RP.
MRIB Src Flags	Displays the MRIB information about the source. If the entry is of type starg or starstarrp, it will contain information about the RP for the group.
Up Time	Displays the time since this source group entry was created.
Resolved By	Displays the route table used for RPF check.
Up JP State	Displays the upstream join prune state for this entry on the interface. PIM join prune messages are sent by the downstream routers towards the RPF neighbor.
Up JP Expiry	Displays the minimum amount of time remaining before this entry will be aged out.
Up JP Rpt	Displays the join prune Rpt state for this entry on the interface. PIM join/prune messages are sent by the downstream routers towards the RPF neighbor. (S,G, rpt) state is a result of receiving (S,G, rpt) JP message from the downstream router on the RP tree.

Label	Description (Continued)
Up JP Rpt Override	Displays the value used to delay triggered Join (S,G, rpt) messages to prevent implosions of triggered messages. If this has a non-zero value, it means that the router was in 'notPruned' state and it saw a prune (S,G, rpt) message being sent to RPF (S,G, rpt). If the router sees a join (S,G, rpt) override message being sent by some other router on the LAN while the timer is still non-zero, it simply cancels the override timer. If it does not see a join (S,G, rpt) message, then on expiry of the override timer, it sends it's own join (S,G, rpt) message to RPF (S,G, rpt). A similar scenario exists when RPF (S,G, rpt) changes to become equal to RPF (*,G).
Register State	Specifies the register state. The register state is kept at the source DR. When the host starts sending multicast packets and if there are no entries programmed for that group, the source DR sends a register packet to the RP (g). Register state transition happen based on the register stop timer and the response received from the RP.
Register Stop Exp	Displays the time remaining before the register state might transition to a different state.
Register from Anycast RP	Displays if the register packet for that group has been received from one of the RP from the anycast-RP set.
RPF Neighbor	Displays the address of the RPF neighbor.
Outgoing Intf List	Displays a list of interfaces on which data is forwarded.
Curr Fwding Rate	Displays the current forwarding rate of the multicast data for this group and source. This forwarding rate is calculated before ingress QoS policing or shaping is applied.
Forwarded Packets	Displays the number of multicast packets that were forwarded to the interfaces in the outgoing interface list. This packet count is before ingress QoS policing or shaping is applied.
Discarded Packets	Displays the number of multicast packets that matched this source group entry but were discarded. For (S,G) entries, if the traffic is getting forwarded on the SPT, the packets arriving from the RPT will be discarded.
Forwarded Octets	Displays the number of octets forwarded.
RPF Mismatches	Displays the number of multicast packets that matched this source group entry but they did not arrive on the interface.
Spt threshold	Displays the value of the SPT threshold configured for that group. 0 Kbps means that the switch to the SP tree will happen immediately.

Sample Output

```

A:NYC>show>router>pim# group
=====
PIM Groups
=====
Group Address      Source Address  RP Address      Type           Spt Incoming    Num
                  Bit           Intf            Oifs
-----
224.24.24.24      *              2.22.187.240   <*,G>          nyc-sjc        1
239.255.255.250  *              2.22.187.240   <*,G>          nyc-sjc        1
-----
Groups : 2
=====
A:NYC>show>router>pim#

A:NYC>show>router>pim# group 239.255.255.250
=====
PIM Groups
=====
Group Address      Source Address  RP Address      Type           Spt Incoming    Num
                  Bit           Intf            Oifs
-----
239.255.255.250  *              2.22.187.240   <*,G>          nyc-sjc        1
-----
Groups : 1
=====
A:NYC>show>router>pim#

A:NYC>show>router>pim# group 239.255.255.250 detail
=====
PIM Source Group
=====
Group Address      : 239.255.255.250 Source Address      : 16.1.1.2
RP Address         : 100.100.100.1 Type                : (S,G)
Flags              : spt, rpt-prn-des Keepalive Timer Exp: 0d 00:03:07
MRIB Next Hop      : 16.1.1.2 MRIB Src Flags      : direct
Up Time            : 0d 00:00:50 Resolved By         : rtable-u

Up JP State        : Joined Up JP Expiry         : 0d 00:00:00
Up JP Rpt          : Pruned Up JP Rpt Override  : 0d 00:00:00

Register State     : Pruned Register Stop Exp   : 0d 00:00:47
Reg From Anycast RP: No

RPF Neighbor       : 16.1.1.2
Incoming Intf      : SOURCE-3
Outgoing Intf List : To-Dut-A

Curr Fwding Rate   : 482.9 kbps
Forwarded Packets  : 1262 Discarded Packets   : 0
Forwarded Octets   : 1269572 RPF Mismatches      : 0
Spt threshold      : 0 kbps
=====
A:NYC>show>router>pim#

```

Show Router PIM Commands

```
B:Dut-C# show router pim group 225.0.0.1 type sg detail
=====
PIM Source Group ipv4
=====
Group Address       : 225.0.0.1
Source Address      : 11.11.0.1
RP Address          : 10.20.1.3
Flags               : rpt-prn-des          Type               : (S,G)
MRIB Next Hop       : 11.11.0.1
MRIB Src Flags      : direct               Keepalive Timer      : Not Running
Up Time             : 0d 00:04:17          Resolved By          : rtable-u

Up JP State         : Joined               Up JP Expiry         : 0d 00:00:00
Up JP Rpt           : Pruned              Up JP Rpt Override   : 0d 00:00:00

Register State      : No Info
Reg From Anycast RP: No

Rpf Neighbor        : 11.11.0.1
Incoming Intf       : svc_itf
Outgoing Host List  : 112.112.1.1

Curr Fwding Rate    : 0.0 kbps
Forwarded Packets   : 0                   Discarded Packets    : 0
Forwarded Octets    : 0                   RPF Mismatches       : 0
Spt threshold       : 0 kbps              ECMP opt threshold   : 7
Admin bandwidth     : 1 kbps              Preference           : 0

=====
PIM Source Group ipv4
=====
Group Address       : 225.0.0.1
Source Address      : 11.11.0.2
RP Address          : 10.20.1.3
Flags               :                     Type               : (S,G)
MRIB Next Hop       : 11.11.0.2
MRIB Src Flags      : direct               Keepalive Timer      : Not Running
Up Time             : 0d 00:04:18          Resolved By          : rtable-u

Up JP State         : Joined               Up JP Expiry         : 0d 00:00:00
Up JP Rpt           : Not Pruned          Up JP Rpt Override   : 0d 00:00:00

Register State      : No Info
Reg From Anycast RP: No

Rpf Neighbor        : 11.11.0.2
Incoming Intf       : svc_itf
Outgoing Host List  : 112.112.1.1, 112.112.1.2

Curr Fwding Rate    : 0.0 kbps
Forwarded Packets   : 0                   Discarded Packets    : 0
Forwarded Octets    : 0                   RPF Mismatches       : 0
Spt threshold       : 0 kbps              ECMP opt threshold   : 7
Admin bandwidth     : 1 kbps              Preference           : 0
-----
Groups : 2
=====
*B:Dut-C#
```



```
A:Dut-A# show router pim group detail
=====
PIM Source Group ipv4
=====
Group Address      : 224.1.1.1
Source Address     : 3.1.1.21
RP Address         : 10.20.1.4
Advt Ruoter       : 10.20.1.3
Flags              :
Type               : (S,G)
MRIB Next Hop     : 10.10.2.3
MRIB Src Flags    : remote
Standby Src Flags : remote
keepalive Timer   : Not Running
Up Time           : 0d 00:01:22
Resolved By       : rtable-u

Up JP State       : Joined
Up JP Rpt         : Pruned
Up JP Expiry      : 0d 00:00:00
Up JP Rpt Override : 0d 00:00:00
Up Stdbby JP State : Joined
Up Stdbby JP Expiry : 0d 00:00:12

Register State    : No Info
Reg From Anycast RP: No

Rpf Neighbor      : 10.10.2.3
Incoming Intf     : ip-10.10.2.1
Outgoing Host List : ix
Stdbby Rpf Neighbor : 10.10.1.2
Stdbby Intf       : ip-10.10.1.1

Curr Fwding Rate  : 0.0 kbps
Forwarded Packets : 0
Discarded Packets : 0
Forwarded Octets  : 0
RPF Mismatches    : 0
Spt threshold     : 0 kbps
ECMP opt threshold : 7
Admin bandwidth   : 1 kbps

=====
PIM Source Group ipv4

A:Dut-A# show router pim group
=====
PIM Group ipv4
=====
Group Address      Type      Spt Bit Inc Intf      no.Oifs
Source Address     RP
-----
224.1.1.1          (S,G)          ip-10.10.2.1      1
3.1.1.2            10.20.1.4      ip-10.10.1*
```

interface

- Syntax** `interface [ip-int-name | mt-int-name |ip-address] [group grp-ip-address | source ip-address [type {starstarrp | starg | sg}]] [detail] [family]`
- Context** `show>router>pim`
- Description** This command displays PIM interface information and the (S,G)/(*,G)/(*, *, rp) state of the interface.
- Parameters** *ip-int-name* — Only displays the interface information associated with the specified IP interface name.

Show Router PIM Commands

ip-address — Only displays the interface information associated with the specified IP address.

group *grp-ip-address* — Specifies the IP multicast group address for which this entry contains information.

source *ip-address* — Specifies the source address for which this entry contains information.

If the type is *starg*, the value of this object will be zero.

If the type is *starstarrp*, the value of this object will be address of the RP.

type — Specifies the type of this entry.

Values *starstarrp*, *starg*, *sg*

detail — Displays detailed interface information.

family — Displays IPv4 or IPv6 information for the interface.

Output **PIM Interface Output** — The following table provides PIM interface field descriptions.

Label	Description
Admin State	Displays the administrative state for PIM protocol on this interface.
Oper State	Displays the current operational state of PIM protocol on this interface.
DR	Displays the designated router on this PIM interface.
DR Priority	Displays the priority value sent in PIM Hello messages and that is used by routers to elect the designated router (DR).
Hello Intvl	Indicates the frequency at which PIM Hello messages are transmitted on this interface.

Sample Output

```
ALA-1# show router pim interface
=====
PIM Interfaces
=====
Interface                Admin  Oper   DR           DR        Hello
                        State  State                Priority  Intvl
-----
system                   Up     Up     N/A          1         30
ip-10.1.7.1              Up     Up     10.1.7.7     5         30
ip-10.1.2.1              Up     Up     10.1.2.2     5         30
ip-100.111.1.1           Up     Up     100.111.1.1  5         30
-----
Interfaces : 4
=====
ALA-1#
```

```
ALA-1# show router pim interface ip-10.1.2.1 detail
```

```

=====
PIM Interface ip-10.1.2.1
=====
Interface                Admin Oper  DR          DR          Hello
                        State State                    Priority Intvl
-----
ip-10.1.2.1              Up    Up    10.1.2.2    5           30
-----
PIM Group Source
-----
Group Address      : 228.101.0.5      Src Address      : 100.111.1.2
Interface          : ip-10.1.2.1      Type             : <S,G>
RP Address         : 200.200.200.4

Join Prune State   : Join                Expires          : 0d 00:03:00
Prune Pend Expires : N/A

Assert State       : No Info
-----
Interfaces : 1
=====
ALA-1#

ALA-1# show router pim interface group
=====
PIM Interface ip-10.1.7.1
=====
Interface                Admin Oper  DR          DR          Hello
                        State State                    Priority Intvl
-----
ip-10.1.7.1              Up    Up    10.1.7.7    5           30
-----
Group Address      Source Address  RP Address      Type      JP      Assert
-----
228.101.0.0        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.1        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.2        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.3        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.4        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.6        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.7        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.8        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
228.101.0.9        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
-----
PIM Interface ip-10.1.2.1
=====
Interface                Admin Oper  DR          DR          Hello
                        State State                    Priority Intvl
-----
ip-10.1.2.1              Up    Up    10.1.2.2    5           30
-----
Group Address      Source Address  RP Address      Type      JP      Assert
-----
228.101.0.5        100.111.1.2    200.200.200.4  <S,G>     Join    No Info
=====
PIM Interface ip-100.111.1.1
=====
Interface                Admin Oper  DR          DR          Hello

```

Show Router PIM Commands

```

                State  State                Priority  Intvl
-----
ip-100.111.1.1      Up    Up      100.111.1.1    5        30
-----
Group Address      Source Address  RP Address      Type    JP      Assert
-----
228.102.0.0        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.1        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.2        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.3        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.4        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.5        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.6        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.7        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.8        *              200.200.200.4  <*,G>   Join    No Info
228.102.0.9        *              200.200.200.4  <*,G>   Join    No Info
-----

```

Interfaces : 3

ALA-1#

ALA-1# show router pim interface group 228.102.0.0 detail

PIM Interface ip-100.111.1.1

```

Interface                Admin  Oper  DR                DR      Hello
                        State  State                Priority Intvl
-----
ip-100.111.1.1          Up    Up    100.111.1.1      5        30
-----

```

PIM Group Source

```

Group Address      : 228.102.0.0      Src Address      : *
Interface          : ip-100.111.1.1  Type             : <*,G>
RP Address         : 200.200.200.4

Join Prune State   : Join              Expires          : 0d 00:02:05
Prune Pend Expires : N/A

```

Assert State : No Info

Interfaces : 1

ALA-1#

ALA-1# show router pim interface type starg

PIM Interface ip-100.111.1.1

```

Interface                Admin  Oper  DR                DR      Hello
                        State  State                Priority Intvl
-----
ip-100.111.1.1          Up    Up    100.111.1.1      5        30
-----

```

```

Group Address      Source Address  RP Address      Type    JP      Assert
-----
228.102.0.0        *              200.200.200.4  <*,G>   Join    No Info

```

```

228.102.0.1      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.2      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.3      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.4      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.5      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.6      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.7      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.8      *           200.200.200.4    <*,G>    Join      No Info
228.102.0.9      *           200.200.200.4    <*,G>    Join      No Info

```

```
-----
Interfaces : 1
=====
```

```
ALA-1#
```

```
A:SetupCLI# show router pim interface detail
```

```
=====
PIM Interface int1
=====
```

```

Interface           : int1
Admin Status         : Up
DR                   : 10.1.1.1
BSM RA Check         : Disabled
Hello Interval       : 30
Multicast Senders    : auto
J/P Tracking Admin   : Disabled
Auto-created         : No
Sticky-DR            : Disabled
Max Groups Allowed   : 0
Num Groups           : 0
Oper Status          : Up
Oper DR Priority      : 1
Cfg DR Priority       : 1
Time for next hello  : 0d 00:00:23
Hello Multiplier     : 35
J/P Tracking Oper    : Disabled
Improved Assert      : Enabled
Sticky-DR Priority    : N/A
Max Groups Till Now  : 0
Bfd Enabled          : No

```

```
=====
PIM Interface sender
=====
```

```

Interface           : sender
Admin Status         : Up
DR                   : 11.1.1.1
Oper Status          : Up
Oper DR Priority      : 1

```

```
=====
A:SetupCLI#
```

neighbor

Syntax **neighbor** [*ip-address* | *ip-int-name* [**address** *ip-address*]] [**detail**] [*family*]

Context show>router>pim

Description This command displays PIM neighbor information.

This can be important if an interface has more than one adjacency. For example, a LAN-interface configuration with three routers connected and all are running PIM on their LAN interfaces. These routers then have two adjacencies on their LAN interface, each with different neighbors. If the **address address** parameter is not defined in this example, then the **show** command output would display two adjacencies.

Parameters **neighbor** *ip-int-name* — Only displays the interface information associated with the specified IP interface name.

neighbor *ip-address* — Only displays the interface information associated with the specified IP address.

address *ip-address* — The ip-address of the neighbor, on the other side of the interface.

detail — Displays detailed neighbor information.

family — Displays either IPv4 or IPv6 information for the specified neighbor.

Output **PIM Neighbor Output** — The following table provides PIM neighbor field descriptions.

Label	Description
Interface	Displays the neighbor's interface name.
Nbr DR Priority	Displays the value of the neighbor's DR priority which is received in the hello message.
Nbr Address	Displays the neighbor's address.
Up Time	Displays the time since this PIM neighbor (last) became a neighbor of the local router.
Expiry Time	Displays the minimum time remaining before this PIM neighbor will be aged out. 0 — Means that this neighbor will never be aged out. This happens when the PIM neighbor sends a Hello message with holdtime set to '0xffff'.
Hold Time	Displays the value of the hold time present in the hello message.
DR Priority	Displays the value of the neighbor's DR priority which is received in the hello message.
Tracking Support	Displays whether the T bit in the LAN prune delay option was present in the hello message. This indicates the neighbor's capability to disable join message suppression.
LAN Delay	Displays the value of the LAN delay field present in the hello message received from the neighbor.
Gen Id	Displays a randomly generated 32-bit value that is regenerated each time PIM forwarding is started or restarted on the interface, including when the router itself restarts. When a hello message with a new GenID is received from a neighbor, any old hello information about that neighbor is discarded and superseded by the information from the new hello message.
Override Intvl (ms)	Displays the value of the override interval present in the Hello message.

Sample Output

```
ALA-1# show router pim neighbor
```

```
=====
```

```

PIM Neighbors
=====
Interface           Nbr DR      Nbr Address  Up Time      Expiry Time  Hold
                    Priority
-----
ip-10.1.7.1         5           10.1.7.7     0d 00:10:39  0d 00:01:36  105
ip-10.1.2.1         5           10.1.2.2     0d 00:10:39  0d 00:01:35  105
ip-100.111.1.1      3           100.111.1.2  0d 00:09:31  0d 00:01:15  105
-----
Neighbors : 3
=====
ALA-1#

ALA-1# show router pim neighbor detail
=====
PIM Neighbor
=====
Interface           : ip-10.1.7.1
Neighbor Addr       : 10.1.7.7          DR Priority        : 5
Tracking Support    : No              LAN Delay(ms)     : 500
Gen Id              : 26470             Override Intvl(ms) : 2500
Up Time             : 0d 00:10:41    Expiry Time       : 0d 00:01:34
Hold Time(sec)      : 105

=====
PIM Neighbor
=====
Interface           : ip-10.1.2.1
Neighbor Addr       : 10.1.2.2          DR Priority        : 5
Tracking Support    : No              LAN Delay(ms)     : 500
Gen Id              : 37928             Override Intvl(ms) : 2500
Up Time             : 0d 00:10:42    Expiry Time       : 0d 00:01:33
Hold Time(sec)      : 105

=====
PIM Neighbor
=====
Interface           : ip-100.111.1.1
Neighbor Addr       : 100.111.1.2      DR Priority        : 3
Tracking Support    : No              LAN Delay(ms)     : 500
Gen Id              : 742098371        Override Intvl(ms) : 2500
Up Time             : 0d 00:09:33    Expiry Time       : 0d 00:01:43
Hold Time(sec)      : 105
-----
Neighbors : 3
=====
ALA-1#

```

rp

Syntax **rp** *ip-address***Context** show>router>pim**Description** This command displays the rendezvous point (RP) set information built by the router.

Show Router PIM Commands

Parameters *ip-address* — Specifies the IP address of the RP.

Output **PIM Neighbor Output** — The following table provides PIM neighbor field descriptions.

Label	Description
Group Address	Displays the multicast group address of the entry.
RP Address	Displays the address of the Rendezvous Point (RP).
Type	Specifies whether the entry was learned through the Bootstrap mechanism or if it was statically configured.
Priority	Displays the priority for the specified group address. The higher the value, the higher the priority.
Holdtime	Displays the value of the hold time present in the BSM message.

Sample Output

```
A:ALA-1# show router pim rp
=====
PIM RP Set
=====
Group Address      RP Address      Type      Priority  Holdtime
-----
224.0.0.0/4        200.200.200.4   Dynamic    192       150
                   10.1.7.1        Static      1         N/A
-----
Group Prefixes : 1
=====
A:ALA-1#

A:ALA-1# show router pim rp 10.1.7.1
=====
PIM RP Set
=====
Group Address      RP Address      Type      Priority  Holdtime
-----
224.0.0.0/4        10.1.7.1        Static      1         N/A
-----
Group Prefixes : 1
=====
A:ALA-1#
```

rp-hash

Syntax **rp-hash** *grp-ip-address*

Context show>router>pim

Description This command hashes the RP for the specified group from the RP set.

Parameters *grp-ip-address* — Displays specific multicast group addresses.

Output **PIM RP-Hash Output** — The following table provides RP-Hash output field descriptions.

Label	Description
Group Address	Displays the multicast group address of the entry.
RP Address	Displays the address of the Rendezvous Point (RP).
Type	Specifies whether the entry was learned through the Bootstrap mechanism or if it was statically configured.

Sample Output

```
A:ALA-1# show router pim rp-hash 228.101.0.0
=====
PIM Group-To-RP mapping
=====
Group Address      RP Address      Type
-----
228.101.0.0        200.200.200.4   Bootstrap
=====
A:ALA-1#
```

```
A:ALA-1# show router pim rp-hash 228.101.0.6
=====
PIM Group-To-RP mapping
=====
Group Address      RP Address      Type
-----
228.101.0.6        200.200.200.4   Bootstrap
=====
A:ALA-1#
```

statistics

Syntax **statistics** [*ip-int-name* | *mt-int-name* | *ip-address*] [*family*]

Context show>router>pim

Description This command displays statistics for a particular PIM instance.

Parameters *ip-int-name* — Only displays the interface information associated with the specified IP interface name.
ip-address — Only displays the interface information associated with the specified IP address.
family — Displays either IPv4 or IPv6 information.

Output **PIM Statistics Output** — The following table provides PIM statistics output field descriptions.

Label	Description
PIM Statistics	The section listing the PIM statistics for a particular interface.
Message Type	Displays the type of message. Hello — Displays the number of PIM hello messages received or transmitted on this interface. Join Prune — Displays the number of PIM join prune messages received or transmitted on this interface. Asserts — Displays the number of PIM assert messages received or transmitted on this interface. Register — Displays the number of register messages received or transmitted on this interface. Null Register — Displays the number of PIM null register messages received or transmitted on this interface. Register Stop — Displays the number of PIM register stop messages received or transmitted on this interface. BSM — Displays the number of PIM Bootstrap messages (BSM) received or transmitted on this interface. Candidate RP Adv — Displays the number of candidate RP advertisements. Total Packets — Displays the total number of packets transmitted and received on this interface.
Received	Displays the number of messages received on this interface.
Transmitted	Displays the number of multicast data packets transmitted on this interface.
Rx Errors	Displays the total number of receive errors.
General Interface Statistics	The section listing the general PIM interface statistics.
Register TTL Drop	Displays the number of multicast data packets which could not be encapsulated in Register messages because the time to live (TTL) was zero.
Tx Register MTU Drop	Displays the number of Bootstrap messages received on this interface but were dropped.
Rx Invalid Register	Displays the number of invalid PIM register messages received on this interface.

Label	Description (Continued)
Rx Neighbor Unknown	Displays the number of PIM messages (other than hello messages) which were received on this interface and were rejected because the adjacency with the neighbor router was not already established.
Rx Bad Checksum Discard	Displays the number of PIM messages received on this interface which were discarded because of bad checksum.
Rx Bad Encoding	Displays the number of PIM messages with bad encodings received on this interface.
Rx Bad Version Discard	Displays the number of PIM messages with bad versions received on this interface.
Rx CRP No Router Alert	Displays the number of candidate-rp advertisements (C-RP-Adv) received on this interface which had no router alert option set.
Rx Invalid Join Prune	Displays the number of invalid PIM join prune messages received on this interface.
Rx Unknown PDU Type	Displays the number of packets received with an unsupported PIM type.
Join Policy Drops	Displays the number of times the join policy match resulted in dropping PIM join-prune message or one of the source group contained in the message.
Register Policy Drops	Displays the number of times the register policy match resulted in dropping PIM register message.
Bootstrap Import Policy Drops	Displays the number of Bootstrap messages received on this interface but were dropped because of Bootstrap import policy.
Bootstrap Export Policy Drops	Displays the number of Bootstrap messages that were not transmitted on this interface because of Bootstrap export policy.
Source Group Statistics	The section listing the source group statistics.
(S,G)	Displays the number of entries in which the type is (S,G).
(*,G)	Displays the number of entries in which the type is (*,G).
(*,*,RP)	Displays the number of entries in which the type is (*,*,rp).

Sample output

```

A:ALA-1# show router pim statistics
=====
PIM Statistics
=====
Message Type      Received      Transmitted    Rx Errors
-----

```

Show Router PIM Commands

Hello	198	200	0
Join Prune	96	75	0
Asserts	0	0	0
Register	0	30	0
Null Register	0	160	0
Register Stop	180	0	0
BSM	34	76	0
Candidate RP Adv	0	0	0
Total Packets	546	541	

----- General Interface Statistics

Register TTL Drop	: 0
Tx Register MTU Drop	: 0
Rx Invalid Register	: 0
Rx Neighbor Unknown	: 0
Rx Bad Checksum Discard	: 0
Rx Bad Encoding	: 0
Rx Bad Version Discard	: 0
Rx CRP No Router Alert	: 0
Rx Invalid Join Prune	: 120
Rx Unknown PDU Type	: 0
Join Policy Drops	: 0
Register Policy Drops	: 0
Bootstrap Import Policy Drops	: 0
Bootstrap Export Policy Drops	: 0

----- Source Group Statistics

(S,G)	: 10
(*,G)	: 10
(*,*,RP)	: 0

=====

A:ALA-1#

A:ALA-1# show router pim statistics 10.1.7.1

=====

PIM Interface 10.1.7.1 Statistics

Message Type	Received	Transmitted	Rx Errors
-----	-----	-----	-----
Hello	62	66	0
Join Prune	36	21	0
Asserts	0	0	0
Register	0	0	0
Null Register	0	0	0
Register Stop	0	0	0
BSM	33	3	0
Total Packets	134	90	

----- General Interface Statistics

Register TTL Drop	: 0
Tx Register MTU Drop	: 0
Rx Invalid Register	: 0
Rx Neighbor Unknown	: 0
Rx Bad Checksum Discard	: 0
Rx Bad Encoding	: 0
Rx Bad Version Discard	: 0

```

Rx CRP No Router Alert          : 0
Rx Invalid Join Prune           : 0
Rx Unknown PDU Type             : 0
Join Policy Drops               : 0
Register Policy Drops           : 0
Bootstrap Import Policy Drops   : 0
Bootstrap Export Policy Drops   : 0
-----
Interface Source Group Statistics
-----
(S,G)                          : 9
(*,G)                          : 0
(*,*,RP)                       : 0
=====
A:ALA-1#

A:ALA-1# show router pim statistics ip-10.1.7.1
=====
PIM Interface ip-10.1.7.1 Statistics
=====
Message Type      Received      Transmitted      Rx Errors
-----
Hello             63             67              0
Join Prune        36             21              0
Asserts           0              0              0
Register          0              0              0
Null Register     0              0              0
Register Stop     0              0              0
BSM               33             3              0
Total Packets     135            91
-----
General Interface Statistics
-----
Register TTL Drop          : 0
Tx Register MTU Drop       : 0
Rx Invalid Register        : 0
Rx Neighbor Unknown        : 0
Rx Bad Checksum Discard    : 0
Rx Bad Encoding            : 0
Rx Bad Version Discard     : 0
Rx CRP No Router Alert     : 0
Rx Invalid Join Prune      : 0
Rx Unknown PDU Type        : 0
Join Policy Drops          : 0
Register Policy Drops      : 0
Bootstrap Import Policy Drops : 0
Bootstrap Export Policy Drops : 0
-----
Interface Source Group Statistics
-----
(S,G)                          : 9
(*,G)                          : 0
(*,*,RP)                       : 0
=====
A:ALA-1#

```

status

Syntax **status** [**detail**] [*family*]

Context show>router>pim

Description This command displays PIM status. The Oper Status reflects the combined operational status of IPv4/IPv6 PIM protocol status. If both are down, then Oper Status will be reflected as down. If IPv4 or IPv6 reflects up, the Oper Status will reflect up.

If PIM is not enabled, the following message appears:

```
A:NYC# show router pim status
MINOR: CLI PIM is not configured.
A:NYC#
```

Parameters **detail** — Displays detailed status information.

family — Displays either IPv4 or IPv6 information.

Output **PIM Status Output** — The following table provides PIM status output field descriptions.

Label	Description
Admin State	Displays the administrative status of PIM.
Oper State	Displays the current operating state of this PIM protocol instance.
BSR State	Displays the state of the router with respect to the Bootstrap mechanism.
Address	Displays the address of the elected Bootstrap router.
Expiry Time	Displays the time remaining before the router sends the next Bootstrap message.
Priority	Displays the priority of the elected Bootstrap router. The higher the value, the higher the priority.
Hash Mask Length	Displays the hash mask length of the Bootstrap router.
Up Time	Displays the time since the current E-BSR became the Bootstrap router.
RPF Intf towards	Displays the RPF interface towards the elected BSR. The value is zero if there is no elected BSR in the network.
Address	Displays the address of the candidate BSR router.
Expiry Time	Displays the time remaining before the router sends the next Bootstrap message.
Priority	Displays the priority of the Bootstrap router. The higher the value, the higher the priority.
Hash Mask Length	Displays the hash mask length of the candidate Bootstrap router.

Label	Description (Continued)
Up Time	Displays the time since becoming the Bootstrap router.
Admin State	Displays the administrative status of CRP.
Oper State	Displays the current operating state of the C-RP mechanism.
Address	Displays the local RP address.
Priority	Displays the CRP's priority for becoming a rendezvous point (RP). A 0 value is the highest priority.
Holdtime	Displays the hold time of the candidate RP. It is used by the Bootstrap router to timeout the RP entries if it does not listen to another CRP advertisement within the holdtime period.
Policy	Displays the PIM policies for a particular PIM instance.
Default Group	Displays the default core group address.
RPF Table	Displays the route table used for RPF check.
MC-ECMP-Hashing	Displays if hash-based multicast balancing of traffic over ECMP links is enabled or disabled.

Sample Output

```
A:dut-d# show router pim status
=====
PIM Status
=====
Admin State           : Up
Oper State            : Up

BSR State              : Accept Any

Elected BSR
  Address              : None
  Expiry Time          : N/A
  Priority              : N/A
  Hash Mask Length     : N/A
  Up Time              : N/A
  RPF Intf towards E-BSR : N/A

Candidate BSR
  Admin State          : Down
  Oper State           : Down
  Address              : None
  Priority              : 0
  Hash Mask Length     : 30

Candidate RP
  Admin State          : Down
  Oper State           : Down
  Address              : None
```

Show Router PIM Commands

```

Priority                : 192
Holdtime                : 150

MC-ECMP-Hashing        : Enabled

Policy                 : None

Default Group           : 239.1.1.1

RPF Table               : rtable-m
=====
A:dut-d#
```

mld

- Syntax** mld
- Context** show>router
- Description** This command displays MLD related information.

group

- Syntax** group [grp-ipv6-address]
- Context** show>router>mld
- Description** This command displays MLD group information.
- Parameters** grp-ipv6-address — Specifies the IPv6 group 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..FFFF]H
 d: [0..255]D

```

Output
*A:SR7# show router mld group
=====
MLD Groups
=====
No Matching Entries
=====
*A:SR7#

*A:SR7# show router mld interface
=====
MLD Interfaces
=====
Interface      Adm  Oper  Cfg/Opr      Num      Policy
Querier                               Groups
-----
Host4_Srcel_IPv6  Up   Up    2/2           0         none
```



```

FE80::216:4DFF:FED4:4D5B
Host1      Up    Up    2/2      0      none
FE80::216:4DFF:FED4:4D5B
Host2      Up    Up    2/2      0      none
FE80::216:4DFF:FE51:3728
Host3_vlan1 Up    Up    2/2      0      none
FE80::216:4DFF:FE51:3729
Host3_vlan2 Up    Up    2/2      0      none
FE80::216:4DFF:FE51:3729
Host3_vlan3 Up    Up    2/2      0      none
FE80::216:4DFF:FE51:3729
Host3_vlan4 Up    Up    2/2      0      none
FE80::216:4DFF:FE51:3729
Host3_vlan5 Up    Up    2/2      0      none
*A:SR7# show router mld ssm-translate
=====
MLD SSM Tranlate Entries
=====
No Matching Entries
=====
*A:SR7#

*A:SR7# show router mld group
=====
MLD Groups
=====
(3FFE:100::2:100,FF05::1:1)
  Up Time : 0d 00:00:31
  Fwd List : Host1

(3FFE:100::2:100,FF05::1:2)
  Up Time : 0d 00:00:31
  Fwd List : Host1

(3FFE:100::2:100,FF05::1:3)
  Up Time : 0d 00:00:31
  Fwd List : Host1

(3FFE:100::2:100,FF05::1:4)
  Up Time : 0d 00:00:31
  Fwd List : Host1

(3FFE:100::2:100,FF05::1:5)
=====
*A:SR7#

*A:SR7# show router mld group ff05::1:1
=====
MLD Groups
=====
(3FFE:100::2:100,FF05::1:1)
  Up Time : 0d 00:00:40
  Fwd List : Host1
-----
(*,G)/(S,G) Entries : 1
=====
*A:SR7#

```

Show Router PIM Commands

```
*A:SR7# show router mld group ff05::1
=====
MLD Groups
=====
No Matching Entries
=====
```

interface

- Syntax** **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-ipv6-address*] [**detail**]
- Context** show>router>mld
- Description** This command displays MLD interface information.
- Parameters**

ip-int-name|*ip-address* — Specifies the IP interface name or interface address.

group *grp-ipv6-address* — Specifies the IPv6 group address.

Values

ipv6-address	x::x::x::x::x (eight 16-bit pieces)
	x::x::x::x::d.d.d.d
	x: [0..FFFF]H
	d: [0..255]D

detail — Displays detailed information.

Output

```
*A:SR7# show router mld interface Host1 detail
=====
MLD Interface Host1
=====
Interface           : Host1
Admin Status        : Up
Oper Status         : Up
Querier             : FE80::216:4DFF:FED4:4D5B
Querier Up Time     : 0d 00:02:18
Querier Expiry Time : N/A
Time for next query : 0d 00:15:25
Admin/Oper version  : 2/2
Num Groups          : 6000
Policy              : none
Max Groups Allowed  : No Limit
Max Groups Till Now : 6000
Query Interval      : 0
Query Resp Interval : 0
Last List Qry Interval : 0
-----
MLD Group
-----
Group Address : FF05::1:1
Last Reporter : FE80::1
Interface     : Host1
Expires      : N/A
Up Time      : 0d 00:00:10
Mode         : include
Vl Host Timer : Not running
Type         : dynamic
Compat Mode  : MLD Version 2
-----
Source
Expires      Type      Fwd/Blk
-----
3FFE:100::2:100
0d 00:34:07   dynamic Fwd
```

MLD Group

```

Group Address : FF05::1:2
Last Reporter : FE80::1
Interface      : Host1                Expires      : N/A
Up Time       : 0d 00:00:11           Mode         : include
V1 Host Timer : Not running           Type         : dynamic
Compat Mode   : MLD Version 2

```

Source

Expires	Type	Fwd/Blk
---------	------	---------

```

3FFE:100::2:100
    0d 00:34:07    dynamic Fwd

```

MLD Group

```

Group Address : FF05::1:3
Last Reporter : FE80::1
Interface      : Host1                Expires      : N/A
Up Time       : 0d 00:00:11           Mode         : include
V1 Host Timer : Not running           Type         : dynamic
Compat Mode   : MLD Version 2

```

Source

Expires	Type	Fwd/Blk
---------	------	---------

```

3FFE:100::2:100
    0d 00:34:07    dynamic Fwd

```

MLD Group

```

Group Address : FF05::1:4
Last Reporter : FE80::1
Interface      : Host1                Expires      : N/A
Up Time       : 0d 00:00:12           Mode         : include
V1 Host Timer : Not running           Type         : dynamic
Compat Mode   : MLD Version 2

```

Source

Expires	Type	Fwd/Blk
---------	------	---------

```

3FFE:100::2:100
    0d 00:34:06    dynamic Fwd

```

MLD Group

```

Group Address : FF05::1:5
Last Reporter : FE80::1
Interface      : Host1                Expires      : N/A
Up Time       : 0d 00:00:12           Mode         : include
V1 Host Timer : Not running           Type         : dynamic
Compat Mode   : MLD Version 2

```

Source

Expires	Type	Fwd/Blk
---------	------	---------

```

3FFE:100::2:100

```

0d 00:34:06 dynamic Fwd

ssm-translate

Syntax	ssm-translate
Context	show>router>mld
Description	This command displays the MLD SSM translate configuration.

static

Syntax	static [<i>ip-int-name</i> <i>ip-address</i>]
Context	show>router>mld
Description	This command displays MLD static group/source configuration.
Parameters	<i>ip-int-name</i> <i>ip-address</i> — iSpecifies the IP interface name or IP address.

Output

```
*A:SR7# show router mld static
=====
MLD Static Group Source
=====
Source                                     Group
  Interface
-----
No Matching Entries
=====
*A:SR7

*A:SR7# show router mld statistics
=====
MLD Interface Statistics
=====
Message Type      Received      Transmitted
-----
Queries           0             640
Report V1         0             0
Report V2        10             0
Dones             0             0
-----
General Interface Statistics
-----
Bad Length       : 0
Bad Checksum     : 0
Unknown Type     : 0
Bad Receive If   : 0
Rx Non Local     : 0
Rx Wrong Version : 0
Policy Drops     : 0
No Router Alert  : 0
```

```
Rx Bad Encodings   : 0
Rx Pkt Drops       : 0
Local Scope Pkts   : 10
Resvd Scope Pkts   : 0
-----
Source Group Statistics
-----
(S,G)               : 0
(*,G)               : 0
=====
*A:SR7#
```

statistics

- Syntax** **statistics** [*ip-int-name* | *ipv6-address*]
- Context** show>router>mld
- Description** This command displays MLD statistics.
ip-int-name|ipv6-address — iSpecifies the IP interface name or IPv6 address.

status

- Syntax** **status**
- Context** show>router>mld
- Description** This command displays the MLD status.

```
Output
*A:SR7# show router mld status
=====
MLD Status
=====
Admin State           : Up
Oper State            : Up
Query Interval        : 1024
Last Listener Query Interval : 1
Query Response Interval : 10
Robust Count          : 2
=====
*A:SR7#

*A:SR7# show router mld interface Host1
=====
MLD Interface Host1
=====
Interface      Adm  Oper  Cfg/Opr      Num      Policy
Querier        Version  Groups
-----
Host1          Up   Up    2/2          5082     none
FE80::216:4DFF:FED4:4D5B
=====
```

Show Router PIM Commands

```
Interfaces : 1
=====
*A:SR7#
```

group

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

Context show>router>msdp

Description This command displays information about MSDP groups.

Parameters *group-name* — Displays information about the specified group name. If no group-name is specified, information about all group names display.

detail — Displays detailed MSDP group information.

Output **MSDP Group Output** — The following table provides MSDP group field descriptions.

Label	Description
Group Name	Displays the MSDP group name.
Mode	Displays the groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers.
Act Srcs	Displays the configured maximum number of active source messages that will be accepted by MSDP.
Local Address	Displays the local end of a MSDP session.
Admin State	Displays the administrative state.
Receive Msg Rate	Displays rate that the messages are read from the TCP session.
Receive Msg Time	Displays the time of MSDP messages that are read from the TCP session within the configured number of seconds.
Receive Msg Thd	Displays the configured threshold number of MSDP messages can be processed before the MSDP message rate limiting function .
SA Limit	Displays the source-active limit.

Sample Output

```
*A:ALA-48>show>router>msdp# group
=====
MSDP Groups
=====
Group Name                Mode      Act Srcs  Local Address
-----
main                      Mesh-group  None  None
loop1                     Mesh-group  None  None
loop2                     Mesh-group  None  None
```

```

loop3                      Mesh-group None None
loop4                      Mesh-group None None
loop5                      Mesh-group None None
-----
Groups : 6
=====
*A:ALA-48>show>router>msdp#

*A:ALA-48>show>router>msdp# group test
=====
MSDP Groups
=====
Group Name                Mode        Act Srcs  Local Address
-----
test                      Mesh-group 50000    10.10.10.103
-----
Groups : 1
=====
*A:ALA-48>show>router>msdp#

*A:ALA-48>show>router>msdp# group test detail
=====
MSDP Groups
=====
Group Name                : test
-----
Local Address             : 10.10.10.103
Admin State               : Up
Receive Msg Rate          : None
Receive Msg Time         : None
Mode                     : Mesh-group
SA Limit                  : 50000
Export Policy             : None Specified / Inherited
Import Policy             : None Specified / Inherited
-----
Groups : 1
=====
*A:ALA-48>show>router>msdp#

```

peer

Syntax	peer [<i>ip-address</i>] [group <i>group-name</i>] [detail]
Context	show>router>msdp
Description	This command displays information about an MSDP peer.
Parameters	<p><i>ip-address</i> — Displays information about the specified IP address. If no IP address specified, information about all MSDP IP addresses display.</p> <p>group <i>group-name</i> — Displays information about the specified group name. If no <i>group-name</i> is specified, information about all MSDP peers display.</p> <p>detail — Displays detailed MSDP peer information.</p>

Output **MSDP Peer Output** — The following table provides MSDP field descriptions.

Label	Description
Peer	Displays the IP address of the peer.
Local Address	Displays the local IP address.
State	Displays the current state of the peer.
Last State Change	Displays the date and time of the peer's last state change.
SA Learn	The number of SAs learned through a peer.

Sample Output

```
A:ALA-48# show router msdp peer
=====
MSDP Peers
=====
Peer           Local Address   State           Last State Change   SA Learnt
-----
10.20.1.1      10.20.1.6       Established      08/30/2002 03:22:13 1008
-----
Peers : 1
=====
A:ALA-48#

A:ALA-48# show router msdp peer detail
=====
MSDP Peers
-----
Peer Address      : 10.20.1.1
-----
Group Name        : None
Local Address     : 10.20.1.6
Last State Change : 08/30/2002 03:22:13 Last Act Src Limit : N/A
Peer Admin State  : Up
Peer Connect Retry : 0
Peer Connect Retry : 0
SA accepted       : 1008
State timer expires: 18
Active Source Limit: None
Receive Msg Time  : 0
Auth Status       : Disabled
Export Policy     : None Specified / Inherited
Import Policy     : None Specified / Inherited
-----
Peers : 1
=====
A:ALA-48#
```


source

Syntax **source** [*ip-address/mask*] [**type** {**configured** | **dynamic** | **both**}] [**detail**]

Context show>router>msdp

Description This command displays the discovery method for this multicast source.

Parameters **configured** — Displays user-created sources.

dynamic — Displays dynamically created sources.

both — Displays both user-configured and dynamically created sources.

detail — Displays detailed MSDP source information.

Output **MSDP Source Output** — The following table provides MSDP source field descriptions.

Label	Description
Source	Displays the IP address of the peer.
Type	Displays the type of peer.
SA limit	Displays the local IP address.
State	Displays the current state of the peer.
Num excd	Indicates the number of times the global active source limit has been exceeded.
Last exceeded	Displays the date and time of the peer's last state change.

source-active

Syntax **source-active** [**group** *ip-address* | **local** | **originator** *ip-address* | **peer** *ip-address* | **source** *ip-address* | (**group** *ip-address* **source** *ip-address*)] [**detail**]

Context show>router>msdp

Description This command displays source active messages accepted by MSDP.

Parameters **group** *ip-address* — Displays information about the specified group IP address.

local — Displays information about local source-active messages.

originator *ip-address* — Displays information about the specified originator IP address.

peer *ip-address* — Displays information about the specified peer IP address.

source *ip-address* — Displays information about the specified source IP address.

group *ip-address* — Displays information about the specified group IP address.

detail Displays detailed MSDP source-active information.

Output **MSDP Source-Active Output** — The following table provides MSDP source-active field descriptions.

Show Router PIM Commands

Label	Description
Grp Address	Displays the IP address of the group.
Src Address	Displays the IP address of the source.
Origin RP	Displays the origination rendezvous point (RP) address.
Peer Address	Displays the address of the peer.
State Timer	The time-out value. If the value reaches zero, the SA is removed.

Sample Output

```
A:ALA-48# show router msdp source-active
=====
MSDP Source Active Info
=====
Grp Address      Src Address      Origin RP      Peer Address State Timer
-----
228.100.0.0      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.1      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.2      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.3      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.4      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.5      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.6      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.7      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.8      100.112.1.2      10.20.1.1      10.20.1.1 69
228.100.0.9      100.112.1.2      10.20.1.1      10.20.1.1 69
-----
MSDP Source Active : 10
=====
A:ALA-48#
```

```
A:ALA-48# show router msdp source-active detail
=====
MSDP Source Active
=====
Group Address      : 228.100.0.0      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
State Timer        : 64          Up Time            : 3d 01:44:25
Group Address      : 228.100.0.1      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
State Timer        : 64          Up Time            : 48d 18:22:29
Group Address      : 228.100.0.2      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
State Timer        : 64          Up Time            : 48d 18:22:29
Group Address      : 228.100.0.3      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
State Timer        : 64          Up Time            : 48d 18:22:29
Group Address      : 228.100.0.4      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
State Timer        : 64          Up Time            : 48d 18:22:29
Group Address      : 228.100.0.5      Source Address      : 100.112.1.2
Origin RP          : 10.20.1.1      Peer Address       : 10.20.1.1
```

```

State Timer      : 64                Up Time           : 48d 18:22:29
Group Address    : 228.100.0.6        Source Address    : 100.112.1.2
Origin RP        : 10.20.1.1          Peer Address      : 10.20.1.1
State Timer      : 64                Up Time           : 48d 18:22:29
Group Address    : 228.100.0.7        Source Address    : 100.112.1.2
Origin RP        : 10.20.1.1          Peer Address      : 10.20.1.1
State Timer      : 64                Up Time           : 48d 18:22:29
Group Address    : 228.100.0.8        Source Address    : 100.112.1.2
Origin RP        : 10.20.1.1          Peer Address      : 10.20.1.1
State Timer      : 64                Up Time           : 48d 18:22:29
Group Address    : 228.100.0.9        Source Address    : 100.112.1.2
Origin RP        : 10.20.1.1          Peer Address      : 10.20.1.1
State Timer      : 64                Up Time           : 48d 18:22:29

```

```
-----
MSDP Source Active : 10
=====
```

```
A:ALA-48#
```

source-active-rejected

Syntax **source-active-rejected** [**peer-group name**] [**group ip-address**] [**source ip-address**] [**originator ip-address**] [**peer ip-address**]

Context show>router>msdp

Description This command displays source active messages rejected by MSDP.

Parameters **group ip-address** — Displays information about the peer group name of the Source Active entry that is rejected.

local — Displays information about local source-active messages.

originator ip-address — Displays information about the specified originator IP address.

peer ip-address — Displays information about the peer from which this rejected source active entry was last received.

source ip-address — Displays information about the source address of the source active entry that is rejected.

group ip-address — Displays information about the specified group IP address.

MSDP Source-Active Output — The following table provides MSDP source-active field descriptions.

Label	Description
Grp Address	Displays the IP address of the group.
Src Address	Displays the IP address of the source.
Origin RP	Displays the origination rendezvous point (RP) address.
Peer Address	Displays the address of the peer.

Label	Description (Continued)
Reject Reason	Displays the reason why this source active entry is rejected.

Sample Output

```
*A:ALA-48# show router msdp source-active-rejected
=====
MSDP Source Active Rejected Info
=====
Grp Address      Src Address      Origin RP        Peer Address      Reject Reason
-----
228.100.0.1      110.0.0.1        10.20.0.1        20.0.0.1          Import Policy
228.100.0.2      110.0.0.2        10.20.0.2        20.0.0.2          Export Policy
228.100.0.3      110.0.0.3        10.20.0.3        20.0.0.3          RPF Failure
228.100.0.4      110.0.0.4        10.20.0.4        20.0.0.4          Limit Exceeded
228.100.0.5      110.0.0.5        10.20.0.5        20.0.0.5          Limit Exceeded
228.100.0.6      110.0.0.6        10.20.0.6        20.0.0.6          Limit Exceeded
228.100.0.7      110.0.0.7        10.20.0.7        20.0.0.7          Limit Exceeded
-----
SA Rejected Entries : 7
=====
*A:ALA-48#
```

statistics

Syntax **statistics** [*peer ip-address*]

Context show>router>msdp

Description This command displays statistics information related to a MSDP peer.

Parameters *peer ip-address* — Displays information about the specified peer IP address

Output **MSDP Statistics Output** — The following table provides MSDP statistics field descriptions.

Label	Description
Last State Change	Displays the date and time the peer state changed.
RPF Failures	Displays the number of reverse path forwarding (RPF) failures.
SA Msgs Sent	Displays the number of source-active messages sent.
SA req. Msgs Sent	Displays the number of source-active request messages sent.
SA res. Msgs Sent	Displays the number of source-active response messages sent.
KeepAlive Msgs Sent	Displays the number of keepalive messages sent.
Unknown Msgs Sent	Displays the number of unknown messages received.

Label	Description (Continued)
Last message Peer	Displays the time the last message was received from the peer.
Remote Closes	Displays the number of times the remote peer close.
SA Msgs Recvd	Displays the number of source-active messages received.
SA req. Msgs Recvd	Displays the number of source-active request messages received.
SA res. Msgs Recvd	Displays the number of source-active response messages received.
KeepAlive Msgs Recd	Displays the number of keepalive messages received.
Error Msgs Recvd	Displays the number of unknown messages received.

Sample Output

```
A:ALA-48# show router msdp statistics
=====
MSDP Statistics
=====
Glo ActSrc Lim Excd: 0
-----
Peer Address      : 10.20.1.1
-----
Last State Change : 0d 11:33:16      Last message Peer : 0d 00:00:17
RPF Failures      : 0                Remote Closes    : 0
SA Msgs Sent      : 0                SA Msgs Recvd    : 709
SA req. Msgs Sent : 0                SA req. Msgs Recvd : 0
SA res. Msgs Sent : 0                SA res. Msgs Recvd : 0
KeepAlive Msgs Sent: 694             KeepAlive Msgs Recd: 694
Unknown Msgs Sent : 0                Error Msgs Recvd  : 0
-----
Peers : 1
=====
A:ALA-48#
```

status

Syntax status

Context show>router>msdp

Description This command displays MSDP status information.

Output **MSDP Status Output** — The following table provides MSDP status field descriptions.

Label	Description
Admin State	Displays the administrative state.
Local Address	Displays the local IP address.

Label	Description (Continued)
Active Src Limit	Displays the active source limit.
Act Src Lim Excd	Displays the active source limit which has been exceeded.
Num. Peers	Displays the number of peers.
Num. Peers Estab	Displays the number of peers established.
Num. Source Active	Displays the number of active sources.
Policies	The policy to export source active state from the source active list into MSDP.
Data Encapsulation	The rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages - enabled or disabled.
Rate	The receive message rate.
Time	The receive message time.
Threshold	The number of MSDP messages that can be processed before the MSDP message rate limiting function is activated.
RPF Table	The name of the reverse path forwarding table.
Last mdsdp Enabled	The time the last MDSP was triggered.

Sample Output

```
A:ALA-48# show router mdsdp status
=====
MSDP Status
=====
Admin State           : Up
Local Address         : None
Global Statistics
Active Src Limit      : None
Act Src Lim Excd      : 0
Num. Peers            : 1
Num. Peers Estab      : 1
Num. Source Active    : 10
Policies              : None
Data Encapsulation    : Enabled
Receive Msg Rate
Rate                  : 0
Time                  : 0
Threshold             : 0
Last Mdsdp Enabled    : 08/30/2002 03:21:43
=====
A:ALA-48#
```

mcac

Syntax	mcac
Context	show>router
Description	This command enables the context to display multicast CAC related information.

policy

Syntax	policy [<i>policy-name</i> [bundle <i>bundle-name</i>] [protocol <i>protocol-name</i>] [interface <i>if-name</i>] [detail]]
Context	show>router>mcac
Description	This command displays MCAC policy information.
Parameters	<p><i>policy-name</i> — Specifies an existing multicast CAC (MCAC) policy name.</p> <p>bundle <i>bundle-id</i> — Specifies an existing multicast bundle name.</p> <p>protocol <i>protocol-name</i> — specifies an applicable protocol to display.</p> <p>Values igmp, pim, igmpSnpg</p> <p>interface <i>if-name</i> — Specifies an interface name to display.</p> <p>detail — Displays detailed information.</p>

Sample Output

```
*A:ALA-48>show>router>mcac# policy
=====
Multicast CAC Policies
=====
Policy                Description
-----
btv_fr                foreign TV offering
btv_vl                eastern TV offering
policy1               this is policy1
policy2               this is policy 2
-----
Policies : 4
=====
*A:ALA-48>show>router>mcac#

*A:ALA-48>show>router>mcac# policy btv_fr
=====
Multicast CAC policy
=====
Policy                : btv_fr
Description           : foreign TV offering
Default Action        : discard
Bundle(s)             : FOR
=====
```

Show Router PIM Commands

*A:ALA-48>show>router>mcac#

statistics

- Syntax

statistics policy *policy-name* [**bundle** *bundle-name*] [**protocol** *protocol-name*] [**interface** *if-name*] **statistics**
- Context

show>router>mcac
- Description

This command displays MCAC statistics.
- Parameters

policy-name — Specifies an existing multicast CAC (MCAC) policy name.

bundle *bundle-id* — Displays statistics for the specified existing multicast bundle name.

protocol *protocol-name* — Displays statistics for the specified applicable protocol.

Values igmp, pim, igmpSnp

interface *if-name* — Displays statistics for the specified interface name.

detail — Displays detailed information.

bindings

- Syntax

bindings active
- Context

show>router>ldp
- Description

This command displays LDP bindings information.

Sample Output

```
*A:Dut-A# show router ldp bindings active

=====
Legend: U - Label In Use,  N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        (S) - Static          (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix                Op   IngLbl  EgrLbl  EgrIntf/LspId  EgrNextHop
-----
10.20.1.1/32          Pop  131071   --      --              --
10.20.1.2/32          Push --      131071  1/1/1          10.10.1.2
10.20.1.2/32          Swap 131070  131071  1/1/1          10.10.1.2
10.20.1.2/32          Push --      262141BU 1/1/2          10.10.2.3
10.20.1.2/32          Swap 131070  262141BU 1/1/2          10.10.2.3
10.20.1.3/32          Push --      131069BU 1/1/1          10.10.1.2
10.20.1.3/32          Swap 131069  131069BU 1/1/1          10.10.1.2
10.20.1.3/32          Push --      262143   1/1/2          10.10.2.3
```


10.20.1.3/32	Swap	131069	262143	1/1/2	10.10.2.3
10.20.1.4/32	Push	--	131068	1/1/1	10.10.1.2
10.20.1.4/32	Swap	131068	131068	1/1/1	10.10.1.2
10.20.1.4/32	Push	--	262140BU	1/1/2	10.10.2.3
10.20.1.4/32	Swap	131068	262140BU	1/1/2	10.10.2.3
10.20.1.5/32	Push	--	131067BU	1/1/1	10.10.1.2
10.20.1.5/32	Swap	131067	131067BU	1/1/1	10.10.1.2
10.20.1.5/32	Push	--	262139	1/1/2	10.10.2.3
10.20.1.5/32	Swap	131067	262139	1/1/2	10.10.2.3
10.20.1.6/32	Push	--	131066	1/1/1	10.10.1.2
10.20.1.6/32	Swap	131066	131066	1/1/1	10.10.1.2
10.20.1.6/32	Push	--	262138BU	1/1/2	10.10.2.3
10.20.1.6/32	Swap	131066	262138BU	1/1/2	10.10.2.3

No. of IPv4 Prefix Active Bindings: 10

=====

=====

LDP IPv6 Prefix Bindings (Active)

=====

Prefix	Op	IngLbl	EgrLbl
EgrNextHop	EgrIf/LspId		

No Matching Entries Found

=====

=====

LDP Generic IPv4 P2MP Bindings (Active)

=====

P2MP-Id	Interface		
RootAddr	Op	IngLbl	EgrLbl
EgrNH	EgrIf/LspId		

No Matching Entries Found

=====

=====

LDP Generic IPv6 P2MP Bindings (Active)

=====

P2MP-Id	Interface		
RootAddr	Op	IngLbl	EgrLbl
EgrNH	EgrIf/LspId		

No Matching Entries Found

=====

=====

LDP In-Band-SSM IPv4 P2MP Bindings (Active)

=====

Source	Interface		
Group	Op	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			

No Matching Entries Found

=====

Show Router PIM Commands

```
=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                Interface
RootAddr                            Op          IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====
```

```
=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group                                RD          Op
RootAddr                            Interface   IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====
```

```
=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                RD          Op
RootAddr                            Interface   IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
-----
No Matching Entries Found
=====
```

*A:Dut-A# show router ldp bindings

```
=====
LDP Bindings (IPv4 LSR ID 1.1.1.1:0)
              (IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP IPv4 Prefix Bindings
=====
```

Prefix	Peer	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop
10.20.1.1/32	10.20.1.2	131071U	--	--	--
10.20.1.1/32	10.20.1.3	131071U	--	--	--
10.20.1.2/32	10.20.1.2	--	131071	1/1/1	10.10.1.2
10.20.1.2/32	10.20.1.3	131070U	262141	1/1/2	10.10.2.3
10.20.1.3/32	10.20.1.2	131069U	131069	1/1/1	10.10.1.2
10.20.1.3/32	10.20.1.3	--	262143	1/1/2	10.10.2.3
10.20.1.4/32	10.20.1.2	131068N	131068	1/1/1	10.10.1.2

10.20.1.4/32	10.20.1.3	131068BU	262140 1/1/2	10.10.2.3
10.20.1.5/32	10.20.1.2	131067U	131067 1/1/1	10.10.1.2
10.20.1.5/32	10.20.1.3	131067N	262139 1/1/2	10.10.2.3
10.20.1.6/32	10.20.1.2	131066N	131066 1/1/1	10.10.1.2
10.20.1.6/32	10.20.1.3	131066BU	262138 1/1/2	10.10.2.3

 No. of IPv4 Prefix Bindings: 12
 =====

=====

LDP IPv6 Prefix Bindings

=====

Prefix	IngLbl	EgrLbl
Peer	EgrIntf/LspId	
EgrNextHop		

No Matching Entries Found
 =====

=====

LDP Generic IPv4 P2MP Bindings

=====

P2MP-Id	Interface	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			
Peer			
100			
1.1.1.1	Unknw	--	131051
90.90.90.2	1/1/6		
2.2.2.2:0			
104			
1.1.1.1	Unknw	--	131050
90.90.90.2	1/1/6		
2.2.2.2:0			
600			
1.1.1.1	Unknw	--	131049
90.90.90.2	1/1/6		
2.2.2.2:0			
700			
1.1.1.1	Unknw	--	131048
90.90.90.2	1/1/6		
2.2.2.2:0			
800			
1.1.1.1	Unknw	--	131047
90.90.90.2	1/1/6		
2.2.2.2:0			
900			
1.1.1.1	Unknw	--	131046
90.90.90.2	1/1/6		
2.2.2.2:0			
1500			
1.1.1.1	Unknw	--	131045

Show Router PIM Commands

```
90.90.90.2          1/1/6
2.2.2.2:0

100
6.6.6.6            Unknw      --      131044
90.90.90.2         1/1/6
2.2.2.2:0

900
6.6.6.6            Unknw      --      131043
90.90.90.2         1/1/6
2.2.2.2:0

-----
No. of Generic IPv4 P2MP Bindings: 9
=====

=====
LDP Generic IPv6 P2MP Bindings
=====
P2MP-Id
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv6 P2MP Bindings
=====
Source
Group
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr          RD
EgrNH             Interface      IngLbl  EgrLbl
                  EgrIf/LspId
```

```

Peer
-----
1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

-----
No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 3
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
Source
Group          RD
RootAddr       Interface
EgrNH          EgrIf/LspId
Peer
-----
1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

-----
No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 3
=====

LDP Service FEC 128 Bindings
=====

```

Show Router PIM Commands

Type Peer	VCId SvcId	SDPId	IngLbl EgrLbl	LMTU RMTU

?-Eth 2.2.2.2:0	100 Ukwn	R. Src	-- 131023D	None 986
?-Eth 2.2.2.2:0	500 Ukwn	R. Src	-- 131022D	None 1386
?-Eth 2.2.2.2:0	2001 Ukwn	R. Src	-- 131019D	None 986
?-Eth 2.2.2.2:0	2003 Ukwn	R. Src	-- 131017D	None 986
?-Ipipe 2.2.2.2:0	1800 Ukwn	R. Src	-- 131014D	None 1486

No. of VC Labels: 5				
=====				
LDP Service FEC 129 Bindings				
=====				
SAII	AGII	IngLbl	LMTU	
TAII	Type	EgrLbl	RMTU	
Peer	SvcId	SDPId		

No Matching Entries Found				
=====				

mvpn

Syntax	mvpn
Context	show>router <i>router-instance</i>
Description	This command displays Multicast VPN related information. The router instance must be specified.

Sample Output

```
*A:Dut-C# show router 1 mvpn
=====
MVPN 1 configuration data
=====
signaling          : Bgp          auto-discovery      : Default
UMH Selection      : Highest-Ip    SA withdrawn        : Disabled
intersite-shared   : Enabled      Persist SA          : Disabled
vrf-import         : N/A
vrf-export         : N/A
vrf-target         : unicast
C-Mcast Import RT  : target:10.20.1.4:105

ipmsi              : rsvp IpmsiTemplate
```

```

i-pmsi P2MP AdmSt : Up
i-pmsi Tunnel Name : IpmsiTemplate-1-74216
enable-bfd-root    : false          enable-bfd-leaf      : false
Mdt-type           : sender-receiver

BSR signalling      : none
Wildcard s-pmsi    : false
spmsi              : rsvp SpmsiTemplate
s-pmsi P2MP AdmSt  : Up
max-p2mp-spmsi     : 4000
data-delay-interval: 3 seconds
enable-asm-mdt     : N/A
data-threshold      : 224.0.0.0/4 --> 1 kbps
rx-threshold        : 224.0.0.0/4 --> pe-thres-add 2 --> pe-thres-delete 4
data-threshold      : ff00::/8 --> 1 kbps
rx-threshold        : ff00::/8 --> pe-thres-add 2 --> pe-thres-delete 4

```

```
=====
```

```
=====
```

```
*A:Dut-D# show router 21 mvpn
```

```
=====
```

```
MVPN 21 configuration data
```

```
=====
```

```

signaling           : Bgp                auto-discovery      : Default
UMH Selection       : Highest-Ip          SA withdrawn        : Disabled
intersite-shared    : Enabled              Persist SA          : Disabled
vrf-import          : N/A
vrf-export          : N/A
vrf-target          : unicast
C-Mcast Import RT   : target:10.20.1.4:106

```

```

ipmsi               : ldp
i-pmsi P2MP AdmSt   : Up
i-pmsi Tunnel Name  : mpls-if-74217
Mdt-type            : sender-receiver

```

```

BSR signalling      : none
Wildcard s-pmsi    : false
spmsi              : ldp
s-pmsi P2MP AdmSt  : Up
max-p2mp-spmsi     : 4000
data-delay-interval: 3 seconds
enable-asm-mdt     : N/A
data-threshold      : 224.0.0.0/4 --> 1 kbps
rx-threshold        : 224.0.0.0/4 --> pe-thres-add 2 --> pe-thres-delete 4
data-threshold      : ff00::/8 --> 1 kbps
rx-threshold        : ff00::/8 --> pe-thres-add 2 --> pe-thres-delete 4

```

```
=====
```

```
*A:Dut-D#
```

tunnel-table

Syntax	tunnel-table [<i>ip-address</i> [/ <i>mask</i>]] [<i>protocol</i> sdp <i>sdp-id</i>] tunnel-table [summary]
Context	show>router
Description	This command displays tunnel table information.
Parameters	<i>protocol</i> — Specifies the protocol. Values bgp ldp rsvp sdp <i>sdp-id</i> — Specifies the SDP ID. Values 1..17407
Output	<pre>*A:Dut-C# show router tunnel-table sdp 17407 ===== Tunnel Table (Router: Base) ===== Destination Owner Encap TunnelId Pref Nexthop Metric ----- 127.0.68.0/32 sdp MPLS 17407 5 127.0.68.0 0 ===== *A:Dut-C>config>router>mpls>lsp# show router tunnel-table detail ===== Tunnel Table (Router: Base) ===== Destination : 1.0.0.2/32 NextHop : 1.1.4.4 Tunnel Flags : exclude-for-lfa Age : 00h17m58s Owner : rsvp Encap : MPLS Tunnel ID : 115 Preference : 7 Tunnel Label : 262054 Tunnel Metric : 9 Tunnel MTU : 1496 LSP ID : 26116 Bypass Label : 0 LSP Bandwidth : 0 LSP Weight : 2 -----</pre>

Clear Commands

database

Syntax **database** [**interface** *ip-int-name*|*ip-address*] **group** *grp-ip-address* [**source** *src-ip-address*]
database **grp-interface** *interface-name* [**fwd-service** *service-id*]
database [**interface** *ip-int-name*|*ip-address*] **group** *grp-ip-address* **source** *src-ip-address*
database **host** [*ip-address*]
database **interface** *ip-int-name*|*ip-address* [**group** *grp-ip-address*] [**source** *src-ip-address*]

Context clear>router>igmp

Description This command clears IGMP or PIM database statistics on a specified interface or IP address.

Parameters **interface** *ip-int-name* — Clears the IGMP or PIM database on the specified interface.
interface *ip-address* — Clears the IGMP or PIM database on the specified IP address.
group *group-ip-address* — Clears the multicast group address(ipv4/ipv6) or zero in the specified address group.
source *ip-address* — Clears the IGMP or PIM database from the specified source IP address.

database

Syntax **database** [**interface** *ip-int-name*|*mt-int-name*|*int-ip-address*] [**group** *grp-ip-address* [**source** *ip-address*]] [*family*]

Context clear>router>pim

Description This command clears IGMP or PIM database statistics on a specified interface or IP address.

Parameters **interface** *ip-int-name* — Clears the IGMP or PIM database on the specified interface.
interface *mt-int-name* — Clears the default core group address of the Multicast Distribution Tree (MDT) for the VPRN instance. The Multicast Tunnel (MT) interface for a VPRN is created when this object is set to a valid group address.
Syntax: *vprn-id-mt-grp-ip-address*
interface *ip-address* — Clears the IGMP or PIM database on the specified IP address.
group *group-ip-address* — Clears the multicast group address(ipv4/ipv6) or zero in the specified address group.
source *ip-address* — Clears the IGMP or PIM database from the specified source IP address.
family — Clears either IPv4 or IPv6 information.
mpls-if-name — Clears the MPLS interface name.

Syntax: *mpls-if-index*

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>igmp

Description This command clears IGMP statistics on a specified interface or IP address.
Note that interface and group/source cannot be specified at the same time.

Parameters **interface** *ip-int-name* — Clears IGMP statistics on the specified interface.
interface *ip-address* — Clears IGMP statistics on the specified IP address.
interface *mt-int-name* — Clears the default core group address of the Multicast Distribution Tree (MDT) for the VPRN instance. The Multicast Tunnel (MT) interface for a VPRN is created when this object is set to a valid group address.

Syntax: *vprn-id-mt-grp-ip-address*

s-pmsi

Syntax **s-pmsi** [*mdSrcAddr*] [*mdGrpAddr*] [*vprnSrcAddr* *vprnGrpAddr*]

Context clear>router>pim

Description This command clears PIM selective provider multicast service interface cache.

Parameters *mdSrcAddr* — Clears the specified source address used for Multicast Distribution Tree (MDT).
mdGrpAddr — Clears the specified group address used for Multicast Distribution Tree (MDT).
vprnSrcAddr — Clears the specified source address of the multicast sender.
vprnGrpAddr — Clears the specified multicast group address.

statistics

Syntax **statistics** [[[**interface** *ip-int-name* | *ip-address* | *mt-int-name*]]] [[**group** *grp-ip-address* [**source** *ip-address*]]] [*family*]

Context clear>router>pim

Description This command clears PIM statistics on a specified interface or IP address.
Note that an interface and group or source cannot be specified at the same time.

Parameters **interface** *ip-int-name* — Clears PIM statistics on the specified interface.
interface *ip-address* — Clears PIM statistics on the specified IP address.

interface *mt-int-name* — Clears the default core group address of the Multicast Distribution Tree (MDT) for the VPRN instance. The Multicast Tunnel (MT) interface for a VPRN is created when this object is set to a valid group address.

syntax: *vprn-id-mt-grp-ip-address*

group *grp-ip-address* — When only the group address is specified and no source is specified, (*,G) statistics are cleared. When the group address is specified along with the source address, then the (S,G) statistics are reset to zero.

source *ip-address* — When the source address is specified along with the group address, then the (S,G) statistics are reset to zero.

family — Clears either IPv4 or IPv6 information.

version

Syntax **version** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>igmp

Description This command clears IGMP statistics on a specified interface or IP address.

Parameters **interface** *ip-int-name* — Clears IGMP or PIM statistics on the specified interface.
interface *ip-address* — Clears IGMP or PIM statistics on the specified IP address.

mld

Syntax **mld**

Context clear>router

Description This command enables the context to to clear and reset Multicast Listener Discovery (MLD) entities.

database

Syntax **database** [**interface** *ip-int-name*|*ipv6-address*] [**group** *ip-address* [**source** *ip-address*]]

Context clear>router>mld

Description This command clears Multicast Listener Discovery (MLD) database parameters.

Parameters **interface** *ip-int-name* — Clears database information for the specified Multicast Listener Discovery (MLD) interface name.
interface *ipv6-address* — Clears database information for the specified Multicast Listener Discovery (MLD) interface IPv6 address.

Clear Commands

group *ip-address* — Clears database information for the specified Multicast Listener Discovery (MLD) group IP address.

source *ip-address* — Clears database information for the specified Multicast Listener Discovery (MLD) source IP address.

statistics

Syntax **statistics** [*ip-int-name* | *ipv6-address*]

Context clear>router>mld

Description This command clears Multicast Listener Discovery (MLD) statistics parameters.

Parameters *ip-int-name* — Clears statistics for the specified Multicast Listener Discovery (MLD) interface name.
ipv6-address — Clears statistics for the specified Multicast Listener Discovery (MLD) IPv6 address.

version

Syntax **version** [*ip-int-name* | *ip-address*]

Context clear>router>mld

Description This command clears Multicast Listener Discovery (MLD) version parameters.

Parameters *ip-int-name* — Clears version information for the specified Multicast Listener Discovery (MLD) interface name.
ip-address — Clears version information for the specified Multicast Listener Discovery (MLD) IP address.

msdp

Syntax **msdp**

Context clear>router

Description This command enables the context to clear and reset Multicast Source Discovery protocol (MSDP) entities and statistics.

cache

Syntax **cache** [**peer** *ip-address*] [**group** *ip-address*] [**source** *ip-address*] [**originrp** *ip-address*]

Context clear>router>msdp

Description This command clears the MSDP cache.

Parameters

peer *ip-address* — Clears the cache of the IP address of the peer to which Multicast Source Discovery protocol (MSDP) source-active (SA) requests for groups matching this entry's group range were sent.

group *ip-address* — Clears the group IP address of the SA entry.

source *ip-address* — Clears the source IP address of the SA entry.

originrp *ip-address* — Clears the origin rendezvous point(RP) address type of the SA entry.

statistics

Syntax **statistics** [**peer** *ip-address*]

Context clear>router>msdp

Description **peer** *ip-address* — Clears the statistics of the IP address of the peer to which Multicast Source Discovery Protocol (MSDP) source-active (SA) requests for groups matching this entry's group range were sent.

neighbor

Syntax **neighbor** [*ip-int-name* | *ip-address*] [*family*]

Context clear>router>pim

Description This command clears PIM neighbor data on a specified interface or IP address.

Parameters

ip-int-name — Clears PIM neighbor on the specified interface.

ip-address — Clears PIM neighbor on the specified IP address.

family — Clears either IPv4 or IPv6 information.

igmp-snooping

Syntax **igmp-snooping**

Context clear>service>id

Description This command enables the context to clear IGMP snooping-related data.

port-db

Syntax **port-db** {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address* [**source** *ip-address*]]

Context clear>service>id>igmp-snooping

Description Clears the information on the IGMP snooping port database.

Clear Commands

Parameters **sap** *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. The *sap-id* can be in one of the following formats:

Encapsulation type	Syntax	Example
null	port-id	1/1/3
dot1q	port-id :qtag1	1/1/3:100
qinq	port-id :qtag1.qtag2	1/1/3:100.200

qtag1, *qtag2* — The encapsulation value on the specified port ID.

Values 0 — 4094

sdp *sdp-id* — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear information.

Default For mesh SDPs only, all VC IDs

Values 1 — 4294967295

group *grp-address* — Clears IGMP snooping statistics matching the specified group address.

source *ip-address* — Clears IGMP snooping statistics matching one particular source within the multicast group.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description Clears information on the IGMP snooping queriers for the VPLS service.

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context clear>service>id>igmp-snooping

Description Clears IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* — Displays IGMP snooping statistics for a specific SAP. The *sap-id* can be in one of the following formats:

Encapsulation type	Syntax	Example
null	port-id	1/1/3
dot1q	port-id :qtag1	1/1/3:100
qinq	port-id :qtag1.qtag2	1/1/3:100.200

qtag1, qtag2 — The encapsulation value on the specified port ID.

Values 0 — 4094

sdp *sdp-id* — Displays the IGMP snooping statistics for a specific spoke or mesh SDP.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Default For mesh SDPs only, all VC IDs

Values 1 — 4294967295

pim-snooping

Syntax **pim-snooping**

Context clear>service>id

Description This command enables the context to clear PIM snooping information.

database

Syntax **database** [[**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** *grp-ip-address*] [**source** *src-ip-address*]]

Context clear>service>id>pim-snooping

Description This command clears PIM snooping source group database information.

Parameters **sap** *sap-id* — Clears PIM snooping SAP information.

sdp *sdp-id* — Clears PIM snooping entries associated with the specified SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 — 17407

group *grp-address* — Clears PIM snooping information matching the specified group address.

source *ip-address* — Clears PIM snooping information matching one particular source within the multicast group.

neighbor

Syntax **neighbor** [*ip-address* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context clear>service>id>pim-snooping

Description This comand clears PIM snooping neighbor information.

Parameters *ip-address* — Clears IP address information.

Clear Commands

sap *sap-id* — Clears PIM snooping SAP information.

sdp *sdp-id* — Clears PIM snooping entries associated with the specified SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 — 17407

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context clear>service>id>pim-snooping

Description This command clears PIM snooping statistics for the specified SAP or SDP.

Parameters **sap** *sap-id* — Clears PIM snooping SAP information.

sdp *sdp-id* — Clears PIM snooping entries associated with the specified SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 — 17407

Debug Commands

Debug IGMP Commands

group-interface

Syntax [no] group-interface [fwd-service service-id] [ip-int-name]

Context debug>router>igmp

Description This command enables debugging for IGMP group-interface.
The **no** form of the command disables debugging.

host

Syntax host [ip-address]
host [fwd-service service-id] group-interface ip-int-name
no host [ip-address]
no host [fwd-service service-id] group-interface ip-int-name

Context debug>router>igmp

Description This command enables debugging for the IGMP host.
The **no** form of the command disables debugging.

interface

Syntax [no] interface [ip-int-name | ip-address]

Context debug>router>igmp

Description This command enables debugging for IGMP interfaces.
The **no** form of the command disables the IGMP interface debugging for the specifies interface name or IP address.

Parameters *ip-int-name* — Only displays the information associated with the specified IP interface name.
ip-address — Only displays the information associated with the specified IP address.

Debug IGMP Commands

mcs

Syntax	mcs <i>[ip-int-name]</i> no mcs
Context	debug>router>igmp
Description	This command enables debugging for IGMP multicast servers (MCS). The no form of the command disables the IGMP interface debugging for the specifies interface name.
Parameters	<i>ip-int-name</i> — Only displays the information associated with the specified IP interface name.

misc

Syntax	[no] misc
Context	debug>router>igmp
Description	This command enables debugging for IGMP miscellaneous. The no form of the command disables the debugging.

Sample Output

```
A:ALA-CA# debug router 100 igmp misc
*A:ALA-CA# show debug
debug
  router "100"
    igmp
      misc
    exit
  exit
exit
*A:ALA-CA#
```

packet

Syntax	packet [query v1-report v2-report v3-report v2-leave] host <i>ip-address</i> packet [query v1-report v2-report v3-report v2-leave] [<i>ip-int-name</i> <i>ip-address</i>] no packet [query v1-report v2-report v3-report v2-leave] [<i>ip-int-name</i> <i>ip-address</i>] no packet [query v1-report v2-report v3-report v2-leave] host <i>ip-address</i>
Context	debug>router>igmp
Description	This command enables/disables debugging for IGMP packets.
Parameters	query — Specifies to log the IGMP group- and source-specific queries transmitted and received on this interface. v1-report — Specifies to log IGMP V1 reports transmitted and received on this interface.

v2-report — Specifies to log IGMP V2 reports transmitted and received on this interface.

v3-report — Specifies to log IGMP V3 reports transmitted and received on this interface.

v2-leave — Specifies to log the IGMP Leaves transmitted and received on this interface.

ip-int-name — Only displays the information associated with the specified IP interface name.

ip-address — Only displays the information associated with the specified IP address.

Debug PIM Commands

adjacency

Syntax	[no] adjacency
Context	debug>router>pim
Description	This command enables/disables debugging for PIM adjacencies.

all

Syntax	all [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no all
Context	debug>router>pim
Description	This command enables/disables debugging for all the PIM modules.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with all PIM modules. Values IPv4 or IPv6 address source <i>ip-address</i> — Debugs information associated with all PIM modules. Values IPv4 or IPv6 address detail — Debugs detailed information on all PIM modules.

assert

Syntax	assert [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no assert
Context	debug>router>pim
Description	This command enables/disables debugging for PIM assert mechanism.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with the PIM assert mechanism. Values multicast group address (ipv4/ipv6) source <i>ip-address</i> — Debugs information associated with the PIM assert mechanism. Values source address (ipv4/ipv6) detail — Debugs detailed information on the PIM assert mechanism.

bsr

Syntax	bsr [detail] no bsr
Context	debug>router>pim
Description	This command enables debugging for PIM Bootstrap mechanism. The no form of the command disables debugging.
Parameters	detail — Debugs detailed information on the PIM assert mechanism.

data

Syntax	data [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no data
Context	debug>router>pim
Description	This command enables/disables debugging for PIM data exception.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with the specified data exception. Values multicast group address (ipv4/ipv6) source <i>ip-address</i> — Debugs information associated with the specified data exception. Values source address (ipv4/ipv6) detail — Debugs detailed IP data exception information.

db

Syntax	db [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no db
Context	debug>router>pim
Description	This command enables/disables debugging for PIM database.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with the specified database. Values multicast group address (ipv4/ipv6) or zero source <i>ip-address</i> — Debugs information associated with the specified database. Values source address (ipv4/ipv6) detail — Debugs detailed IP database information.

interface

Syntax	interface [<i>ip-int-name</i> <i>mt-int-name</i> <i>ip-address</i>] [detail] no interface
Context	debug>router>pim
Description	This command enables/disables debugging for PIM interface.
Parameters	<i>ip-int-name</i> — Debugs the information associated with the specified IP interface name. Values IPv4 or IPv6 interface address <i>mt-int-address</i> — Debugs the information associated with the specified VPRN ID and group address. <i>ip-address</i> — Debugs the information associated with the specified IP address. detail — Debugs detailed IP interface information.

jp

Syntax	jp [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no jp
Context	debug>router>pim
Description	This command enables/disables debugging for PIM Join-Prune mechanism.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with the specified Join-Prune mechanism. Values multicast group address (ipv4/ipv6) or zero source <i>ip-address</i> — Debugs information associated with the specified Join-Prune mechanism. Values source address (ipv4/ipv6) detail — Debugs detailed Join-Prune mechanism information.

mrib

Syntax	mrib [group <i>grp-ip-address</i>] [source <i>ip-address</i>] [detail] no mrib
Context	debug>router>pim
Description	This command enables/disables debugging for PIM MRIB.
Parameters	group <i>grp-ip-address</i> — Debugs information associated with the specified PIM MRIB. Values multicast group address (ipv4/ipv6) source <i>ip-address</i> — Debugs information associated with the specified PIM MRIB. Values source address (ipv4/ipv6)

detail — Debugs detailed MRIB information.

msg

Syntax	msg [detail] no msg
Context	debug>router>pim
Description	This command enables/disables debugging for PIM messaging.
Parameters	detail — Debugs detailed messaging information.

packet

Syntax	packet [hello register register-stop jp bsr assert crp] [ip-int-name ip-address] no packet
Context	debug>router>pim
Description	This command enables/disables debugging for PIM packets.
Parameters	hello register register-stop jp bsr assert crp — PIM packet types. <i>ip-int-name</i> — Debugs the information associated with the specified IP interface name. <div style="margin-left: 40px;">Values IPv4 or IPv6 interface address</div> <i>ip-address</i> — Debugs the information associated with the specified IP address of a particular packet type.

register

Syntax	register [group grp-ip-address] [source ip-address] [detail] no register
Context	debug>router>pim
Description	This command enables/disables debugging for PIM Register mechanism.
Parameters	group grp-ip-address — Debugs information associated with the specified PIM register. <div style="margin-left: 40px;">Values multicast group address (ipv4/ipv6)</div> source ip-address — Debugs information associated with the specified PIM register. <div style="margin-left: 40px;">Values source address (ipv4/ipv6)</div> detail — Debugs detailed register information.

Debug PIM Commands

rtm

Syntax	rtm [detail] no rtm
Context	debug>router>pim
Description	This command enables/disables debugging for PIM RTM.
Parameters	detail — Debugs detailed RTM information.

s-pmsi

Syntax	s-pmsi [{vpnSrcAddr [vpnGrpAddr]} [mdSrcAddr]] [detail] no s-pmsi
Context	debug>router>pim
Description	This command enables debugging for PIM selective provider multicast service interface. The no form of the command disables the debugging.
Parameters	<i>vpnSrcAddr</i> — Specifies the VPN source address. <i>vpnGrpAddr</i> — Specifies the VPN group address <i>mdSrcAddr</i> — Specifies the source address of the multicast sender. detail — Displays detailed information for selective PMSI.

msdp

Syntax	[no] msdp
Context	debug>router
Description	This command enables debugging for Multicast Source Discovery Protocol (MSDP). The no form of the command disables MSDP debugging.

packet

Syntax	packet [pkt-type] [peer ip-address]
Context	debug>router>msdp
Description	This command enables debugging for Multicast Source Discovery Protocol (MSDP) packets. The no form of the command disables MSDP packet debugging.

Parameters *pkt-type* — Debugs information associated with the specified packet type.

Values keep-alive, source-active, sa-request, sa-response

peer *ip-address* — Debugs information associated with the specified peer IP address.

pim

Syntax **pim** [*grp-address*]
no pim

Context debug>router>msdp

Description This command enables debugging for Multicast Source Discovery Protocol (MSDP) PIM. The **no** form of the command disables MSDP PIM debugging.

Parameters *grp-address* — Debugs the IP multicast group address for which this entry contains information.

rtm

Syntax **rtm** [*rp-address*]
no rtm

Context debug>router>msdp

Description This command enables debugging for Multicast Source Discovery Protocol (MSDP) route table manager (RTM). The **no** form of the command disables MSDP RTM debugging.

Parameters *rp-address* — Debugs the IP multicast address for which this entry contains information.

sa-db

Syntax **sa-db** [**group** *grpAddr*] [**source** *srcAddr*] [**rp** *rpAddr*]
no sadb

Context debug>router>msdp

Description This command enables debugging for Multicast Source Discovery Protocol (MSDP) source-active requests. The **no** form of the command disables the MSDP source-active database debugging.

Parameters **group** *grpAddr* — Debugs the IP address of the group.
source *srcAddr* — Debugs the source IP address.
rp *rpAddr* — Debugs the specified rendezvous point RP address.

In This Chapter

This chapter provides information about configuring Routing Information Protocol (RIP) parameters.

Topics in this chapter include:

- [RIP Overview on page 248](#)
 - [RIP Features on page 249](#)
 - [RIP Version Types on page 249](#)
 - [RIPv2 Authentication on page 249](#)
 - [RIP Packet Format on page 250](#)
- [RIPng on page 252](#)
 - [RIPng Protocol on page 252](#)
- [Common Attributes on page 253](#)
 - [Metrics on page 253](#)
 - [Timers on page 253](#)
 - [Import and Export Policies on page 254](#)
 - [RIP Packet Format on page 250](#)
- [RIP Configuration Process Overview on page 255](#)
- [Configuration Notes on page 256](#)

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, 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.

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 address. 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.

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 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.

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.

RIP Packet Format

The RIP packet format is displayed in [Figure 6](#):




Figure 6: 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.

RIPv1 Format

There can be between 1 and 25 (inclusive) RIP entries. [Figure 7](#) displays RIPv1 format:

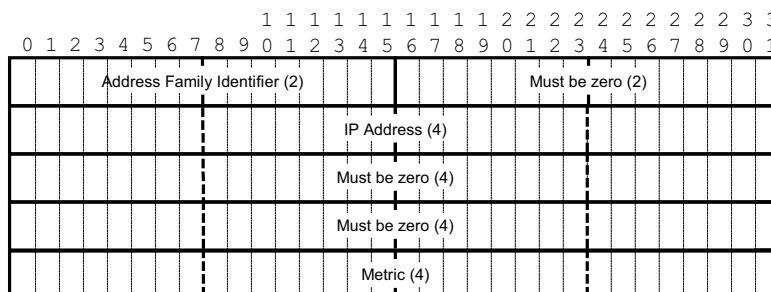


Figure 7: RIPv1 Format

RIPv2 Format

The RIP packet format is displayed in [Figure 8](#):

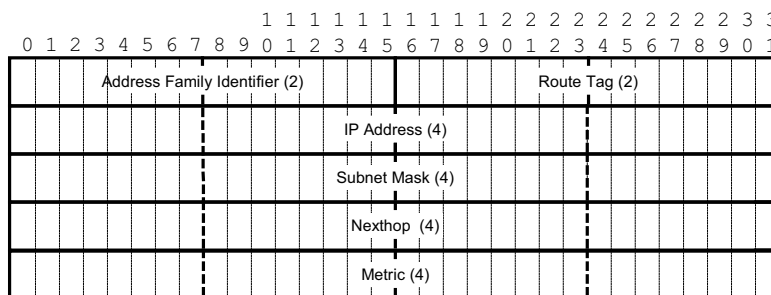


Figure 8: RIPv2 Format

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.

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.

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
 - Ø The destination IP for any periodic or triggered update should be sent to the multicast group FF02::9, (all-rip-routers multicast group)
 - Ø When sending responses to an RIPng request, the RIPng response is sent to the unicast IP address of the requestor.
- Each route entry in an update message contains the following:
 - Ø IPv6 prefix
 - Ø Prefix length
 - Ø Route metric
 - Ø (Optional) route tag

Common Attributes

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).

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.

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.

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.

RIP Configuration Process Overview

Figure 9 displays the process to configure RIP parameters.

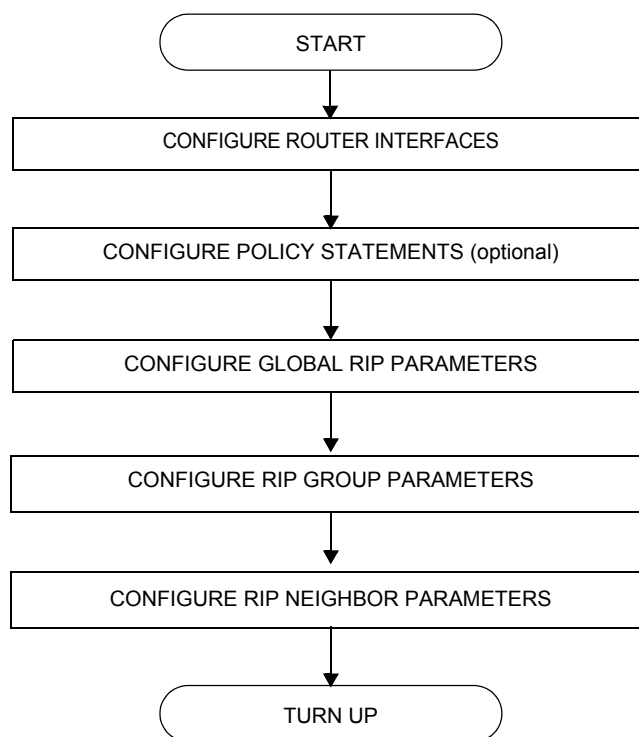


Figure 9: RIP Configuration and Implementation Flow

Configuration Notes

This section describes RIP configuration caveats.

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.

Configuring RIP with CLI

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

Topics in this section include:

- [RIP & RIPng Configuration Overview on page 258](#)
- [Basic RIP Configuration on page 259](#)
- [Common Configuration Tasks on page 260](#)
 - [Configuring Interfaces on page 261](#)
 - [Configuring a Route Policy on page 262](#)
 - [Configuring RIP Parameters on page 264](#)
 - [Configuring Global-Level Parameters on page 266](#)
 - [Configuring Group-Level Parameters on page 267](#)
 - [Configuring Neighbor-Level Parameters on page 268](#)
- [RIP Configuration Management Tasks on page 269](#)
 - [Modifying RIP Parameters on page 269](#)
 - [Deleting a Group on page 270](#)
 - [Deleting a Neighbor on page 270](#)

RIP & RIPng Configuration Overview

Preconfiguration Requirements

Configure the following entities before beginning the RIP configuration:

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

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.

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#
```

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:

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

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 *IP Router Configuration Overview* chapter.

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#
```

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|reject}
    entry entry-id
      description text
      action {accept|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#

Configuring RIP Parameters

Use the CLI syntax displayed below for:

- [Configuring Global-Level Parameters on page 266](#)
- [Configuring Group-Level Parameters on page 267](#)
- [Configuring Neighbor-Level Parameters on page 268](#)

CLI Syntax: `config>router`
`rip`
 `authentication-key [authentication-key|hash-key`
 `[hash|hash2]`
 `authentication-type {none|password|message-digest|mes-`
 `sage-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
```

Configuring Global-Level Parameters

Once the RIP protocol instance is created, the `no shutdown` command is not required since RIP is administratively enabled upon creation. Minimally, 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 "ac18651vz1d" hash
timers 300 600 600
-----
ALA-A>config>router>rip#
```

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 "ac1865lvzld" hash
timers 300 600 600
group "headquarters"
      description "Mt. View"
exit
-----
ALA-A>config>router>rip#
```

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#
```


RIP Configuration Management Tasks

Examples are provided for the following RIP configuration management tasks:

- [Modifying RIP Parameters on page 269](#)
 - [Deleting a Group on page 270](#)
 - [Deleting a Neighbor on page 270](#)
-

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 on page 264](#).

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 "ac1865lvzld" 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#
```

Deleting a Group

A group must be shut down first in order to delete it.

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"

If you try to delete the group without shutting it down first, the following message appears:

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

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"

If you try to delete the neighbor before it is shut down, the following message appears:

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

RIP Command Reference

Command Hierarchies

- [Configuration Commands on page 271](#)
 - [Group Commands on page 272](#)
 - [Neighbor Commands on page 273](#)
- [RIPng Configuration Commands on page 274](#)
 - [Group Commands on page 275](#)
 - [Neighbor Commands on page 276](#)
- [Show RIP Commands on page 277](#)
- [Clear RIP Commands on page 277](#)
- [Debug RIP Commands on page 277](#)

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
      — check-zero {enable | disable}
      — no check-zero
      — description string
      — no description
      — export policy-name [policy-name ...(up to 5 max)]
      — no export
      — export-limit number [log percentage]
      — no export-limit
      — import policy-name [policy-name ...(up to 5 max)]
      — 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

```

Command Hierarchies

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

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**
 - **check-zero** {enable | disable}
 - **no check-zero**
 - **description** *description-string*
 - **no description**
 - **export** *policy-name* [*policy-name* ... (up to 5 max)]
 - **no export**
 - **import** *policy-name* [*policy-name* ... (up to 5 max)]
 - **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**

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}
                — no authentication-type
                — check-zero {enable | disable}
                — no check-zero
                — description description-string
                — no description
                — export policy-name [policy-name ...(up to 5 max)]
                — no export
                — import policy-name [policy-name ...(up to 5 max)]
                — 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

```

RIPng Configuration Commands

```
config
  — router router-name
    — [no] ripng
      — check-zero {enable | disable}
      — no check-zero
      — description string
      — no description
      — export policy-name [policy-name ...(up to 5 max)]
      — no export
      — export-limit number [log percentage]
      — no export-limit
      — import policy-name [policy-name ...(up to 5 max)]
      — 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
```

Group Commands

```

config
— router router-name
— [no] ripng
— [no] group group-name
— check-zero {enable | disable}
— no check-zero
— description description-string
— no description
— export policy-name [policy-name ...(up to 5 max)]
— no export
— import policy-name [policy-name ...(up to 5 max)]
— 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

```

Command Hierarchies

Neighbor Commands

```
config
— router router-name
  — [no] ripng
    — [no] group group-name
      — [no] neighbor ip-int-name
        — check-zero {enable | disable}
        — no check-zero
        — description description-string
        — no description
        — export policy-name [policy-name ...(up to 5 max)]
        — no export
        — import policy-name [policy-name ...(up to 5 max)]
        — 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
```


Show RIP Commands

```
show
  — router
    — rip
    — ripng
      — database [ip-prefix [/mask>] [longer] [peer ip-address] [detail]
      — group [name] [detail]
      — neighbors [ip-int-name | ip-addr] [detail] [advertised-routes]
      — peer [interface-name]
      — statistics [ip-int-name | ip-addr]
```

Clear RIP Commands

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

Debug RIP Commands

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

debug
  — router
    — ripng
      — [no] error [neighbor ip-int-name | ip-addrres]
      — [no] events [neighbor ip-int-name | ip-addrres]
      — [no] holddown [neighbor ip-int-name | ip-addrres]
      — [no] packets [neighbor ip-int-name | ip-addrres]
      — [no] request [neighbor ip-int-name | ip-addrres]
      — [no] trigger [neighbor ip-int-name | ip-addrres]
      — [no] updates [neighbor ip-int-name | ip-addrres]
```

RIP Configuration Commands

Generic Commands

description

Syntax	description <i>string</i> no description
Context	config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of the command removes any description string from the context.
Default	no description — no description associated with the configuration context.
Parameters	<i>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, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax	[no] shutdown
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	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. The shutdown 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. Unlike other commands and parameters where the default state will not be indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

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.

RIP Commands

rip

Syntax	[no] rip
Context	config>router
Description	<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 [no] shutdown command.</p> <p>The no form of the command deletes the RIP protocol instance removing all associated configuration parameters.</p>
Default	no rip — No RIP protocol instance defined.

ripng

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

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<p>This command sets the authentication password to be passed between RIP neighbors.</p> <p>The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed.</p>

RIP Commands

The **no** form of the command removes the authentication password from the configuration and disables authentication.

Default **no authentication-key** — No authentication key configured.

Parameters *authentication-key* — 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.

hash-key — 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 (“ ”).

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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax **authentication-type {none|password|message-digest|message-digest-20}**
no authentication-type

Context config>router>rip
config>router>rip>group *group-name*
config>router>rip>group *group-name*>neighbor *ip-int-name*

Description This command sets the type of authentication to be used between RIP neighbors.

The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication type from the configuration and effectively disables authentication.

Default **no authentication-type** — No authentication enabled.

Parameters **none** — The **none** parameter explicitly disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.

password — Specify password to enable simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest — Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one **message-digest-key** must be configured.

message-digest-20 — Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, *RIP-2 MD5 Authentication*. If this option is configured, then at least one **message-digest-key** must be configured.

check-zero

Syntax	check-zero {enable disable} no check-zero
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<p>This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.</p> <p>The check-zero enable command enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages.</p> <p>The check-zero disable command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.</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 (no check-zero), the setting from the less specific level is inherited by the lower level.</p> <p>The no form of the command removes the check-zero command from the configuration.</p>
Special Cases	RIP Global — By default, check-zero is disabled at the global RIP instance level.
Parameters	<p>enable — Specifies reject RIP messages which do not have zero in the RIPv1 and RIPv2 mandatory fields.</p> <p>disable — Specifies allows receipt of RIP messages which do not have the mandatory zero fields reset.</p>

export

Syntax	export <i>policy-name</i> [<i>policy-name</i> ...up to 5 max] no export
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<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>

RIP Commands

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.

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

Default **no export** — No export route policies specified.

Parameters *policy-name* — The export 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, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

export-limit

Syntax **export-limit** *number* [**log** *percentage*]
no export-limit

Context config>router>rip
config>router>ripng

Description This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

The **no** form of the command removes the parameters from the configuration.

Default no export-limit, the export limit for routes or prefixes is disabled..

Parameters *number* — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 — 4294967295

log percentage — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 — 100

group

Syntax [**no**] **group** *group-name*

Context config>router>rip
config>router>ripng

Description This command creates a context for configuring a RIP group of neighbor interfaces.

RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** 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.

Default **no group** — No group of RIP neighbor interfaces defined.

Parameters *group-name* — 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, etc.), the entire string must be enclosed within double quotes.

import

Syntax **import** *policy-name* [*policy-name* ...up to 5 max]
no import

Context config>router>rip
 config>router>rip>group *group-name*
 config>router>rip>group *group-name*>neighbor *ip-int-name*
 config>router>ripng
 config>router>ripng>group *group-name*
 config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description 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.

Default **no import** — No import route policies specified.

Parameters *policy-name* — 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, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

message-size

Syntax **message-size** *max-num-of-routes*
no message-size

Context config>router>rip
 config>router>rip>group *group-name*
 config>router>rip>group *group-name*>neighbor *ip-int-name*
 config>router>ripng
 config>router>ripng>group *group-name*
 config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description This command configures the maximum number of routes per RIP update message.

RIP Commands

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

Default **message-size 25** — A maximum of 25 routes per RIP update message.

Parameters *max-num-of-routes* — The maximum number of RIP routes per RIP update message expressed as a decimal integer.

Values 25 — 255

metric-in

Syntax **metric-in** *metric*
no metric-in

Context config>router>rip
config>router>rip>group *group-name*
config>router>rip>group *group-name*>neighbor *ip-int-name*
config>router>ripng
config>router>ripng>group *group-name*
config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description This command configures the metric added to routes received from a RIP neighbor.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

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

Default **metric-in 1** — Add 1 to the metric of routes received from a RIP neighbor.

Parameters *metric* — The value added to the metric of routes received from a RIP neighbor expressed as a decimal integer.

Values 1 — 16

metric-out

Syntax **metric-out** *metric*
no metric-out

Context config>router>rip
config>router>rip>group *group-name*
config>router>rip>group *group-name*>neighbor *ip-int-name*
config>router>ripng
config>router>ripng>group *group-name*
config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description This command configures the metric assigned to routes exported into RIP and advertised to RIP neighbors.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

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

Default metric-out 1 — Routes exported from non-RIP sources are given a metric of 1.

Parameters *metric* — The value added to the metric for routes exported into RIP and advertised to RIP neighbors expressed as a decimal integer.

Values 1 — 16

neighbor

Syntax [**no**] **neighbor** *ip-int-name*

Context config>router>rip>group *group-name*
config>router>ripng>group *group-name*

Description This command creates a context for configuring a RIP neighbor interface.

By default, interfaces are not activated in any interior gateway protocol, such as RIP, unless explicitly configured.

The **no** form of the command deletes the RIP interface configuration for this interface. The **shutdown** command in the config>router>rip>group *group-name*>neighbor *ip-int-name* context can be used to disable an interface without removing the configuration for the interface.

Default **no neighbor** — No RIP interfaces defined.

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, etc.), 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.

preference

Syntax **preference** *preference*
no preference

Context config>router>rip
config>router>rip>group *group-name*
config>router>rip>group *group-name*>neighbor *ip-int-name*
config>router>ripng
config>router>ripng>group *group-name*
config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description This command configures the preference for RIP routes.

RIP Commands

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 [Table 4](#). 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 **preference 100** — Preference of 100 for RIP routes.

Parameters *preference* — The preference for RIP routes expressed as a decimal integer. Defaults for different route types are listed in [Table 4](#).

Table 4: 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
TMS	167	No
Values	0 — 255	

receive

Syntax	receive {both none version-1 version-2} no receive
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<p>This command configures the type(s) of RIP updates that will be accepted and processed.</p> <p>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.</p> <p>If version-1 is specified, the router only listens for and accept packets sent to the broadcast address.</p> <p>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.</p> <p>The no form of the command reverts to the default value.</p>
Default	receive both
Parameters	<p>both — Specifies that RIP updates in either version 1 or version 2 format will be accepted.</p> <p>none — Specifies that RIP updates will not be accepted.</p> <p>version-1 — Specifies that RIP updates in version 1 format only will be accepted.</p> <p>version-2 — Specifies that RIP updates in version 2 format only will be accepted.</p>

send

Syntax	send {broadcast multicast none version-1} no send
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<p>This command specifies the type of RIP messages sent to RIP neighbors.</p> <p>If version-1 is specified, the router need only listen for and accept packets sent to the broadcast address.</p> <p>This control can be issued at the global, group or interface level.</p> <p>The no form of the command reverts to the default value.</p>
Default	send broadcast — RIPv2 formatted messages will be sent to the broadcast address.
Parameters	<p>broadcast — Specifies send RIPv2 formatted messages to the broadcast address.</p> <p>multicast — Specifies send RIPv2 formatted messages to the multicast address.</p>

RIP Commands

none — Specifies not to send any RIP messages (i.e. silent listener).

version-1 — Specifies send RIPv1 formatted messages to the broadcast address.

send

Syntax	send {none ripng unicast} no send
Context	config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<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 no form of the command reverts to the default value.</p>
Default	send ripng — RIPng formatted messages will be sent to the RIPng IPv6 multicast address.
Parameters	<p>ripng — Specifies RIPng messages to be sent to the standard multicast address (FF02::9).</p> <p>none — Specifies not to send any RIPng messages (i.e. silent listener).</p> <p>unicast — Specifies to send RIPng updates as unicast messages to the defined unicast address configured through the unicast-address command. This option is only allowed within the neighbor context.</p>

split-horizon

Syntax	split-horizon {enable disable} no split-horizon
Context	config>router>rip config>router>rip>group <i>group-name</i> config>router>rip>group <i>group-name</i> >neighbor <i>ip-int-name</i> config>router>ripng config>router>ripng>group <i>group-name</i> config>router>ripng>group <i>group-name</i> >neighbor <i>ip-int-name</i>
Description	<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 split-horizon disable 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>

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 split-horizon**), the setting from the less specific level is inherited by the lower level.

The **no** form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.

Default **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 *group-name*
 config>router>rip>group *group-name*>neighbor *ip-int-name*
 config>router>ripng
 config>router>ripng>group *group-name*
 config>router>ripng>group *group-name*>neighbor *ip-int-name*

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** — RIP update timer set to 30 seconds, timeout timer to 180 seconds and flush timer to 120 seconds.

Parameters *update* — The RIP update timer value in seconds expressed as a decimal integer.

Values 1 — 600

timeout — The RIP timeout timer value in seconds expressed as a decimal integer.

Values 1 — 1200

flush — The RIP flush timer value in seconds expressed as a decimal integer.

Values 1 — 1200

unicast-address

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

Context config>router>ripng>group *group-name*>neighbor *ip-int-name*

Description This command configures the unicast IPv6 address, RIPng updates messages will be sent to if the RIPng **send** command is set to **send unicast**.

Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address.

The **no** form of the command deletes the specified IPv6 unicast address from the configuration.

Default *ipv6-address* — IPv6 unicast address to which unicast RIPng updates should be sent.

Show Commands

database

Syntax `database [ip-prefix [/mask] [longer] [peer ip-address]`

Context `show>router>rip`
`show>router>ripng`

Description This command displays the routes in the RIP database.

Output **RIP Database Output** — The following table describes the RIP route 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. Yes — The route is valid.

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	0x0000	163	No
180.0.0.10/32	180.1.8.14	0.0.0.0	2	0x0000	179	No
180.0.0.14/32	180.1.8.14	0.0.0.0	1	0x0000	179	Yes
180.0.6.0/24	180.1.7.15	0.0.0.0	11	0x2002	163	No
180.0.6.0/24	180.1.8.14	0.0.0.0	11	0x2002	179	No
180.0.7.0/24	180.1.7.15	0.0.0.0	11	0x2002	163	No
180.1.5.0/24	180.1.7.15	0.0.0.0	2	0x0000	151	Yes
180.1.5.0/24	180.1.8.14	0.0.0.0	1	0x0000	167	No
180.100.17.16/30	180.1.7.15	0.0.0.0	2	0x0000	151	No
180.100.17.16/30	180.1.8.14	0.0.0.0	2	0x0000	167	No

Show Commands

```
-----  
No. of Routes: 10  
=====
```

```
A:ALA-A#
```

group

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

Context show>router>rip
show>router>ripng

Description Display RIP group information.

Parameters *group-name* — Displays RIP group information for the specified group.
detail — Displays detailed RIP group information.

Output **Standard RIP Group Output** — The following table describes the standard command output fields for a RIP group.

Label	Description
Group	The RIP group name.
Adm	Down — The RIP group is administratively down. Up — The RIP group is administratively up.
Opr	Down — The RIP group is operationally down. Up — The RIP group is operationally up.
Send Mode	Bcast — Specifies that RIPv2 formatted messages are sent to the broadcast address. Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address. None — Specifies that no RIP messages are sent (i.e., silent listener) 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. None — Specifies that RIP updates will not be accepted. RIPv1 — Specifies that RIP updates in version 1 format only will be accepted. RIPv2 — Specifies that RIP updates in version 2 format only will be accepted.

Label	Description (Continued)
Metric In	The metric value added to routes received from a RIP neighbor.

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#
```

neighbors

Syntax	neighbors [<i>ip-addr</i> <i>ip-int-name</i>] [advertised-routes detail]
Context	show>router>rip show>router>ripng
Description	Displays RIP neighbor interface information.
Parameters	<i>ip-addr</i> <i>ip-int-name</i> — Displays information for the specified IP interface.
Default	all neighbor interfaces

advertised-routes — 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.

Default **display RIP information**

Output **Standard RIP Neighbor Output** — The following table describes the standard command output fields for a RIP group.

Table 5: RIP Neighbor Standard Output Fields

Label	Description
Neighbor	The RIP neighbor interface name.
Adm	Down — The RIP neighbor interface is administratively down. Up — The RIP neighbor interface is administratively up.
Opr	Down — The RIP neighbor interface is operationally down. 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. Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address. None — Specifies that no RIP messages are sent (i.e., silent listener). 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. None — Specifies that RIP updates will not be accepted. RIPv1 — Specifies that RIP updates in version 1 format only are accepted. RIPv2 — Specifies that RIP updates in version 2 format only are accepted.
Metric In	The metric added to routes received from a RIP neighbor.

Sample Output

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

```

Output Detailed RIP Neighbor Output — The following table describes the standard command output fields for a RIP group.

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. Up — The RIP neighbor interface is administratively up.
Oper State	Down — The RIP neighbor interface is operationally down. Up — The RIP neighbor interface is operationally up.
Send Mode	BCast — Specifies that RIPv2 formatted messages are sent to the broadcast address. Mcast — Specifies that RIPv2 formatted messages are sent to the multicast address. None — Specifies that no RIP messages are sent (i.e., silent listener). 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. None — Specifies that RIP updates will not be accepted. RIPv1 — Specifies that RIP updates in version 1 format only will be accepted. 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.

Label	Description (Continued)
Split Horizon	Disabled — Split horizon disabled for the neighbor. Enabled — Split horizon and poison reverse enabled for the neighbor.
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. 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.

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
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      0x2002   n/a
180.0.0.5/32         180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.8/32         180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.9/32         180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.10/32        180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.11/32        180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.12/32        180.1.8.12        0.0.0.0           1       0x0000   n/a
180.0.0.13/32        180.1.8.12        0.0.0.0           10      0x2002   n/a
180.0.0.14/32        180.1.8.12        0.0.0.0           16      0x0000   n/a
180.0.0.15/32        180.1.8.12        0.0.0.0           2       0x0000   n/a
180.0.0.16/32        180.1.8.12        0.0.0.0           3       0x0000   n/a
-----
No. of Advertised Routes: 11
=====
A:ALA-A#
```

peer**Syntax** `peer [ip-int-name]`**Context** show>router>rip
show>router>ripng**Description** Displays RIP peer information.**Parameters** *ip-int-name* — Displays peer information for peers on the specified IP interface.**Default** `display peers for all interfaces`**Output** **RIP Peer Output** — The following table describes the command output fields for a RIP peer.

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.
No. of Peers	The number of RIP peers.

Sample Output

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

Show Commands

```
=====
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
show>router>ripng

Description Display interface level statistics for the RIP protocol

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 IP interface.

Output **RIP Statistics Output** — The following table describes the output fields for RIP statistics.

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

Label	Description (Continued)
Triggered Updates	Last 5 Min — The number of RIP updates that were sent in the last 5 minutes.
	Last 1 Min — The number of RIP updates that were sent in the last 1 minute.
	Total — The total number of triggered updates sent. These updates are sent before the entire RIP routing table is sent.
	Last 5 Min — The number of triggered updates that were sent in the last 5 minutes.
Bad Packets Received	Last 1 Min — The number of triggered updates that were sent in the last 1 minute.
	Total — The total number of RIP updates received on this interface that were discarded as invalid.
	Last 5 Min — The number of RIP updates received on this interface that were discarded as invalid in the last 5 minutes.
	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 RIPv1 updates received.
	Last 5 Min — The number of RIPv1 updates received in the last 5 minutes.
	Last 1 Min — The number of RIPv1 updates received in the last 1 minute.
	Total — The total number of RIPv1 updates ignored.
RIPv1 Updates Ignored	Last 5 Min — The number of RIPv1 updates ignored in the last 5 minutes.
	Last 1 Min — The number of RIPv1 updates ignored in the last 1 minute.
	Total — The total number of bad routes received from the peer.
	Last 5 Min — The number of bad routes received from the peer in the last 5 minutes.
RIPv1 Bad Routes	Last 1 Min — The number of bad routes received from the peer in the last minute.
	Total — The total number of times the router received RIPv1 route requests from other routers.
RIPv1 Requests Received	

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

Label	Description (Continued)
	Last 5 Min — The number of times the router ignored RIPv2 route requests from other routers in the last 5 minutes.
	Last 1 Min — The number of times the router ignored RIPv2 route requests from other routers in the last minute.
Authentication Errors	Total — The total number of authentication errors to secure table updates.
	Last 5 Min — The number of authentication errors to secure table updates in the last 5 minutes.
	Last 1 Min — The number of authentication errors to secure table updates in the last minute.

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#
```

Clear Commands

database

Syntax	database
Context	clear>router>rip show>router>ripng
Description	Flush all routes in the RIP database.

statistics

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

Debug RIP Commands

auth

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

error

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

events

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

holddown

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

Debug RIP Commands

packets

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

request

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

trigger

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

updates

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

In This Chapter

This chapter provides information about configuring the Open Shortest Path First (OSPF) protocol.

Topics in this chapter include:

- [Configuring OSPF on page 308](#)
 - [OSPF Areas on page 309](#)
 - [Backbone Area on page 309](#)
 - [Stub Area on page 310](#)
 - [Not-So-Stubby Area on page 311](#)
 - [OSPF Super Backbone on page 312](#)
 - [Virtual Links on page 318](#)
 - [Neighbors and Adjacencies on page 319](#)
 - [Link-State Advertisements on page 320](#)
 - [Metrics on page 320](#)
 - [Authentication on page 321](#)
 - [IP Subnets on page 322](#)
 - [Preconfiguration Recommendations on page 322](#)
- [Loop-Free Alternate Shortest Path First \(LFA SPF\) Policies on page 335](#)
- [OSPF Configuration Process Overview on page 344](#)
- [Configuration Notes on page 345](#)

Configuring OSPF

OSPF (Open Shortest Path First) 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. Alcatel-Lucent's implementation of OSPF conforms to OSPF Version 2 specifications presented in RFC 2328, *OSPF Version 2*. Routers running OSPF can be enabled with minimal configuration. All default and command parameters can be modified.

Key OSPF features are:

- Backbone areas
- Stub areas
- Not-So-Stubby areas (NSSAs)
- Virtual links
- Authentication
- Route redistribution
- Routing interface parameters
- OSPF-TE extensions (Alcatel-Lucent's implementation allows MPLS fast reroute)

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.

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 10](#)) 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 10: Backbone Area

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 10](#), 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.

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 10](#), area 0.0.0.3 could be configured as a NSSA area.

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

VPRN 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.

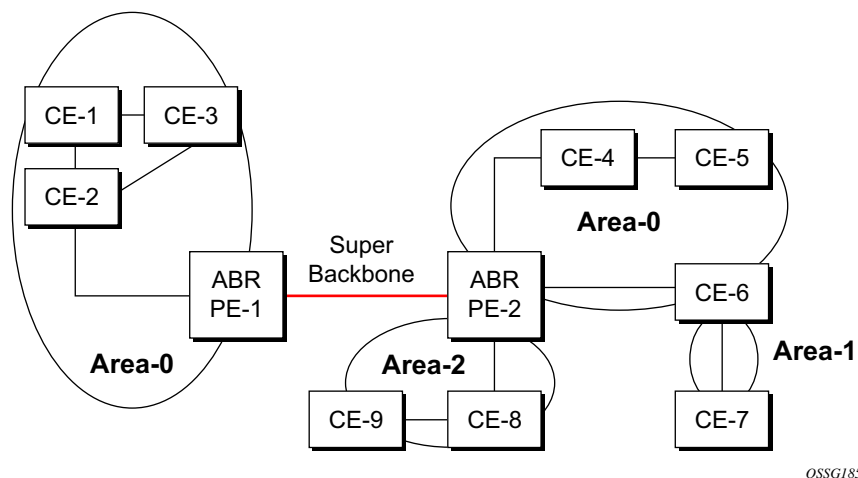


Figure 11: PEs Connected to an MPLS-VPN Super Backbone

In [Figure 11](#), 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.

Sham Links




Figure 12: Sham Links

Figure 12 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-VPRN. 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 shamlink forms an adjacency over the MPLS-VPRN 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.

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-VPN 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.

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.

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.

OSPF Areas

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.

Sham Links

A sham link is only required if a backdoor link (shown as the red link in [Figure 12](#)) is present, otherwise configuring an OSPF super backbone will probably suffice.

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 the SR-OS platforms can support a fully non-stop routing model for control plane high availability the SR-OS node has no need for graceful restart as defined by the IETF in various RFCs for each routing protocol. However, since the 7x50 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 the SROS based system to provide other routers which have requested it, a grace period, 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.

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 10](#) 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.

Neighbors and Adjacencies

A router uses the OSPF Hello protocol to discover neighbors. A neighbor is a router configured with an interface to a common network. 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 media 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.

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.

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.

Authentication

All OSPF protocol exchanges can be authenticated. This means that only trusted routers can participate in autonomous system routing. Alcatel-Lucent'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. Alcatel-Lucent'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.

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).

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.

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.

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.

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.

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 the SROS 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.

For example,

```
config
  router
    ospf3 64 10.20.1.3
```

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 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”.

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:

```
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 IOM a LFA next-hop, when found by SPF, in addition to the primary next-hop for each prefix in the FIB.

```
config>router>ip-fast-reroute
```

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.

```
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:

```
config>router>isis>interface> loopfree-alternate-exclude
```

```
config>router>ospf>area>interface> loopfree-alternate-exclude
```

Note that 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:

```
config>service>vprn>ospf>area>loopfree-alternate-exclude
```

```
config>service>vprn>ospf>area>interface>loopfree-alternate-exclude
```

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.

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 on page 328](#).


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 thus remain up when the FIB switches from the primary IGP next-hop to the LFA IGP next-hop.

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



OSSG712

Figure 13: Example Topology with Primary and LFA Routes

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.

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

Consider now the case where the primary next-hop uses a broadcast interface as illustrated in [Figure 14](#).

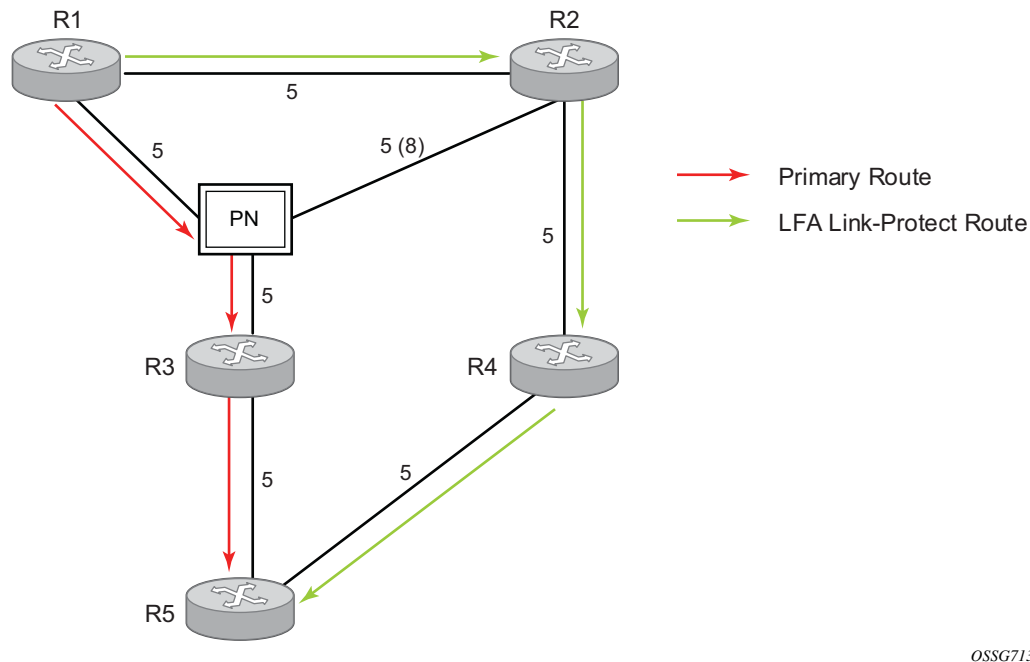


Figure 14: Example Topology with Broadcast Interfaces

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):
 $\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, R1) + \text{Distance_opt}(R1, R5)$
 and,
 $\text{Distance_opt}(R2, R5) \geq \text{Distance_opt}(R2, R3) + \text{Distance_opt}(R3, R5)$
- **Rule 2:** Node-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):
 $\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, R1) + \text{Distance_opt}(R1, R5)$
 and,
 $\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, R3) + \text{Distance_opt}(R3, R5)$
- **Rule 3:** Link-protect LFA backup next-hop (primary next-hop R1-R3 is a broadcast interface):
 $\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, R1) + \text{Distance_opt}(R1, R5)$
 and,
 $\text{Distance_opt}(R2, R5) < \text{Distance_opt}(R2, \text{PN}) + \text{Distance_opt}(\text{PN}, R5)$
 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:

- A) It will pick the node-protect type in favor of the link-protect type.
- B) If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
- C) 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:

- A) 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.
- B) 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.

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

- C) If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
- D) 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.
- E) If no LFA next-hop results from Step D, SPF will rerun Steps B-D using the second set.

Note 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.

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 the selection algorithm in Section 1.3 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.

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.

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.

Configuration of 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:

```
configure>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.

Once 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.

Configuring Affinity or Admin Group Constraint in Route Next-Hop Policy

Configuration of Route Next-Hop Policy Template

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:

```
config>router>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.

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

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

```
config>service>vprn>interface>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 that 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:

```
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. The modified LFA next-hop selection algorithm is shown in Section 7.5.

Configuring SRLG Group Constraint in Route Next-Hop Policy

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:

```
configure>router>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.

```
config>router>interface>if-attribute>srlg-group group-name [group-name...(up to 5 max)]
```

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

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

Configuration of Route Next-Hop Policy Template

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 that 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:

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. The modified LFA next-hop selection algorithm is shown in Section 7.5.

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 new IP admin groups and SRLGs be compatible with the ones already supported in MPLS. The following rules are implemented:

- The definition of admin groups and SRLGs are moved under the new ‘config>router>if-attribute’ context. When upgrading customers to R12, all user configured admin groups and SRLGs under ‘config>router>mpls’ context will automatically be moved into the new context. The configuration of admin groups and SRLGs under the ‘config>router>mpls’ context in CLI is deprecated.
- The binding of an MPLS interface to a group, i.e., configuring membership of an MPLS interface in a group, continues to be performed under ‘config>router>mpls>interface’ context.
- The binding of a local or remote MPLS interface to an SRLG in the SRLG database continues to be performed under the ‘config>router>mpls>srlg-database’ context.
- The binding of an ISIS/OSPF interface to a group is performed in the ‘config>router>interface>if-attribute’ or ‘config>service>vpn>if>if-attribute’ or ‘config>service>ies>if>if-attribute’ contexts. This is used by ISIS 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.

Configuring Protection Type and Next-Hop Type Preference in Route next-hop policy template

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 SROS 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 SROS 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:

configure router route-nh-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.

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 is achieves that:

config>router>isis>interface>lfa-policy-map route-nh-template *template-name*

config>router>ospf(3)>area>interface>lfa-policy-map route-nh-template *template-name*

config>service>vprn>ospf(3)>area>interface>lfa-policy-map route-nh-template *template-name*

Excluding Prefixes from LFA SPF

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.

Excluding Prefixes from LFA SPF

In the current SROS 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:

```
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 SROS implementation:

```
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
default-action reject

exit

If the user enabled the R12 IS-IS prefix prioritization based on tag, it will also apply to SPF LFA. However, 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 that 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.

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.

Note that 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 instance, 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 ‘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. Note however the admin-group constraints will 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. The changes are highlighted in yellow color.

- 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.

Modification to LFA Next-Hop Selection Algorithm

- 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**.

OSPF 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 (e.g. 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.

OSPF Configuration Process Overview

Figure 15 displays the process to provision basic OSPF parameters.

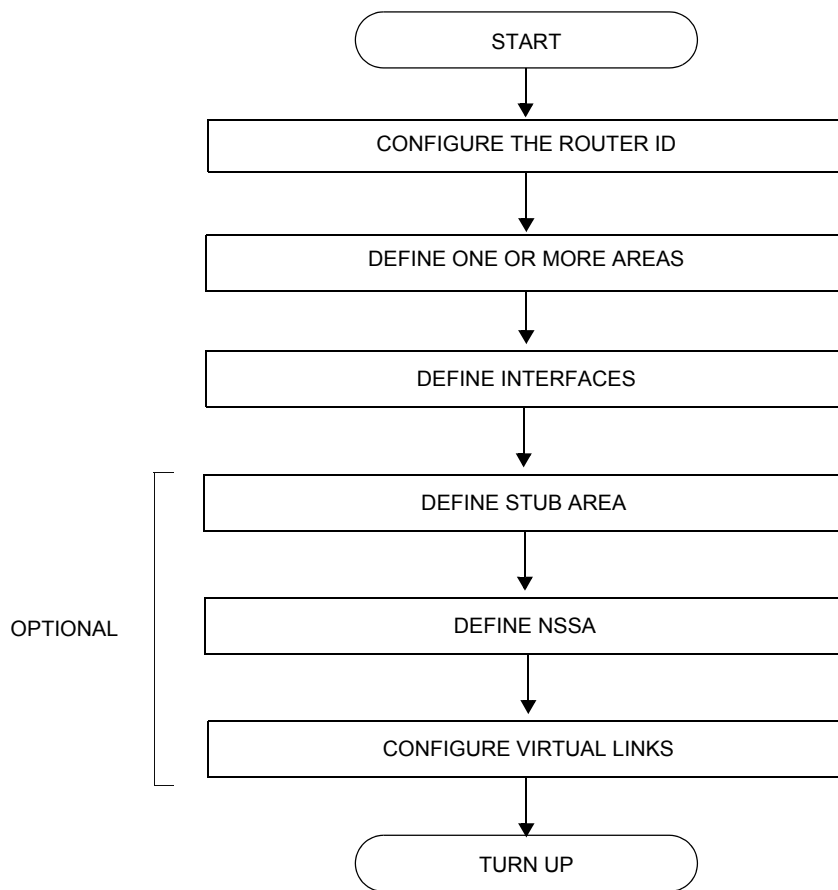


Figure 15: OSPF Configuration and Implementation Flow

Configuration Notes

This section describes OSPF configuration caveats.

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.
-

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.

Configuring OSPF with CLI

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

Topics in this section include:

- [OSPF Configuration Guidelines on page 348](#)
- [Basic OSPF Configuration on page 349](#)
- [Configuring the Router ID on page 350](#)
- [Configuring OSPF Components on page 351](#)
 - [Configuring the Router ID on page 350](#)
 - [Configuring an OSPF Area on page 352](#)
 - [Configuring a Stub Area on page 353](#)
 - [Configuring a Not-So-Stubby Area on page 354](#)
 - [Configuring a Virtual Link on page 355](#)
 - [Configuring an Interface on page 357](#)
 - [Configuring Authentication on page 359](#)
 - [Assigning a Designated Router on page 363](#)
 - [Configuring Route Summaries on page 365](#)
 - [Configuring Route Preferences on page 367](#)
- [OSPF Configuration Management Tasks on page 369](#)
 - [Modifying a Router ID on page 369](#)
 - [Deleting a Router ID on page 371](#)
 - [Modifying OSPF Parameters on page 372](#)

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.

Basic OSPF Configuration

This section provides information to configure OSPF 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"`).

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#
```

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 node 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#
```

Configuring OSPF Components

Use the CLI syntax displayed below for:

- [Configuring OSPF Parameters on page 351](#)
 - [Configuring a Stub Area on page 353](#)
 - [Configuring a Not-So-Stubby Area on page 354](#)
 - [Configuring a Virtual Link on page 355](#)
 - [Configuring an Interface on page 357](#)
 - [Configuring Authentication on page 359](#)
 - [Assigning a Designated Router on page 363](#)
 - [Configuring Route Summaries on page 365](#)
-

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"
graceful-restart
    helper-disable
exit
-----
A:ALA-49>config>router>ospf# ex
```

Configuring an OSPF 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`

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:
```

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.

Use the following CLI syntax to configure virtual links:

CLI Syntax: ospf

```

    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#
  
```

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`

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"
graceful-restart
    helper-disable
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#
```

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"
graceful-restart
  helper-disable
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
```

Configuring a Virtual Link

```
        nssa
        exit
    exit
    area 1.2.3.4
    exit
-----
A:ALA-49>config>router>ospf#
```


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 RFC5185. 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}
        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
```

Configuring an Interface

```
traffic-engineering
export "OSPF-Export"
graceful-restart
  helper-disable
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
```

Configuring Authentication

This section includes the following topics:

- [Overview on page 359](#)
- [Configuring Authentication Keys and Algorithms on page 359](#)
- [Configuring Authentication using Keychains on page 362](#)

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.

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 and 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 on page 357](#) and [Configuring a Virtual Link on page 355](#).

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"
graceful-restart
    helper-disable
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"
    graceful-restart
        helper-disable
    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#

```

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.

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"
graceful-restart
  helper-disable
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"
```

Assigning a Designated Router

```
        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#
```


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 on page 353](#) and [Configuring a Not-So-Stubby Area on page 354](#).

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"
graceful-restart
    helper-disable
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
```

Configuring Route Summaries

```
        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#
```

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 6](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

Table 6: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ^a
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
TMS	167	No
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. 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 on page 351](#).

Use the following CLI syntax to configure a route preference:

CLI Syntax: *ospf ospf-instance*
 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"
graceful-restart
    helper-disable
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
-----
```

OSPF Configuration Management Tasks

This section discusses the following OSPF configuration management tasks:

- [Modifying a Router ID on page 369](#)
- [Deleting a Router ID on page 371](#)
- [Modifying OSPF Parameters on page 372](#)

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
```

Modifying a Router ID

```
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#
```

Deleting a Router ID

You can modify a router ID, but you cannot delete the parameter. When the `no 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.

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"
graceful-restart
  helper-disable
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#
```

OSPF Command Reference

Command Hierarchies

- [Configuration Commands on page 375](#)
- [Show Commands on page 378](#)
- [Clear Commands on page 378](#)
- [Debug Commands on page 378](#)

Configuration Commands

```

config
— router
  — [no] ospf [ospf-instance] [router-id]
  — ospf3 [instance-id] [router-id]
  — [no] ospf3 instance-id
    — advertise-router-capability {link | area | as}
    — no advertise-router-capability
    — advertise-tunnel-link {link | area | as}
    — no advertise-tunnel-link
    — [no] area area-id
      — area-range ip-prefix/mask [advertise | not-advertise]
      — no area-range ip-prefix/mask
      — [no] blackhole-aggregate
      — [no] interface ip-int-name [secondary]
        — [no] advertise-subnet
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — authentication-type {password | message-digest}
        — no authentication-type
        — auth-keychain name
        — no auth-keychain
        — bfd-enable
        — no bfd-enable
        — dead-interval seconds
        — no dead-interval
        — hello-interval seconds
        — no hello-interval
        — interface-type {broadcast | point-to-point}
        — no interface-type
        — [no] loopfree-alternate-exclude
        — lfa-policy-map route-nh-template template-name
        — no lfa-policy-map
        — 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*] [**hash** | **hash2**]
- **no message-digest-key** *key-id*
- **metric** *metric*
- **no metric**
- **mtu** *bytes*
- **no mtu**
- [no] **passive**
- **priority** *number*
- **no priority**
- **retransmit-interval** *seconds*
- **no retransmit-interval**
- [no] **shutdown**
- **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*
 - **area-range** *ip-prefix/prefix-length* [**advertise** | **not-advertise**]
 - **no area-range** *ip-prefix/prefix-length*
 - **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*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - **no authentication-key**
 - **authentication-type** {**password** | **message-digest**}
 - **no authentication-type**
 - **auth-keychain** *name*
 - **no auth-keychain**
 - **dead-interval** *seconds*
 - **no dead-interval**
 - **hello-interval** *seconds*
 - **no hello-interval**
 - **message-digest-key** *key-id* **md5** [*key* | *hash-key*] [**hash** | **hash2**]
 - **no message-digest-key** *key-id*
 - **retransmit-interval** *seconds*
 - **no retransmit-interval**
 - [no] **shutdown**
 - **transit-delay** *seconds*
 - **no transit-delay**
- [no] **asbr** [**trace-path** *domain-id*]
- [no] **compatible-rfc1583**
- [no] **disable-ldp-sync**
- **export** *policy-name* [*policy-name...*(up to 5 max)]
- **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] ldp-over-rsvp**
- **[no] loopfree-alternate**
- **loopfree-alternate-exclude** **prefix-policy** *prefix-policy* [*prefix-policy...* up to 5]
- **no loopfree-alternate-exclude**
- **[no] multicast-import**
- **[no] opaque-lsa**
- **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**
- **router-id** *ip-address*
- **no router-id**
- **[no] rsvp-shortcut**
- **[no] shutdown**
- **timers**
 - **[no] incremental-spf-wait** *inc-spf-wait*
 - **[no] lsa-accumulate** *lsa-accum-time*
 - **[no] lsa-arrival** *lsa-arrival-time*
 - **[no] lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
 - **[no] redistribute-delay** *redist-wait*
 - **[no] spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
- **[no] traffic-engineering**
- **[no] unicast-import-disable**

Show Commands

```
show
  — router
    — ospf [ospf-instance]
    — ospf3 [ospf-instance]
      — 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]
      — interface [area area-id] [detail]
      — interface [ip-int-name | ip-address] [detail]
      — lfa-coverage
      — neighbor [remote ip-address] [detail]
      — neighbor [ip-int-name] [router-id] [detail]
      — opaque-database [link link-id | area area-id | as] [adv-router router-id] [ls-id] [detail]
      — range [area-id]
      — routes [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary] [exclude-short-cut]
      — spf [lfa]
      — statistics
      — status
      — virtual-link [detail]
      — virtual-neighbor [remote ip-address] [detail]
```

Clear Commands

```
clear
  — router
    — ospf [ospf-instance]
    — ospf3 [ospf-instance]
      — database [purge]
      — export
      — neighbor [ip-int-name | ip-address]
      — statistics
```

Debug Commands

```
debug
  — router
    — ospf [ospf-instance]
    — ospf3 [ospf-instance]
      — area [area-id]
      — no area
      — area-range [ip-address]
      — no area-range
      — cspf [ip-addr]
      — 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* | *router-id*]
- **no neighbor**
- **nssa-range** [*ip-address*]
- **no nssa-range**
- **packet** [*packet-type*] [*ip-address*]
- **no packet**
- **rtm** [*ip-addr*]
- **no rtm**
- **spf** [*type*] [*dest-addr*]
- **no spf**
- **virtual-neighbor** [*ip-address*]
- **no virtual-neighbor**

Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>router>ospf
config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description The **shutdown** 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 **no shutdown** command.

The **shutdown** 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default **OSPF Protocol** — The Open Shortest Path First (OSPF) protocol is created in the **no shutdown** state.

OSPF Interface — When an IP interface is configured as an OSPF interface, OSPF on the interface is in the **no shutdown** state by default.

OSPF Global Commands

ospf

Syntax `[no] ospf ospf-instance [instance-id] [router-id]`

Context config>router

Description This command creates an OSPF routing instance and then enters the associated context to configure the associated protocol parameters.

Additionally, the router ID can be specified as another parameter of the OSPF command. This parameter is required for all non-base OSPF instances.

The default value for the base instance is inherited from the configuration in the config>router context. When that is not configured the following applies:

1. The system uses the system interface address (which is also the loopback address).
2. If a system interface address is not configured, use 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 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 avoid 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.

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.

Issue the shutdown and no shutdown commands for the instance for the new router ID to be used, or reboot the entire router.

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

Default `no ospf`

Parameters *instance-id* — 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 — 31

router-id — Specifies the OSPF router ID to be used with the associated OSPF instance. The router-id must be given a dot decimal notation format.


Values 1 — 31

ospf3

Syntax	ospf3 [<i>instance-id</i>] [<i>router-id</i>] [no] ospf3 <i>instance-id</i>
Context	config>router
Description	<p>This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.</p> <p>When an OSPFv3 instance is created, the protocol is enabled. To start or suspend execution of the OSPF.</p> <p>The no form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.</p>
Default	no default
Parameters	<p><i>instance-id</i> — Specify 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.</p> <p>Values 0 — 31: IPV6 unicast</p> <p>Values 64—95: IPV4 unicast</p> <p><i>router-id</i> — Specifies the OSPF router ID to be used with the associated OSPF instance. The router-id must be given a dot decimal notation format.</p>

asbr

Syntax	[no] asbr [<i>trace-path domain-id</i>]
Context	config>router>ospf
Description	<p>This command configures the router as a 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.</p> <p>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.</p> <p>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 avoid 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>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</p>



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.

Default **no asbr** — The router is not an ASBR.

Parameters *domain-id* — Specifies the domain ID.

Values 1 — 31

Default 0

compatible-rfc1583

Syntax [no] compatible-rfc1583

Context config>router>ospf

Description This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.

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.

Although it would be favorable to require all routers to run a more current compliancy level, this command allows the router to use obsolete methods of calculation.

The **no** form of the command enables the post-RFC1583 method of summary and external route calculation.

Default **compatible-rfc1583** — RFC1583 compliance is enabled.

disable-ldp-sync

Syntax	[no] disable-ldp-sync
Context	config>router>ospf
Description	<p>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 advertized cost is different. It will then disable IGP-LDP synchronornization 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.</p> <p>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.</p>
Default	no disable-ldp-sync

export

Syntax	export <i>policy-name</i> [<i>policy-name</i>...] no export
Context	config>router>ospf
Description	<p>This command associates export route policies to determine which routes are exported from the route table to OSPF. Export polices 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 no form of the command removes all policies from the configuration.</p>
Default	no export — No export route policies specified.
Parameters	<p><i>policy-name</i> — The export 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, etc.), the entire string must be enclosed within double quotes.</p> <p>The specified name(s) must already be defined.</p>

export-limit

Syntax	export-limit <i>number</i> [<i>log percentage</i>] no export-limit
Context	config>router>ospf config>router>ospf3 config>service>vpn>ospf config>service>vpn>ospf3
Description	<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 no form of the command removes the parameters from the configuration.</p>
Default	no export-limit, the export limit for routes or prefixes is disabled.
Parameters	<p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into OSPF from the route table.</p> <p>Values 1 — 4294967295</p> <p><i>log percentage</i> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p>Values 1 — 100</p>

external-db-overflow

Syntax	external-db-overflow <i>limit interval</i> no external-db-overflow
Context	config>router>ospf
Description	<p>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.</p> <p>The <i>limit</i> 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 <i>limit</i>, 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.</p> <p>The <i>interval</i> 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.</p> <p>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.</p> <p>The no form of the command disables limiting the number of non-default AS-external-LSA entries.</p>

Default **no external-db-overflow** — No limit on non-default AS-external-LSA entries.

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 0 — 2147483674

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 — 2147483674

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>router>ospf

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 7, “Route Preference Defaults by Route Type,” on page 388. 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** — OSPF external routes have a default preference of 150.

OSPF Global Commands

Parameters *preference* — The preference for external routes expressed as a decimal integer. Defaults for different route types are listed in Table 7.

Table 7: 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
TMS	167	No
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

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

Values 1 — 255

graceful-restart

Syntax **[no] graceful-restart**

Context config>router>ospf
config>router>ospf3

Description This command enables graceful-restart for OSPF. When the control plane of a GR-capable router fails, the neighboring routers (GR helpers) temporarily preserve adjacency information, so packets continue to be forwarded through the failed GR router using the last known routes. If the control plane of the GR router comes back up within the GR timer, then the routing protocols would re-converge to minimize service interruption.

The **no** form of the command disables graceful restart and removes all graceful restart configurations in 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	<p>This command disables the helper support for graceful restart.</p> <p>When graceful-restart is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the router will not help the neighbors to restart).</p> <p>The no helper-disable command enables helper support and is the default when graceful-restart is enabled.</p>
Default	disabled

ldp-over-rsvp

Syntax	[no] ldp-over-rsvp
Context	config>router>ospf
Description	This command allows LDP-over-RSVP processing in this OSPF instance.

loopfree-alternate

Syntax	[no] loopfree-alternate
Context	config>router>ospf>area config>router>ospf3>area
Description	<p>This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level.</p> <p>When this command is enabled, it instructs 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 routing table along with the primary next-hop for the prefix.</p> <p>The no form of this command disables the LFA computation by IGP SPF.</p>
Default	no loopfree-alternate

lfa-policy-map

Syntax	lfa-policy-map route-nh-template <i>template-name</i> no lfa-policy-map
---------------	--

OSPF Global Commands

Context	config>router>ospf>area>interface config>router>ospf3>area>interface
Description	<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 loopfree-alternate-exclude, 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 no form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.</p>
Parameters	<i>template-name</i> — Specifies the name of the template, up to 32 characters.

loopfree-alternate-exclude

Syntax	loopfree-alternate-exclude prefix-policy <i>prefix-policy</i> [<i>prefix-policy</i>... up to 5] no loopfree-alternate-exclude
Context	config>router>ospf config>router>ospf3
Description	<p>This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.</p> <p>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.</p> <p>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 that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.</p> <p>The default action of the 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 that did not match any entry in the policy will be accepted into LFA SPF.</p> <p>The no form deletes the exclude prefix policy.</p>
Parameters	prefix-policy <i>prefix-policy</i> — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

multicast-import

Syntax	[no] multicast-import
Context	config>router>ospf
Description	This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF. The no form of the command disables the submission of routes into the multicast RTM.
Default	no multicast-import

opaque-lsa

Syntax	[no] opaque-lsa
Context	config>router>ospf
Description	This command enables the router's support for opaque LSA types. The no form of the command disables the support.

overload

Syntax	overload [timeout seconds] no overload
Context	config>router>ospf config>router>ospf3
Description	<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 no overload command is executed.</p> <p>If the overload command is encountered during the execution of an overload-on-boot 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 overload-on-boot command is saved after the overload command. However, when overload-on-boot is configured under OSPF with no timeout value configured, the router will remain in overload state indefinitely after a reboot.</p> <p>Use the no form of this command to return to the default. When the no overload command is executed, the overload state is terminated regardless of the reason the protocol entered overload state.</p>
Default	no overload

OSPF Global Commands

Parameters **timeout** *seconds* — Specifies the number of seconds to reset overloading.

Values 1 — 1800

Default 60

overload-include-ext-2

Syntax **[no] overload-include-ext-2**

Context config>router>ospf

Description 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.

Default **no overload-include-ext-2**

overload-include-stub

Syntax **[no] overload-include-stub**

Context config>router>ospf

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

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.

Parameters **timeout** *seconds* — Specifies the number of seconds to reset overloading.

Values 1 — 1800

Default indefinitely in overload.

preference

Syntax **preference** *preference*
no preference

Context config>router>ospf

This command configures the preference for OSPF internal 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 [Table 8](#). 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 **preference 10** — OSPF internal routes have a preference of 10.

Parameters *preference* — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in Table 8.

Table 8: 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
TMS	167	No
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

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

Values 1 — 255

reference-bandwidth

Syntax **reference-bandwidth** *bandwidth-in-kbps*
reference-bandwidth [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]
no reference-bandwidth

Context config>router>ospf

Description This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

$$\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$$

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100

- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric *metric*** command in the **config>router>ospf>area>interface *ip-int-name*** context.

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

Default **reference-bandwidth 100000000** — Reference bandwidth of 100 Gbps.

Parameters *bandwidth-in-kbps* — The reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 1000000000

tbps *Tera-bps* — The reference bandwidth in terabits per second expressed as a decimal integer.

Values 1 — 4

gbps *Giga-bps* — The reference bandwidth in gigabits per second expressed as a decimal integer.

Values 1 — 999

mbps *Mega-bps* — The reference bandwidth in megabits per second expressed as a decimal integer.

Values 1 — 999

kbps *Kilo-bps* — reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 999

router-id

Syntax **router-id *ip-address***
no router-id

Context config>router>ospf

Description This command configures the router ID for the OSPF instance. This command configures the router ID for the OSPF instance.

When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the **config>router** context.

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **tconfig>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, use 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.

OSPF Global Commands

It is possible to configure an SR OS node 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.

advertise-router-capability

Syntax	advertise-router-capability {link area as} no advertise-router-capability
Context	config>router>ospf config>router>ospf3
Description	<p>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:</p> <ul style="list-style-type: none">• 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 <p>The parameters (link, area and as) control the scope of the capability advertisements..</p> <p>The no form of this command, disables this capability.</p>
Default	no advertise-router-capability
Parameters	<p>link — are only advertised over local link and not flooded beyond</p> <p>area — are only advertised within the area of origin.</p> <p>as — are only advertised throughout the entire autonomous system</p>

rsvp-shortcut

Syntax	[no] rsvp-shortcut
Context	config>router>ospf
Description	<p>This command enables the use of an RSVP-TE shortcut for resolving IGP routes by IS-IS or OSPF routing protocols.</p> <p>This command instructs IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. If the user enabled the relative-metric option for this LSP, IGP will apply the shortest IGP cost</p>

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 which 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 LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When `rsvp-shortcut` is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in `configure>router>mpls>lsp>to`, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP 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 OSPF or IS-IS will only use 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 the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

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 next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
 - select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
 - exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
 - prefer tunnel next-hop over ip next-hop
 - within tunnel next-hops:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index
 - within ip next-hops:
 - i. select lowest downstream router-id
 - ii. if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies 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.

The ingress IOM will spray 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 `rsvp-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.

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.

The `no` form of this command disables the resolution of IGP routes using RSVP shortcuts.

OSPF Global Commands

Default **no rsvp-shortcut**

advertise-tunnel-link

Syntax **advertise-tunnel-link {link | area | as}**
no advertise-tunnel-link

Context config>router>ospf
config>router>ospf3

Description This command enables the advertisement of router as defined in the IETF RFC 4970. This adds a new TLV based mechanism, allowing OSPF (OSPFv2 and OSPFv3) router to advertise specific including Traffic Engineering capability, graceful restart helper support and stub router support.

The parameters (**link**, **area**, and **as**) control the scope of the capabilities advertisements.

The no form of this command, disables this capability.

Default **no advertise-tunnel-link**

Parameters **link** — are only advertised over local link and not flooded beyond.

area — are only advertised within the area of origin.

as — are only advertised throughout the entire autonomous system.

super-backbone

Syntax **[no] super-backbone**

Context config>service>vpn>ospf

Description This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Refer to the OS Services Guide for syntax and command usage information.

The **no** form of the command disables the super-backbone functionality.

Default **no super-backbone**

timers

Syntax	timers
Context	config>router>ospf
Description	<p>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.</p> <p>Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.</p>
Default	none

incremental-spf-wait

Syntax	incremental-spf-wait <i>inc-spf-wait</i> no incremental-spf-wait
Context	config>router>ospf>timers
Description	<p>This command sets the internal OSPF delay before an incremental SPF calculation is performed.</p> <p>The no incremental-spf-wait form of the command resets the timer value back to the default value.</p>
Default	1000ms (1 second)
Parameters	<i>inc-spf-wait</i> — Specifies the OSPF incremental SPF recalculation delay.
Values	0 — 10000

lsa-accumulate

Syntax	lsa-accumulate <i>lsa-accum-time</i> no lsa-accumulate
Context	config>router>ospf>timers
Description	<p>This commands sets the internal OSPF delay to allow for the accumulation of multiple LSA so OSPF messages can be sent as efficiently as possible.</p> <p>Shorting this delay can speed up the advertisement of LSAs to OSPF neighbors but may increase the number of OSPF messages sent.</p>
Default	1000ms (1 second)
Parameters	<i>lsa-accum-time</i> — Specifies the LSA acculation delay in milliseconds.
Values	0 — 10000

lsa-arrival

Syntax	lsa-arrival <i>lsa-arrival-time</i> no lsa-arrival
Context	config>router>ospf>timers
Description	<p>This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.</p> <p>It is recommended that the neighbors configured (lsa-generate) <i>lsa-second-wait</i> interval is equal or greater than the lsa-arrival timer configured here.</p> <p>Use the no form of this command to return to the default.</p>
Default	no lsa-arrival
Parameters	<p><i>lsa-arrival-time</i> — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected.</p> <p>Values 0 — 600000</p>

lsa-generate

Syntax	lsa-generate <i>max-lsa-wait</i> [<i>lsa-initial-wait</i> [<i>lsa-second-wait</i>]] no lsa-generate-interval
Context	config>router>ospf>timers
Description	<p>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 <i>lsa-second-wait</i> timer until a maximum value is reached.</p> <p>Configuring the lsa-arrival interval to equal or less than the <i>lsa-second-wait</i> interval configured in the lsa-generate command is recommended.</p> <p>Use the no form of this command to return to the default.</p>
Default	no lsa-generate
Parameters	<p><i>max-lsa-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated.</p> <p>Values 10 — 600,000</p> <p>Default 5,000 milliseconds</p> <p><i>lsa-initial-wait</i> — 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.</p>

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 — 600000

Default 5,000 milliseconds

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 — 600000

Default 5,000 milliseconds

redistribute-delay

Syntax **redistribute-delay** *redist-wait*
no redistribute-delay

Context config>router>ospf>timers

Description This command sets the internal OSPF hold down timer for external routes being redistributed into OSPF. 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. The **no redistribute-delay** form of the command resets the timer value back to the default value.

Default 1000ms (1 second)

Parameters *redist-wait* — Specifies the OSPF redistribution hold down timer for external routes being advertised into OSPF.

Values 0 — 1000

spf-wait

Syntax **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>router>ospf>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 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*.

OSPF Global Commands

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.

Default **no spf-wait**

Parameters *max-spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 10 — 120000

Default 10000

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 — 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 10 — 100000

Default 1000

traffic-engineering

Syntax **[no] traffic-engineering**

Context config>router>ospf

Description This command enables traffic engineering route calculations constrained by nodes or 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 traffic engineered route calculations.

Default **no traffic-engineering** — Traffic engineered route calculations is disabled.

unicast-import-disable

Syntax **[no] unicast-import-disable**

Context config>router>ospf

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 **disabled**

OSPF Area Commands

area

Syntax	[no] area <i>area-id</i>
Context	config>router>ospf
Description	<p>This command creates the context to configure an OSPF 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.</p> <p>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 etc., that are currently assigned to this area.</p>
Default	no area — No OSPF areas are defined.
Parameters	<p><i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.</p> <p>Values 0.0.0.0 — 255.255.255.255 (dotted decimal), 0 — 4294967295 (decimal integer)</p>

area-range

Syntax	area-range <i>ip-prefix/mask</i> [advertise not-advertise] no area-range <i>ip-prefix/mask</i>
Context	config>router>ospf>area config>router>ospf>area>nssa
Description	<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 no form of the command deletes the range (non) advertisement.</p>
Default	no area-range — No range of addresses are defined.
Special Cases	<p>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.</p> <p>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.</p>

OSPF Area Commands

Parameters	<i>ip-prefix</i> — 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) <i>mask</i> — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation. Values 0 — 32 (mask length), 0.0.0.0 — 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. The default is advertise .
-------------------	--

blackhole-aggregate

Syntax	[no] blackhole-aggregate
Context	config>router>ospf>area
Description	<p>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.</p> <p>It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.</p> <p>The no form of this command removes this option.</p>
Default	blackhole-aggregate

default-metric

Syntax	default-metric metric no default-metric
Context	config>router>ospf>area>stub
Description	<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 stub area.</p> <p>The no form of the command reverts to the default value.</p>
Default	default-metric 1

Parameters	<i>metric</i> — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.
Values	1 — 16777215

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate
Context	config>router>ospf>area config>router>ospf>area>interface
Description	<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 no form of this command re-instates the default value for this command.</p>
Default	no loopfree-alternate-exclude.

lsa-filter-out

Syntax	lsa-filter-out [all except-own-rtrlsa except-own-rtrlsa-and-defaults] no lsa-filter-out
Context	config>router>ospf>area>interface config>router>ospf3>area>interface config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface
Description	<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"> • 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 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 loop-back or system interface address (contained in the router-LSA) as next-hop. • 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. <p>The no form of this command disables OSPF LSA filtering (normal operation).</p>

OSPF Area Commands

Default no lsa-filter-out

nssa

Syntax [no] nssa

Context config>router>ospf>area

Description This command creates the context to configure an OSPF 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 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** — The OSPF area is not an NSSA.

originate-default-route

Syntax **originate-default-route [type-7] [no-adjacency-check]**
no originate-default-route

Context config>router>ospf>area>nssa

Description 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) or Autonomous System Border Router (ASBR).

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.

The **no** form of the command disables origination of a default route.

Default **no originate-default-route** — A default route is not originated.

Parameters **type-7** — Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default Type 3 LSA for the default route.

no-adjacency-check — Specifies whether or not adjacency checks shall be performed for the NSSA.

redistribute-external

Syntax	[no] redistribute-external
Context	config>router>ospf>area>nssa
Description	<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 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 domain.</p> <p>The no form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.</p>
Default	redistribute-external — External routes are redistributed into the NSSA.

stub

Syntax	[no] stub
Context	config>router>ospf>area
Description	<p>This command enables access to the context to configure an OSPF 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 stub command. An OSPF 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 no form of the command removes the stub designation and configuration context from the area.</p>
Default	no stub — The area is not configured as a stub area.

summaries

Syntax	[no] summaries
Context	config>router>ospf>area>stub config>router>ospf>area>nssa
Description	<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>

OSPF Area Commands

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default **summaries** — Summary routes are advertised by the ABR into the stub area or NSSA.

Interface/Virtual Link Commands

advertise-subnet

Syntax	[no] advertise-subnet
Context	config>router>ospf>area>interface <i>ip-int-name</i>
Description	<p>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.</p> <p>The no form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.</p>
Default	advertise-subnet — Advertises point-to-point interfaces as subnet routes.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>ospf>area>interface <i>ip-int-name</i> config>router>ospf>area>virtual-link > <i>if</i> >
Description	<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 authentication-type is configured as password, then this key must be configured.</p> <p>By default, no authentication key is configured.</p> <p>The no form of the command removes the authentication key.</p>
Default	no authentication-key — No authentication key is defined.
Parameters	<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>hash — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.</p>

Interface/Virtual Link Commands

hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.

authentication-type

Syntax	authentication-type { password message-digest } no authentication-type
Context	config>router>ospf>area>interface <i>ip-int-name</i> config>router>ospf>area>virtual-link <i>router-id</i>
Description	<p>This command enables authentication and specifies the type of authentication to be used on the OSPF interface.</p> <p>Both simple password and message-digest authentication are supported.</p> <p>By default, authentication is not enabled on an interface.</p> <p>The no form of the command disables authentication on the interface.</p>
Default	no authentication — No authentication is enabled on an interface.
Parameters	<p>password — This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.</p> <p>message-digest — 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>

auth-keychain

Syntax	auth-keychain
Context	config>router>ospf>areas>interface config>router>ospf>areas>virtual-link config>service>vprn>ospf>areas>interface config>service>vprn>ospf>areas>sham-link config>service>vprn>ospf>areas>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	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

bfd-enable

Syntax	[no] bfd-enable [remain-down-on-failure]
Context	config>router>ospf>area>interface config>router>ospf3>area>interface
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 protocol adjacency.
Default	no bfd-enable
Parameters	remain-down-on-failure — Forces adjacency down on BFD failure.

dead-interval

Syntax	dead-interval <i>seconds</i> no dead-interval
Context	config>router>ospf>area>interface config>router>ospf>area>virtual-link
Description	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. The no form of the command reverts to the default value.
Default	40 seconds
Special Cases	OSPF Interface — If the dead-interval configured applies to an interface, then all nodes on the subnet must have the same dead interval. Virtual Link — If the dead-interval configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.
Parameters	<i>seconds</i> — The dead interval expressed in seconds. Values 1 — 65535

hello-interval

Syntax	hello-interval <i>seconds</i> no hello-interval
Context	config>router>ospf>area>interface config>router>ospf>area>virtual-link
Description	This command configures the interval between OSPF hellos issued on the interface or virtual link.

Interface/Virtual Link Commands

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.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

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

Default **hello-interval 10** — A 10-second hello interval.

Special Cases **OSPF Interface** — If the **hello-interval** configured applies to an interface, then all nodes on the subnet must have the same hello interval.

Virtual Link — If the **hello-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.

Parameters *seconds* — The hello interval in seconds expressed as a decimal integer.

Values 1 — 65535

interface

Syntax [no] **interface** *ip-int-name* [**secondary**]

Context config>router>ospf>area

Description This command creates a context to configure an OSPF interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.

The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the config>router>ospf>interface context can be used to disable an interface without removing the configuration for the interface.

Default **no interface** — No OSPF interfaces are defined.

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, etc.), 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.

interface-type

Syntax	interface-type { broadcast point-to-point } no interface-type
Context	config>router>ospf>area>interface
Description	<p>This command configures the interface type to be 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 broadcast adjacency maintenance overhead of the Ethernet link provided 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 OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.</p> <p>The no form of the command reverts to the default value.</p>
Default	<p>point-to-point if the physical interface is SONET.</p> <p>broadcast if the physical interface is Ethernet or unknown.</p>
Special Cases	Virtual-Link — A virtual link is always regarded as a point-to-point interface and not configurable.
Parameters	<p>broadcast — Configures the interface to maintain this link as a broadcast network. 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 media such as Ethernet.</p> <p>point-to-point — Configures the interface to maintain this link as a point-to-point link.</p>

message-digest-key

Syntax	message-digest-key <i>keyid</i> md5 [<i>key</i> <i>hash-key</i>] [hash] no message-digest-key <i>keyid</i>
Context	config>router>ospf>area>interface config>router>ospf>area>virtual-link
Description	<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 no form of the command removes the message digest key identified by the <i>key-id</i>.</p>
Default	No message digest keys are defined.
Parameters	<p>keyid — The <i>keyid</i> is expressed as a decimal integer.</p> <p>Values 1 — 255</p> <p>md5 key — The MD5 key. The <i>key</i> can be any alphanumeric string up to 16 characters in length.</p> <p>md5 hash-key — The MD5 hash key. The key can be any combination of ASCII characters up to 32 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>

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

metric

Syntax	metric <i>metric</i> no metric
Context	config>router>ospf>area>interface
Description	<p>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.</p> <p>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.</p>
Default	no metric — The metric is based on reference-bandwidth setting and the link speed.
Parameters	<p><i>metric</i> — The metric to be applied to the interface expressed as a decimal integer.</p> <p>Values 1 — 65535</p>

mtu

Syntax	mtu <i>bytes</i> no mtu
Context	config>router>ospf>area>interfac
Description	<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> <p>config>port>ethernet config>port>sonet-sdh>path</p> <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>Use the no form of this command to revert to default.</p>
Default	no mtu — Uses the value derived from the MTU configured in the config>port context.
Parameters	<p><i>bytes</i> — The MTU to be used by OSPF for this logical interface in bytes.</p> <p>Values 512 — 9198 (9212 — 14) (Depends on the physical media)</p>

passive

Syntax	[no] passive
Context	config>router>ospf>area>interface
Description	<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 passive 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>The no form of the command removes the passive property from the OSPF interface.</p>
Default	<p>Service interfaces defined in config>router>service-prefix are passive.</p> <p>All other interfaces are not passive.</p>

priority

Syntax	priority <i>number</i> no priority
Context	config>router>ospf>area>interface
Description	<p>This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.</p> <p>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.</p> <p>The no form of the command reverts the interface priority to the default value.</p>
Default	priority 1
Parameters	<p><i>number</i> — 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.</p> <p>Values 0 — 255</p>

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> no retransmit-interval
Context	config>router>ospf>area>interface config>router>ospf>area>virtual-link
Description	<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 acknowledgement has been received, the LSA will be retransmitted.</p> <p>The no form of this command reverts to the default interval.</p>
Default	retransmit-interval 5
Parameters	<i>seconds</i> — The retransmit interval in seconds expressed as a decimal integer.
	Values 1 — 1800

transit-delay

Syntax	transit-delay <i>seconds</i> no transit-delay
Context	config>router>ospf>area>interface config>router>ospf>area>virtual-link
Description	<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 no form of this command reverts to the default delay time</p>
Default	transit-delay 1
Parameters	<i>seconds</i> — The transit delay in seconds expressed as a decimal integer.
	Values 1 — 1800

virtual-link

Syntax	[no] virtual-link <i>router-id</i> transit-area <i>area-id</i>
Context	config>router>ospf>area
Description	<p>This command configures a virtual link to connect area border routers to the backbone via a virtual link.</p> <p>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</p>

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. (*Default: none specified*)

Default No virtual link is defined.

Parameters *router-id* — The router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area 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 16](#)) 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 16: OSPF Areas

Show Commands

ospf

Syntax	ospf [<i>ospf-instance</i>]
Context	show>router
Description	This command enables the context to display OSPF information.
Parameters	<i>ospf-instance</i> — Clears the specified OSPF instance.. Values 1 — 31

ospf3

Syntax	ospf3 [<i>ospf-instance</i>]
Context	show>router
Description	This command enables the context to display OSPF3 information.
Parameters	<i>ospf-instance</i> — Clears the specified VR-ID. Values 0 — 31 64..95 0 — 31 ipv6-unicast address-family 64 — 95 ipv4-unicast address-family

area

Syntax	area [<i>area-id</i>] [detail] [lfa]
Context	show>router>ospf
Description	This command displays configuration information about all areas or the specified area. When detail is specified operational and statistical information will be displayed.
Parameters	<i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer. detail — Displays detailed information about the specified area. lfa — Displays Loop-Free Alternate (LFA) next-hop information.

Output **OSPF Area Output** — The following table describes the standard and detailed command output fields for an OSPF area.

Label	Description
Area Id	A 32 bit integer uniquely identifying an area.
Type	NSSA — This area is configured as an NSSA area. Standard — This area is configured as a standard area (not NSSA or Stub). 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 LSA's.
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.
Blackhole Range	False — No blackhole route is installed for aggregates configured in this area.

Label	Description (Continued)
	True — A lowest priority blackhole route is installed for aggregates configured in this area.

Sample Output

```
A:SetupCLI# show router ospf area detail
=====
OSPF Areas (Detailed)
=====
-----
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 : True          Unknown 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              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 : 0xf493
Blackhole Range : False         Unknown LSAs : 0
-----
Area Id: 2.2.2.2
-----
Area Id      : 2.2.2.2          Type      : Standard
Key Rollover Int.: 10          LFA       : Include
Virtual Links : 1              Total Nbrs  : 0
Active IFs    : 0              Total IFs   : 0
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 : True          Unknown LSAs : 0
=====
A:SetupCLI#
```

Show Commands

```
A:SR# show router ospf area detail
=====
OSPF Areas (Detailed)
=====
-----
Area Id: 0.0.0.0
-----
Area Id      : 0.0.0.0          Type      : Standard
Virtual Links : 0              Total Nbrs   : 2
Active IFs    : 3              Total IFs    : 3
Area Bdr Rtrs : 0              AS Bdr Rtrs  : 0
SPF Runs      : 7              Last SPF Run : 10/26/2006 10:09:18
Router LSAs   : 3              Network LSAs : 3
Summary LSAs  : 0              Asbr-summ LSAs : 0
Nssa ext LSAs : 0              Area opaque LSAs : 3
Total LSAs    : 9              LSA Cksum Sum : 0x28b62
Blackhole Range : True          Unknown LSAs  : 0
=====
*A:Bombadil# show router ospf area 0.0.0.0 detail

=====
OSPF Area (Detailed) : 0.0.0.0
=====
-----
Configuration
-----
Area Id      : 0.0.0.0          Type      : Standard
-----
Statistics
-----
Virtual Links : 0              Total Nbrs   : 2
Active IFs    : 3              Total IFs    : 3
Area Bdr Rtrs : 0              AS Bdr Rtrs  : 0
SPF Runs      : 7              Last SPF Run : 10/26/2006 10:09:18
Router LSAs   : 3              Network LSAs : 3
Summary LSAs  : 0              Asbr-summ LSAs : 0
Nssa ext LSAs : 0              Area opaque LSAs : 3
Total LSAs    : 9              LSA Cksum Sum : 0x28b62
Blackhole Range : True          Unknown LSAs  : 0
=====

*A:Dut-B# show router ospf area 0.0.0.0 lfa
=====
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 area 0.0.0.0 lfa detail
=====
Path Table
=====
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 `area [router-id]`

Context `show>router>ospf`
`show>router>ospf3`

Description This command displays the entries in the Router Information (RI) LSAs.

Parameters *router-id* — Only the LSAs related to that router-id are listed. 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:cses-V95# show router ospf capabilities
=====
vrId 1 OSPFv2 instance 0 router capabilities
=====

```

Show Commands

```
scope      Router Id      Capabilities
-----
Area       0.0.0.95       0x14: Stub  P2P-VLAN
-----
No. of LSAs: 1
=====
```

database

Syntax **database** [**type** {**router** | **network** | **summary** | **asbr-summary** | **external** | **nssa** | **all**}] [**area** *area-id*] [**adv-router** *router-id*] [*link-state-id*] [**detail**]

Context show>router>ospf

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 *ospf-instance* — The OSPF instance.

Values 1 — 4294967295

type *keyword* — Specifies to filter the OSPF LSDB information based on the type specified by *keyword*.

type **router** — Display only router (Type 1) LSAs in the LSDB.

type **network** — Display only network (Type 2) LSAs in the LSDB.

type **summary** — Display only summary (Type 3) LSAs in the LSDB.

type **asbr-summary** — Display only ASBR summary (Type 4) LSAs in the LSDB.

type **external** — Display 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.

type **nssa** — Displays only NSSA area-specific AS external (Type 7) LSAs in the LSDB.

type **all** — Display 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* — Display LSDB information associated with the specified OSPF *area-id*.

adv-router *router-id* [*link-state-id*] — Display 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.

Output **OSPF Database Output** — The following table describes the standard and detailed command output fields for an OSPF database.

Label	Description
Area Id	The OSPF area identifier.
Type LSA Type	Router — LSA type of router (OSPF)

Label	Description (Continued)
	Network — LSA type of network (OSPF)
	Summary — LSA type of summary (OSPF)
	ASBR Summary — LSA type of ASBR summary (OSPF)
	Nssa-ext — LSA area-specific, NSSA external (OSPF)
	Area opaque — LSA type of area opaque (OSPF)
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 Adv Router Id	The router identifier of the router advertising the LSA.
Age	The age of the link state advertisement in seconds.
Sequence Sequence No	The signed 32-bit integer sequence number.
Cksum 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 DC — Demand Circuit Support R — If clear, a node can participates in OSPF topology distribution without being used to forward transit traffic. N — Type 7 LSA Support MC — Multicast Support E — External Routes Support
Prefix Options	P — Propagate NSSA LSA. MC — Multicast support.
Flags	None — No flags set V — The router is an endpoint for one or more fully adjacent Virtual Links having the described area as the transit area E — The router is an AS Boundary Router B — The router is an Area Border Router
Link Count	The number of links advertised in the LSA.
Link Type (n)	The link type of the <i>n</i> th link in the LSA.

Show Commands

Label	Description (Continued)
Network (<i>n</i>)	The network address 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.

Sample Output

```
A:ALA-A# show router ospf 1 database
=====
OSPF Link State Database (Type : All)
=====
Area Id          Type      Link State Id  Adv Rtr Id    Age  Sequence      Cksun
-----
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 0xcfbf
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#

A:ALA-A# show router ospf database detail
=====
OSPF Link State Database (Type : All) (Detailed)
-----
Router LSA for Area 0.0.0.0
-----
Area Id          : 0.0.0.0          Adv Router Id    : 180.0.0.2
Link State Id     : 180.0.0.2          LSA Type         : Router
Sequence No      : 0x800000b7       Checksum         : 0xd55
Age              : 155           Length           : 192
Options          : E
Flags            : None
Link Type (1)    : Point To Point
Nbr Rtr Id (1)   : 180.0.0.13     I/F Address (1)  : 180.0.22.2
No of TOS (1)    : 0
Link Type (2)    : Stub Network
Metric-0 (1)     : 25
```

Network (2)	: 180.0.22.0	Mask (2)	: 255.255.255.0
No of TOS (2)	: 0	Metric-0 (2)	: 25
Link Type (3)	: Point To Point		
Nbr Rtr Id (3)	: 180.0.0.12	I/F Address (3)	: 180.0.5.2
No of TOS (3)	: 0	Metric-0 (3)	: 25
Link Type (4)	: Stub Network		
Network (4)	: 180.0.5.0	Mask (4)	: 255.255.255.0
No of TOS (4)	: 0	Metric-0 (4)	: 25
Link Type (5)	: Point To Point		
Nbr Rtr Id (5)	: 180.0.0.8	I/F Address (5)	: 180.0.13.2
No of TOS (5)	: 0	Metric-0 (5)	: 6
Link Type (6)	: Stub Network		
Network (6)	: 180.0.13.0	Mask (6)	: 255.255.255.0
No of TOS (6)	: 0	Metric-0 (6)	: 6
Link Type (7)	: Point To Point		
Nbr Rtr Id (7)	: 180.0.0.5	I/F Address (7)	: 180.0.14.2
No of TOS (7)	: 0	Metric-0 (7)	: 6
Link Type (8)	: Stub Network		
Network (8)	: 180.0.14.0	Mask (8)	: 255.255.255.0
No of TOS (8)	: 0	Metric-0 (8)	: 6
Link Type (9)	: Point To Point		
Nbr Rtr Id (9)	: 180.0.0.11	I/F Address (9)	: 180.0.17.2
No of TOS (9)	: 0	Metric-0 (9)	: 25
Link Type (10)	: Stub Network		
Network (10)	: 180.0.17.0	Mask (10)	: 255.255.255.0
No of TOS (10)	: 0	Metric-0 (10)	: 25
Link Type (11)	: Stub Network		
Network (11)	: 180.0.0.2	Mask (11)	: 255.255.255.255
No of TOS (11)	: 0	Metric-0 (11)	: 1
Link Type (12)	: Stub Network		
Network (12)	: 180.0.18.0	Mask (12)	: 255.255.255.0
No of TOS (12)	: 0	Metric-0 (12)	: 24
Link Type (13)	: Point To Point		
Nbr Rtr Id (13)	: 180.0.0.10	I/F Address (13)	: 180.0.3.2
No of TOS (13)	: 0	Metric-0 (13)	: 25
Link Type (14)	: Stub Network		
Network (14)	: 180.0.3.0	Mask (14)	: 255.255.255.0
No of TOS (14)	: 0	Metric-0 (14)	: 25

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#

interface

Syntax `interface [ip-addr | ip-int-name | area area-id] [detail]`

Context `show>router>ospf`

Description Displays the details of the OSPF interface, this interface can be identified by ip-address or ip interface name. When neither is specified, all in-service interfaces are displayed.

The **detail** option produces a great amount of data. It is recommended to detail only when requesting a specific interface.

Parameters *ip-addr* — Display only the interface identified by this IP address.

ip-int-name — Display only the interface identified by this interface name.

area *area-id* — Display all interfaces configured in this area.

detail — Displays detailed information on the interface.

Output **Standard OSPF Interface Output** — The following table describes the standard command output fields for an OSPF interface.

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. wait — The router is trying to determine the identity of the (Backup) Designated router for the network. PTOP — The interface is operational, and connects either to a physical point-to-point network or to a virtual link. DR — This router is the Designated Router for this network. BDR — This router is the backup Designated Router for this network.

Label	Description (Continued)
	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 "DUTC_TO_DUTB.1.0"
=====
OSPFv2 (0) 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 interface "DUTC_TO_DUTB.1.0" detail
=====
OSPFv2 (0) interface "DUTC_TO_DUTB.1.0" (detailed)
=====
-----
Configuration
-----
IP Address       : 1.0.23.3
Area Id          : 0.0.0.0
Hello Intrvl     : 2 sec
Retrans Intrvl   : 5 sec
Cfg Metric       : 7000
Transit Delay    : 1
Passive          : False
LSA-filter-out   : None
LFA              : Include
Auth Type        : None
Priority          : 1
Rtr Dead Intrvl  : 10 sec
Poll Intrvl      : 120 sec
Advert Subnet    : True
Cfg IF Type      : Point To Point
Cfg MTU          : 0
Adv Rtr Capab    : Yes
LFA NH Template  : template1
-----
State
-----
Admin Status     : Enabled
Designated Rtr   : 0.0.0.0
IF Type          : Point To Point
Oper MTU         : 1500
Oper Metric      : 7000
Te Metric        : 7000
Admin Groups     : None
Ldp Sync         : outOfService
Ldp Timer State  : Disabled
Oper State       : Point To Point
Backup Desig Rtr : 0.0.0.0
Network Type     : Transit
Last Enabled     : 01/14/2014 14:33:07
Bfd Enabled      : No
Te State         : Down
Ldp Sync Wait    : Disabled
Ldp Tm Left      : 0
-----
Statistics
-----
Nbr Count        : 1
Tot Rx Packets    : 603
Rx Hellos         : 576
If Events         : 1
Tot Tx Packets    : 602
Tx Hellos         : 577
```

Show Commands

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 Hello Int.	: 1
Bad Dead Int.	: 1	Bad Options	: 0
Bad Versions	: 0	Bad Checksums	: 0
LSA Count	: 0	LSA Checksum	: 0x0

*A:Dut-C#

A:SetupCLI# show router ospf interface "ip_if_1" detail

=====

OSPF Interface (Detailed) : ip_if_1

=====

Configuration

IP Address	: 10.10.1.1	Priority	: 10
Area Id	: 0.0.0.0	Rtr Dead Intrvl	: 45 sec
Hello Intrvl	: 9 sec	Poll Intrvl	: 120 sec
Retrans Intrvl	: 10 sec	Advert Subnet	: True
Cfg Metric	: 11	Auth Type	: MD5
Transit Delay	: 2	Cfg MTU	: 9198
Passive	: False		
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	: Secondary	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 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

 Configuration

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		
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 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 interface area 1.1.1.1 detail

=====

OSPF Interfaces in Area (Detailed) : 1.1.1.1

=====

Show Commands

```
-----
Interface : ip_if_1
-----
IP Address      : 10.10.1.1
Area Id         : 1.1.1.1
Hello Intrvl    : 9 sec
Retrans Intrvl  : 10 sec
Cfg Metric      : 11
Transit Delay   : 2
Passive         : False
LFA             : Exclude
IPsec InStatSA  :
IPsec InStatSAmp:
Admin Status    : Enabled
Designated Rtr  : 0.0.0.0
IF Type         : Point To Point
Oper MTU        : 1576
Oper Metric     : 11
Te Metric       : 16777215
Admin Groups    : None
Ldp Sync        : outOfService
Ldp Timer State : Disabled
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
Priority         : 10
Rtr Dead Intrvl : 45 sec
Poll Intrvl     : 120 sec
Advert Subnet   : False
Auth Type       : MD5
Cfg MTU         : 9198
IPsec OutStatSA :
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
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
=====
A:SetupCLI#

A:SetupCLI# show router ospf 1 interface detail
=====
OSPF Interfaces (Detailed)
-----
Interface : system
-----
IP Address      : 9.1.255.255
Area Id         : 0.0.0.0
Hello Intrvl    : 10 sec
Retrans Intrvl  : 5 sec
Cfg Metric      : 0
Transit Delay   : 1
Passive         : True
Admin Status    : Enabled
Designated Rtr  : 2.2.2.2
IF Type         : Broadcast
Oper MTU        : 1500
Priority         : 1
Rtr Dead Intrvl : 40 sec
Poll Intrvl     : 120 sec
Advert Subnet   : True
Auth Type       : None
Cfg MTU         : 0
Oper State      : Designated Rtr
Backup Desig Rtr : 0.0.0.0
Network Type    : Transit
Last Enabled    : 05/14/2006 09:16:26
```

```

Oper Metric      : 0
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

Bfd Enabled     : No
If Events       : 5
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 Addr   : 0
Auth Failures   : 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.

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.

Label	Description (Continued)
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.
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	<p>Disabled — OSPF on this interface is administratively shut down.</p> <p>Enabled — OSPF on this interface is administratively enabled.</p>

Label	Description (Continued)
Oper State	<p data-bbox="634 302 1382 365">Down — This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.</p> <p data-bbox="634 386 1333 449">Waiting — The router is trying to determine the identity of the (Backup) Designated router for the network.</p> <p data-bbox="634 470 1349 533">Point To Point — The interface is operational, and connects either to a physical point-to-point network or to a virtual link.</p> <p data-bbox="634 554 1373 617">Designated Rtr — This router is the Designated Router for this network.</p> <p data-bbox="634 638 1354 735">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 data-bbox="634 756 1333 819">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 data-bbox="634 1075 1081 1106">Broadcast — LANs, such as Ethernet.</p> <p data-bbox="634 1127 1219 1159">NBMA — X.25, Frame Relay and similar technologies.</p> <p data-bbox="634 1180 1333 1211">Point-To-Point — Links that are definitively point to point.</p>
Network Type	<p data-bbox="634 1232 1382 1329">Stub — OPSF 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.</p> <p data-bbox="634 1350 1382 1446">Transit — OPSF 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.</p>
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.

Label	Description (Continued)
If Events	The number of times this OSPF interface has changed its state, or an error has occurred since this interface was last enabled.
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 Acknowledgements received on this interface since this interface was last enabled.
Tx LS Acks	The total number of Link State Acknowledgements 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.

Label	Description (Continued)
Bad Areas	The total number of OSPF packets received with an area mismatch since this interface was last enabled.
Bad Dest Addr	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.

Label	Description (Continued)
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 interface area 1 detail
=====
OSPF Interfaces in Area (Detailed) : 1
=====
Interface : ip-10.10.1.1
-----
IP Address      : 10.10.1.1
Area Id         : 0.0.0.1
Hello Intrvl    : 5 sec
Retrans Intrvl  : 5 sec
Cfg Metric      : 0
Transit Delay   : 1
Passive         : False
Admin Status    : Enabled
Designated Rtr  : 10.20.1.1
IF Type         : Broadcast
Oper MTU        : 1500
Oper Metric     : 1000
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
TE Metric       : 678
Priority         : 1
Rtr Dead Intrvl : 15 sec
Poll Intrvl     : 120 sec
Advert Subnet    : True
Auth Type       : None
Cfg MTU         : 0
Oper State      : Designated Rtr
Backup Desig Rtr : 0.0.0.0
Network Type     : Transit
Last Enabled     : 04/11/2007 16:06:27
Bfd Enabled      : No
If Events        : 5
Tot Tx Packets   : 1116
Tx Hellos        : 1116
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
=====
*A:JC-NodeA#
```

lfa-coverage

Syntax **lfa-coverage**

Context show>router>ospf

Description This command displays OSPF Loop-Free Alternate (LFA) next-hop information.

Sample Output

```
*A:Dut-A# show router ospf lfa-coverage
=====
LFA coverage ospfv2 instance 0
=====
Area                Node                Prefix
-----
0.0.0.0              4/4 (100%)          8/8 (100%)
=====
*A:Dut-A#
```

neighbor

Syntax **neighbor** [*ip-int-name*] [*router-id*]

Context show>router>ospf

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* — Display neighbor information only for neighbors of the interface identified by the interface name

router-id — Display neighbor information for the neighbor identified by the the specified router ID.

Output **Standard OSPF Neighbor Output** — The following table describes the standard command output fields for an OSPF neighbor.

Label	Description
Nbr IP Addr	The IP address this neighbor is using in its IP Source Address. Note that, 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	Down — This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor.

Label	Description (Continued)
	<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 (i.e., 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>
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
=====
OSPF 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.

Label	Description
Neighbor IP Addr	The IP address this neighbor is using in its IP source address. Note that, 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.
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 (i.e., 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>

Label	Description (Continued)
	Full — In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.
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	E — External Routes Support MC — Multicast Support N/P — Type 7 LSA Support EA — External Attribute LSA Support DC — Demand Circuit Support O — Opaque LSA Support
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.

Label	Description (Continued)
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:ALA-A# show router ospf neighbor detail

=====

OSPF Neighbors

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

Show Commands

```
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 [link link-id | area area-id |as] [adv-router router-id] [ls-id] [detail]`

Context `show>router>ospf`

Description This command displays OSPF opaque database information.

Output **OSPF Opaque Database Output** — The following table describes the 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.
Type	NSSA — This area is configured as a NSSA area. Area — This area is configured as a standard area (not NSSA or stub). Stub — This area is configured as a 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.

Label	Description (Continued)
Sequence	The signed 32-bit integer sequence number.
Cksum	The 32-bit unsigned sum of the link-state advertisements' LS check-sums.

Sample Output

```
A:ALA-A# show router ospf opaque-database
=====
OSPF Opaque Link State Database (Type : All)
=====
Area Id          Type  Link State Id    Adv Rtr Id      Age  Sequence      Cksum
-----
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-A# show router ospf opaque-database adv-router 10.20.1.1 detail
=====
OSPF Opaque Link State Database (Type : All) (Detailed)
=====
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      : 0x80000028      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
```

Show Commands

```
Link State Id      : 1.0.0.2          LSA Type          : Area Opaque
Sequence No       : 0x8000000d        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#
```

range

Syntax	range [<i>area-id</i>]
Context	show>router>ospf
Description	This command displays ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression.
Parameters	<i>area-id</i> — Display the configured ranges for the specified area.
Output	OSPF Range Output — The following table describes the 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.
Advertise	False — The specified address/mask is not advertised outside the area. True — The specified address/mask is advertised outside the area.

Label	Description (Continued)
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.
	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
=====
OSPF 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
=====
OSPF Ranges for Area Id : 180.0.7.9
=====
Area Id          Address/Mask      Advertise    LSDb Type
-----
No. of Ranges: 0
=====
A:ALA-A#
```

routes

Syntax	routes [<i>ip-prefix[/prefix-length]</i>] [type] [detail] [alternative] [summary] [exclude-shortcut]
Context	show>router>ospf
Description	This command information about OSPF routes.
Parameters	<p><i>ip-prefix[/prefix-length]</i> — Displays information about the specified IP prefix and length.</p> <p>type — Displays information about the specified type.</p> <p>Values intra-area, inter-area, external-1, external-2, nssa-1, nssa-2</p> <p>detail — Displays detailed information about the routes.</p> <p>alternative — Displays the level of protection per prefix (ref. show router isis routes alternative)</p> <p>summary — Displays summarized information about the routes.</p>

Show Commands

Sample Output

```
*A:Dut-C# show router ospf routes exclude-shortcut alternative detail
```

```
=====
OSPFv2 (0) Routing Table excluding shortcuts (detailed)
=====
```

Destination	Type (Dest)	Stat		
NHIF	NHIF	Cost [E2]	Area	Tunnel-Information
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
```

```
=====
OSPFv2 Routing Table (detailed)
```

```
=====
Destination      Type (Dest)      Stat
NHIP             NHIF      Cost [E2]   Area
A-NHIP (L)      A-NHIF    A-Cost [E2]  A-Type      Tunnel-Information
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
10.20.1.6/32	IA (HOST)	N (R)		
1.1.3.3	3	30	0.0.0.0	
1.1.2.2	2	30	0.0.0.0	
10.20.1.2/0	IA (RTR)	N (N)		
1.1.2.2	2	10	0.0.0.0	
10.20.1.3/0	IA (RTR)	N (N)		
1.1.3.3	3	10	0.0.0.0	

Show Commands

```
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-A#

*A:Dut-C>config>router>mpls>lsp# show router ospf 1 routes 10.0.0.2/32 detail

=====
OSPFv2 (1) Routing Table (detailed)
=====
Destination      Type (Dest)      Stat
  NHIP           NHIF      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
=====
```

spf

Syntax **spf [lfa]****Context** show>router>ospf**Description** This command displays statistics of shortest-path-first (SPF) calculations.**Parameters** **lfa** — Displays Loop-Free Alternate (LFA) next-hop information.**Output** **SPF Output Fields** — The following table describes SPF output fields.

Label	Description
Total SPF Runs	The total number of incremental SPF runs triggered by new or updated LSAs.
Last Full SPF run @	The date and time when the external OSPF Dijkstra (SPF) was last run.
Last Full SPF Time	The length of time, in seconds, when the last full SPF was run.
Intra SPF Time	The time when intra-area SPF was last run on this area.
Inter SPF Time	The total number of incremental SPF runs triggered by new or updated type-3 and type-4 summary LSAs.
Extern SPF Time	The total number of incremental SPF runs triggered by new or updated type-5 external LSAs.
RTM Updt Time	The time, in hundredths of seconds, used to perform a total SPF calculation.
Min/Avg/Max Full SPF Time	Min — The minimum time, in hundredths of seconds, used to perform a total SPF calculation. Avg — The average time, in hundredths of seconds, of all the total SPF calculations performed by this OSPF router. Max — The maximum time, in hundredths of seconds, used to perform a total 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-A# show router ospf spf lfa
=====
OSPF SPF Statistics
=====
Total SPF Runs           :    6
Last Full SPF run @      : 02/20/2012 09:19:35
Last Full SPF Time       : < 0.01 secs

```

Show Commands

```
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:Dut-A#
```

```
A:ALA-A#  show router ospf 1 spf
=====
OSPF 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#
```


statistics

Syntax **statistics****Context** show>router>ospf**Description** This command displays the global OSPF statistics.**Output** **OSPF Statistics Output Fields** — The following table describes the command output fields for OSPF statistics.

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 Acknowledgements (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.

Show Commands

Label	Description
Ext SPF Runs	The total number of incremental SPF runs triggered by new or updated type-5 external LSAs.
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 Addrs	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.

Sample Output

```
A:ALA-A# show router ospf 1 statistics
```

```
=====
```

OSPF 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 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
Failed SPF Attempts: 0
CSPF Requests   : 0           CSPF Request Drops : 0
CSPF Path Found : 0           CSPF Path Not Found: 0
=====
A:ALA-A#

```

status

Syntax **status****Context** show>router>ospf**Description** Displays the general status of OSPF.**Output** **OSPF Status Output Fields** — The following table describes the command output fields for OSPF status.

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. 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. 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. 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. True — This router is an area border router.
AS Border Router	False — This router is not configured as an Autonomous System border router. True — This router is configured as an Autonomous System border router.
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-C# show router ospf status
=====
OSPFv2 (0) Status
=====
OSPF Cfg Router Id       : 10.20.1.3
OSPF Oper Router Id     : 10.20.1.3
OSPF Version             : 2
OSPF Admin Status       : Enabled
OSPF Oper Status        : Enabled
Graceful Restart        : Disabled
GR Helper Mode          : Disabled
Preference               : 10
External Preference     : 150
Backbone Router         : True
Area Border Router      : False
AS Border Router        : True
Opaque LSA Support      : True
Traffic Engineering Support : False
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           : 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       : 0x21502
OSPF Last Enabled      : 01/14/2014 14:33:07
Unicast Import          : True
Multicast Import        : False
Export Policies         : static
Import Policies         : None
Lfa Policies            : pol1
                        : pol2
                        : pol3
                        : pol4
                        : pol5
OSPF Ldp Sync Admin Status : Enabled
LDP-over-RSVP          : Disabled
RSVP-Shortcut          : Disabled
Advertise-Tunnel-Link   : Disabled
LFA                    : Enabled
Export Limit            : 0
Export Limit Log Percent : 0
Total Exp Routes        : 1

```

Show Commands

```
=====*A:Dut-C#

A:SetupCLI# show router ospf status
=====
OSPF Status
=====
OSPF Cfg Router Id       : 255.255.255.255
OSPF Oper Router Id     : 10.20.30.40
OSPF Version            : 2
OSPF Admin Status       : Enabled
OSPF Oper Status        : Enabled
Graceful Restart        : Disabled
GR Helper Mode          : Disabled
Preference              : 11
External Preference     : 150
Backbone Router         : True
Area Border Router     : False
AS Border Router        : True
Opaque LSA Support      : True
Traffic Engineering Support : True
RFC 1583 Compatible     : True
Demand Exts Support     : False
In Overload State       : True (Indefinitely in overload)
In External Overflow State : False
Exit Overflow Interval  : 0
Last Overflow Entered   : Never
Last Overflow Exit      : 10/01/2011 07:34:03
External LSA Limit      : -1
Reference Bandwidth     : 10,000 Kbps
Init SPF Delay          : 10000 msec
Sec SPF Delay           : 10000 msec
Max SPF Delay           : 12000 msec
Min LS Arrival Interval : 600000 msec
Init LSA Gen Delay      : 100 msec
Sec LSA Gen Delay       : 400 msec
Max LSA Gen Delay       : 600000 msec
Last Ext SPF Run        : Never
Ext LSA Cksum Sum       : 0x0
OSPF Last Enabled      : 10/01/2011 07:34:03
Multicast Import        : False
Export Policies         : test 567890123456789012345678901
                        : test 567890123456789012345678902
                        : test 567890123456789012345678903
                        : test 567890123456789012345678904
                        : test 567890123456789012345678905
Import Policies         : test 567890123456789012345678901
                        : test 567890123456789012345678902
                        : test 567890123456789012345678903
                        : test 567890123456789012345678904
                        : test 567890123456789012345678905
OSPF Ldp Sync Admin Status : Disabled
LDP-over-RSVP           : Enabled
RSVP-Shortcut           : Enabled
Advertise-Tunnel-Link   : Enabled
LFA                     : Enabled
Export Limit            : 0
Export Limit Log Percent : 0
Total Exp Routes        : 0
```

```

=====
A:SetupCLI#

A:ALA-A#  show router ospf 1 status
=====
OSPF Status
=====
OSPF Router Id           : 10.13.7.165
OSPF Version             : 2
OSPF Admin Status        : Enabled
OSPF Oper Status         : Enabled
Graceful Restart         : Enabled
GR Helper Mode           : Disabled
Preference                : 10
External Preference      : 150
Backbone Router          : True
Area Border Router       : True
AS Border Router         : True
Opaque LSA Support       : True
Traffic Engineering Support : True
RFC 1583 Compatible      : True
TOS Routing 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            : 500 msec
Sec SPF Delay             : 2000 msec
Max SPF Delay             : 15000 msec
Min LS Arrival Interval   : 500 msec
Max LSA Gen Delay         : 5000 msec
Last Ext SPF Run          : Never
Ext LSA Cksum Sum         : 0x2afce
OSPF Last Enabled        : 05/23/2006 23:34:36
Export Policies           : export-static
=====
A:ALA-A#

```

virtual-link

Syntax `virtual-link [detail]`**Context** `show>router>ospf`**Description** This command displays information for OSPF virtual links.**Parameters** **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.

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.
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.

Label	Description (Continued)
Rx LS Acks	Specifies the total number of Link State Acknowledgements 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.
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 Acknowledgements (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 Addr	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.

Label	Description (Continued)
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.
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
=====
OSPF 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#
```

```
A:ALA-A# show router ospf virtual-link detail
=====
OSPF Virtual Links (detailed)
=====
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 Pkt Types    : 0
Bad Lengths     : 0                    Bad Hello Int.   : 0
Bad Dead Int.   : 0                    Bad Options      : 0
Retrans Intrvl  : 5 sec                 Transit Delay    : 1 sec
Last Event      : 11/07/2006 17:11:56 Authentication : None

```

```
-----
Neighbor Router Id : 180.0.0.10

```

```

Nbr Router Id   : 180.0.0.10          Area Id         : 0.0.0.2
Local Interface : 180.2.7.12           Metric          : 300
State           : Point To Point      Admin State      : Up
Hello Intrvl    : 10 sec               Rtr Dead Intrvl : 60 sec
Tot Rx Packets  : 43073                Tot Tx Packets   : 43034
Rx Hellos       : 24851                Tx Hellos        : 24844
Rx DBDs         : 3                    Tx DBDs          : 2
Rx LSRs         : 1                    Tx LSRs          : 1
Rx LSUs         : 18071                Tx LSUs          : 17853
Rx LS Acks      : 147                 Tx LS Acks       : 334
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 Pkt Types    : 0
Bad Lengths     : 0                    Bad Hello Int.   : 0
Bad Dead Int.   : 0                    Bad Options      : 0
Retrans Intrvl  : 5 sec                 Transit Delay    : 1 sec
Last Event      : 11/07/2006 17:12:00 Authentication : MD5

```

```
=====
A:ALA-A#

```

virtual-neighbor

Syntax `virtual-neighbor [remote router-id] [detail]`

Context `show>router>ospf`

Description This command displays virtual neighbor information.

Parameters **remote** *router-id* — Displays the specified router ID. This reduces the amount of output displayed.
detail — Produces detailed information on the virtual neighbor. This option produces a large amount of data. It is recommended to use **detail** only when requesting information for a specific neighbor.

Output **OSPF Virtual Neighbor Output** — The following table describes OSPF virtual neighbor output fields.

Label	Description
Nbr IP Addr	The IP address this neighbor is using in its IP source address. Note that, 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.

Label	Description (Continued)
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.
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
=====
OSPF 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
=====
OSPF Virtual Neighbors
=====
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--

```

Show Commands

```
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 MTUs       : 0
Bad Packets     : 0                      LSA not in LSDB : 0
Option Mismatches: 0                    Nbr Duplicates  : 0
-----
Virtual Neighbor Router Id : 180.0.0.10
-----
Neighbor IP Addr : 180.2.9.10           Neighbor Rtr Id  : 180.0.0.10
Neighbor State   : Full                  Transit Area     : 0.0.0.2
Retrans Q Length : 0                    Options          : -E--
Events          : 4                      Last Event Time  : 11/07/2006 17:11:59
Up Time         : 2d 17:47:14           Time Before Dead : 59 sec
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
=====
A:ALA-A#
```

Clear Commands

ospf

Syntax	ospf [<i>ospf-instance</i>]
Context	clear>router
Description	This command clears and resets OSPF protocol entities.
Parameters	<i>ospf-instance</i> — Clears the specified OSPF instance.
Values	1 — 31

ospf3

Syntax	ospf [<i>ospf-instance</i>]
Context	clear>router
Description	This command clears and resets OSPF3 protocol entities.
Parameters	<i>ospf-instance</i> — Clears the specified OSPF3 instance.
Values	0 — 31 64..95 0 — 31 ipv6-unicast address-family 64 — 95 ipv4-unicast address-family

database

Syntax	database [purge]
Context	clear>router>ospf
Description	This command clears all LSAs received from other nodes. Sets all adjacencies better then two way to one way. Refreshes all self originated LSAs
Parameters	purge — The purge parameter also clears all self-originated LSAs and re-originates all self-originated LSAs

Clear Commands

export

Syntax	export
Context	clear>router>ospf
Description	Re-evaluates all effective export policies

neighbor

Syntax	neighbor [<i>ip-int-name</i> <i>ip-address</i>]
Context	clear>router>ospf
Description	Marks the neighbor as dead and re-initiates the affected adjacencies.
Parameters	<i>ip-int-name</i> — Clear all neighbors for the interface specified by this interface name. <i>ip-address</i> — Clear all neighbors for the interface specified by this IP-address

statistics

Syntax	statistics
Context	clear>router>ospf
Description	Clears all neighbor, router, interface, SPF and global statistics of this OSPF instance.

OSPF Debug Commands

ospf

Syntax	ospf [<i>ospf-instance</i>]
Context	debug>router
Description	Indicates the OSPF instance for debugging purposes.
Parameters	<i>ospf-instance</i> — The OSPF instance. Values 1 — 31

ospf3

Syntax	ospf3 [<i>ospf-instance</i>]
Context	debug>router
Description	Indicates the OSPF3 instance for debugging purposes.
Parameters	<i>ospf-instance</i> — Clears the specified OSPF3 instance. Values 0 — 31 64..95 0 — 31 ipv6-unicast address-family 64 — 95 ipv4-unicast address-family

area

Syntax	area [<i>area-id</i>] no area
Context	debug>router>ospf
Description	This command enables debugging for an OSPF area.
Parameters	<i>area-id</i> — Specify the OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

OSPF Debug Commands

area-range

Syntax	area-range [<i>ip-address</i>] no area-range
Context	debug>router>ospf
Description	This command enables debugging for an OSPF area range.
Parameters	<i>ip-address</i> — Specify the IP address for the range used by the ABR to advertise the area into another area.

cspf

Syntax	cspf [<i>ip-address</i>] no cspf
Context	debug>router>ospf
Description	This command enables debugging for an OSPF constraint-based shortest path first (CSPF).
Parameters	<i>ip-address</i> — Specify the IP address for the range used for CSPF.

graceful-restart

Syntax	[no] graceful-restart
Context	debug>router>ospf
Description	This command enables debugging for OSPF graceful-restart.

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-address</i>] no interface
Context	debug>router>ospf
Description	This command enables debugging for an OSPF interface.
Parameters	<i>ip-int-name</i> — Specify 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. <i>ip-address</i> — Specify the interface's IP address.

leak

Syntax	leak [<i>ip-address</i>] no leak
Context	debug>router>ospf
Description	This command enables debugging for OSPF leaks.
Parameters	<i>ip-address</i> — Specify the IP address to debug OSPF leaks.

lsdb

Syntax	lsdb [type] [<i>ls-id</i>] [<i>adv-rtr-id</i>] [area <i>area-id</i>] no lsdb
Context	debug>router>ospf
Description	This command enables debugging for an OSPF link-state database (LSDB).
Parameters	<i>type</i> — Specifies the OSPF link-state database (LSDB) type. <div style="margin-left: 20px;">Values router, network, summary, asbr, extern, nssa, area-opaque, as-opaque, link-opaque</div> <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. <i>adv-rtr-id</i> — Specifies the router identifier of the router advertising the LSA. <i>area-id</i> — Specifies a 32-bit integer uniquely identifying an area.

misc

Syntax	[no] misc
Context	debug>router>ospf
Description	This command enables debugging for miscellaneous OSPF events.

neighbor

Syntax	neighbor [<i>ip-int-name</i> <i>ip-address</i>] no neighbor
Context	debug>router>ospf
Description	This command enables debugging for an OSPF neighbor.
Parameters	<i>ip-int-name</i> — Specifies the neighbor interface name.

OSPF Debug Commands

ip-address — Specifies neighbor information for the neighbor identified by the the specified router ID.

nssa-range

Syntax	nssa-range [<i>ip-address</i>] no nssa-range
Context	debug>router>ospf
Description	This command enables debugging for an NSSA range.
Parameters	<i>ip-address</i> — Specifies the IP address range to debug.

packet

Syntax	packet [<i>packet-type</i>] [<i>ip-address</i>] no packet
Context	debug>router>ospf
Description	This command enables debugging for OSPF packets.
Parameters	<i>packet-type</i> — Specifies the OSPF packet type to debug. Values hello, dbdescr, lsrequest, lsupdate, lsack <i>ip-address</i> — Specifies the IP address to debug. Values ipv4-address: a.b.c.d

rtm

Syntax	rtm [<i>ip-address</i>] no rtm
Context	debug>router>ospf
Description	This command enables debugging for OSPF RTM.
Parameters	<i>ip-address</i> — Specifies the IP address to debug. Values ipv4-address: a.b.c.d

spf

Syntax	spf [<i>type</i>] [<i>dest-addr</i>] no spf
Context	debug>router>ospf
Description	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.
Parameters	<i>type</i> — Specifies the area to debug Values intra-area, inter-area, external <i>dest-addr</i> — Specifies the destination IP address to debug.

virtual-neighbor

Syntax	virtual-neighbor [<i>ip-address</i>] no virtual-neighbor
Context	debug>router>ospf
Description	This command enables debugging for an OSPF virtual neighbor.
Parameters	<i>ip-address</i> — Specifies the IP address of the virtual neighbor.

In This Chapter

This chapter provides information to configure Intermediate System to Intermediate System (IS-IS).

Topics in this chapter include:

- [Configuring IS-IS on page 476](#)
 - [Routing on page 477](#)
 - [IS-IS Frequently Used Terms on page 479](#)
 - [ISO Network Addressing on page 480](#)
 - [ISO Network Addressing on page 480](#)
 - [IS-IS PDU Configuration on page 481](#)
 - [IS-IS Operations on page 482](#)
 - [IS-IS Route Summarization on page 483](#)
 - [IS-IS Administrative Tags on page 484](#)
 - [Segment Routing in Shortest Path Forwarding on page 485](#)
- [IS-IS Configuration Process Overview on page 512](#)
- [Configuration Notes on page 513](#)

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 17 displays an example of an IS-IS routing domain.

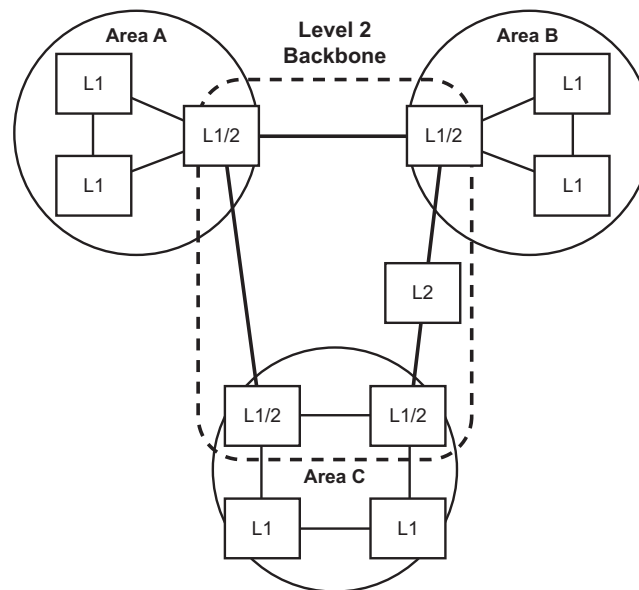


Figure 17: IS-IS Routing Domain

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 on page 480](#)). 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 will accept 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.

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.

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, an 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 18](#)).

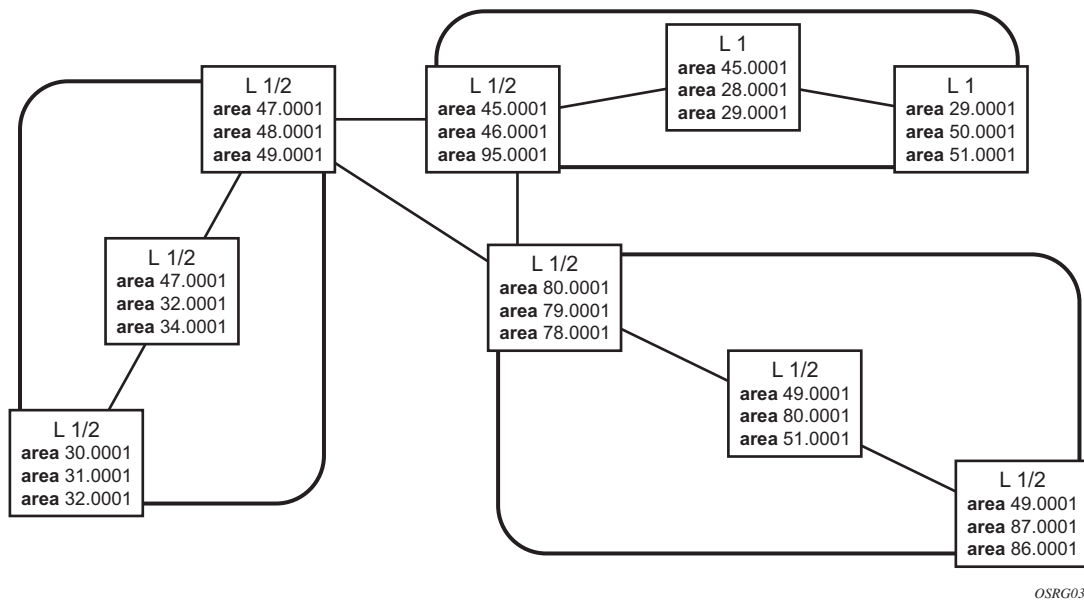


Figure 18: Using Area Addresses to Form Adjacencies

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.

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.
- ESSs can build a link-state PDU based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- ESSs 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.

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 will be the smallest metric of all the more specific IPv4 routes.
-

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 will not perform full SPF but will instead perform 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.

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.
-

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.

Using Route Tags

Although an operator on this or another (neighboring) IS-IS router has configured setting of the IS-IS administrative tags it will not have any effect unless policies are configured 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
```

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.

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).

Segment Routing in Shortest Path Forwarding

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 prefix SID is resolved, 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 IPv4 SR tunnel to the prefix destination is also added to the TTM.

The SR tunnel in TTM is available to be used in the following contexts:

- VLL, LDP VPLS, IES/VRPN spoke-interface, R-VPLS, BGP EVPN.
- BGP-AD VPLS, BGP-VPLS, BGP VPWS when the **use-provisioned-sdp** option is enabled in the binding to the PW template.
- Intra-AS BGP VPRN for vpn-ipv4 and vpn-ipv6 prefixes with both auto-bind and explicit SDP.
- Multicast over IES/VRPN spoke interface with spoke-sdp riding a SR tunnel.

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 a LFA backup next-hop. If the prefix of a given SR tunnel is not protected by the base LFA, the remote LFA will automatically compute a backup next-hop using a SR tunnel if the **remote-lfa** option is also enabled in the IGP instance.

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 will be 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:

```
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:

```
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 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, i.e., 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) 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.

The following is an example of a router advertising its loopback address and the resulting packet label encapsulation throughout the network.

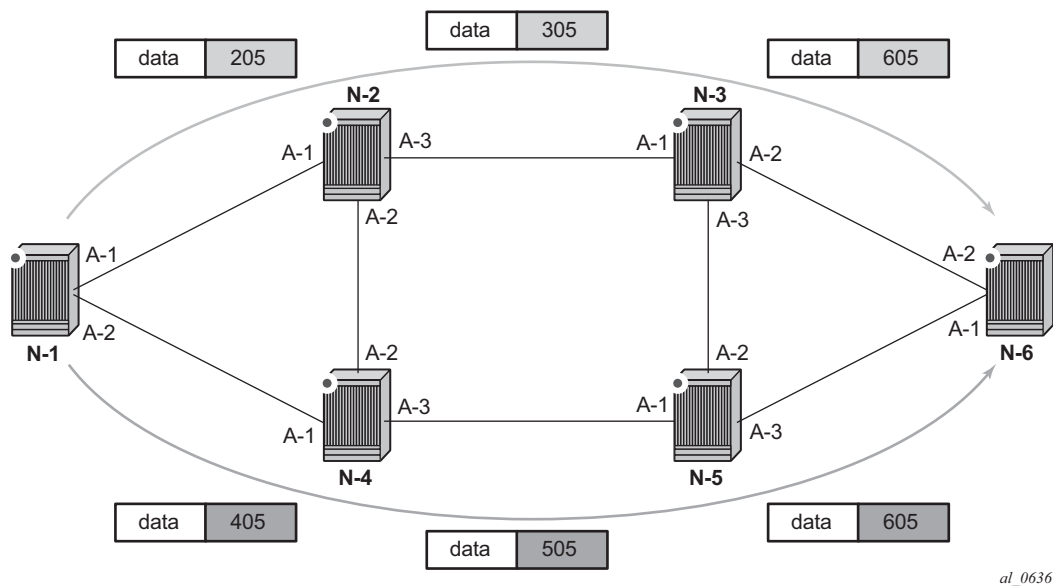


Figure 19: Packet Label Encapsulation using Segment Routing Tunnel

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 a 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:

```
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 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. Note that all resulting net label values (start-label + 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.

Finally, the user brings up segment routing on that IGP instances by un-shutting the context:

```
configure> router>isis>segment-routing>no shutdown
```

```
configure> router>ospf>segment-routing>no shutdown
```

This command will be failed 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 will only program the node SID in the data path towards neighbors which support it.

```
configure> router>isis>advertise-router-capability {area|as}
```

```
configure> router>ospf>advertise-router-capability {link|area|as}
```

Finally, it should be noted that 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 or un-shutting the **segment-routing** node will be failed.

The **segment-routing** command is also mutually exclusive with the **rsvp-shortcut** and **advertise-tunnel-link** options under IGP since a SR tunnel cannot resolve to a RSVP tunnel next-hop.

Next, the user assigns a node SID index or label to the prefix representing the primary address of an IPv4 network interface of type **loopback** using one of the following commands:

- **configure> router>isis>interface>ipv4-node-sid index *value***
- **configure> router>ospf>interface>node-sid index *value***
- **configure> router>isis>interface>ipv4-node-sid label *value***
- **configure> router>ospf>interface>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.

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 since the index and thus label ranges of the various IGP instance are not allowed to overlap.

Segment Routing Operational Procedures

Prefix Advertisement and Resolution

Once segment routing is successfully enabled in the IS-IS or OSPF instance, the router will perform the following operations. See [IS-IS Control Protocol Changes](#) for details of all TLVs and sub-TLVs for each of 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 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. Note the following points:
 - Adjacency SID is advertised for both numbered and unnumbered network IP interface.
 - Adjacency SID for parallel adjacencies between two IGP neighbors is not supported.
 - Adjacency SID will not be advertised for an IES interface since 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 will assign 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 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. A SR tunnel is also added to the TTM. Note that if a node SID resolves over an IES interface, the data path will not be programmed and a trap will be 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 will be dropped.

5. LSA filtering will cause SIDs not to be sent in one direction which means some node SIDs will not be resolved in parts of the network upstream of the advertisement suppression

Note that the SID/Label Binding TLV is supported in receive side and processed. It will however not be generated by the router.

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.

Error and Resource Exhaustion Handling

When the prefix corresponding to a node SID is being resolved, the following procedures are followed:

1. SR OS supports assigning different prefix-SID indices and labels to the same prefix in different IGP instances. While other routers that receive these prefix SIDs will program a single route into RTM, based on the winning instance ID as per RTM route type preference. SR OS will add 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 20](#)), Router D has the same prefix destination, with different SIDs (SIDx and SIDy).

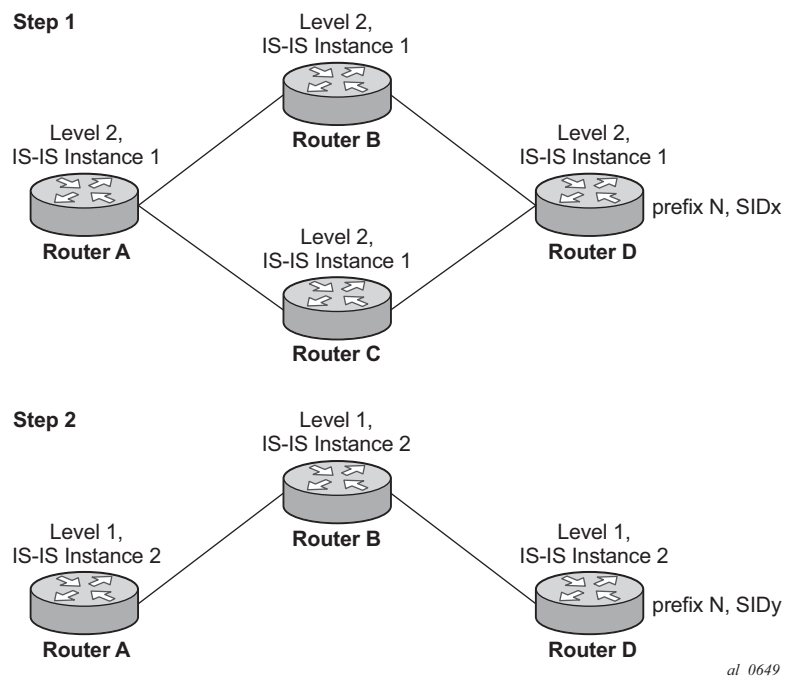


Figure 20: Programming multiple tunnels to the same destination

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 that the TTM tunnel type preference is not used by the SR module. It is put in the TTM and will be used by other applications such a 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 will prefer 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
 - nhop1 = B
 - tunl-pref 10
2. While 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 will resolve only one of them if received from another SR implementation and based on the RTM active route selection.

3. While SR OS does not allow assigning different SID indices or labels to different routes of the same prefix within the same IGP instance, it will resolve only one of them if received from another SR implementation and based on the RTM active route selection.
4. If any of the following conditions are true, the router logs a trap and an 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.
5. Received duplicate prefix-SID index or label for different prefixes within the same IGP instance
 - Program ILM/NHLFE for the first one, log a trap and a syslog error message, and do not program the subsequent one in data path.
6. Received duplicate prefix-SID index for different prefixes across IGP instances
 - In global SID index range mode of operation, the resulting ILM label values will be 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 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.
7. Received duplicate prefix-SID index for the same prefix across IGP instances
 - In global SID index range mode of operation, the resulting ILM label values will be 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 will have different values across the IGP instances. The router programs ILM/NHLFE for each prefix as expected.

The behavior in the case of a global SID index range is illustrated by the IS-IS example in [Figure 21](#).

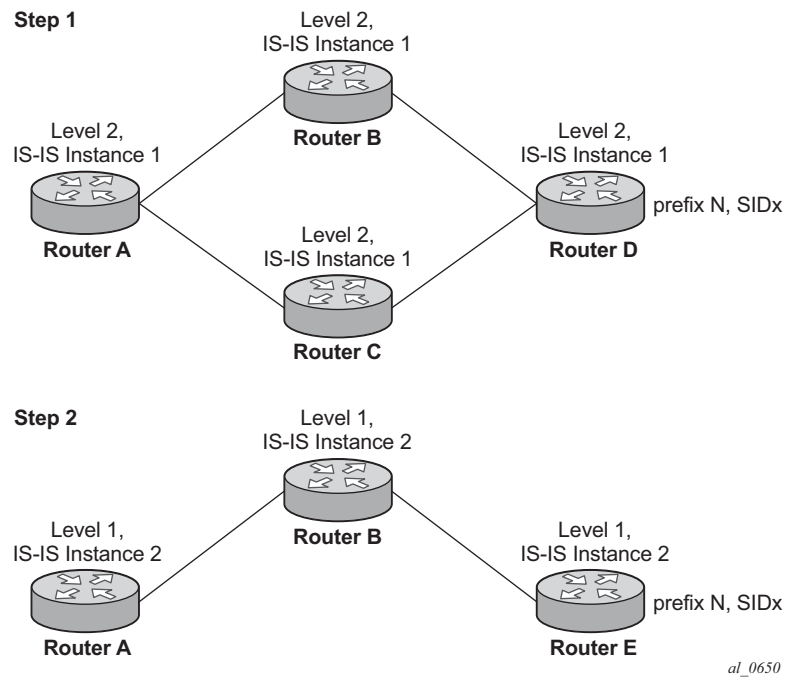


Figure 21: Handling of Same Prefix and SID in different IS-IS Instances

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

Segment Routing in Shortest Path Forwarding

Note that the TTM tunnel type preference is not used by the SR module. It is put in the TTM and will be used by other applications such a 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 will prefer L1 route over L2 route.
 - For prefix N, there is one SR tunnel entry in TTM:
 - SR entry for L2:
 - tunnel-id 2: prefix N-SIDx
 - nhop1 = B
 - rtm-pref 15
 - tunl-pref 10
- SR makes similar decision as RTM is based on rtm-pref 15, which is better than rtm-pref 18, so tunnel-id 2 is chosen.
8. System exhausted ILM resource while assigning an SID index/label to a local loopback interface.
 - Index allocation is failed and an error is returned in CLI. In addition, log a trap and a syslog error message.

9. System exhausted ILM, NHLFE, or any other IOM or CPM resource while resolving and programming a received prefix SID or programming a local adjacency SID.
 - 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. Once IGP is brought back up, it will attempt to program at the next SPF all tunnels which previously failed the programming operation.

Segment Routing Tunnel Management

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 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 will thus be automatically protected with FRR when a 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 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-OSPF and SR-ISIS).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- 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-OSPF (SR-ISIS) is the same regardless if one or more OSPF (IS-IS) 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.

The TTM 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 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. 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-OSPF and SR-ISIS, the user can configure the preference of each specific IGP instance away from the above default values.

- **configure>router>isis>segment-routing>tunnel-table-pref preference <1..255>**
- **configure>router>ospf>segment-routing>tunnel-table-pref preference <1..255>**

The SR tunnel in TTM is available to all users of the TTM: BGP routes, VPRN auto-bind and explicit SDP binding, EVPN MPLS auto-bind, and L2 service with PW template auto-bind and with explicit SDP binding.

Local adjacency SIDs are not programmed into TTM but the remote ones can be used together with a node SID in a tunnel configuration in directed LFA feature.

Tunnel MTU Determination

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. Segment routing however supports remote LFA which programs a LFA backup next-hop that adds another label to the tunnel for a total of two. Finally, directed LFA, if implemented by other routers in the network, can push additional labels but most of the common topologies will not exceed a total of three labels.

Based on the above, the user is provided with a CLI to configure the MTU of all SR tunnels within each IGP instance:

configure> router>isis (ospf)>segment-routing>tunnel-mtu *bytes*

There is no default value for this new command. If the user does not configure a SR tunnel MTU, the MTU will be fully determined by IGP as explained below.

The MTU of the SR tunnel is then determined as follows:

$$SR_Tunnel_MTU = MIN \{Cfg_SR_MTU, IGP_Tunnel_MTU - 3 \text{ labels}\}$$

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 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.

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.

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:

- **configure> router>isis>loopfree-alternate remote-lfa**
- **configure> router>ospf>loopfree-alternate remote-lfa**

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. A repair tunnel can in theory be an RSVP LSP, a LDP-in-LDP tunnel, or a SR tunnel. In SR OS, this feature is restricted to use SR repair tunnel to the remote LFA node.

The remote LFA algorithm for link protection is described in *draft-ietf-rtgwg-remote-lfa*. It 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. Assume the topology in [Figure 22](#).




Figure 22: 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 will have no LFA next-hop due to the ring topology. Node D would loop a packet back to node C if it used it as a backup next-hop. The remote LFA will then run the following algorithm, referred to as the “PQ Algorithm” in *draft-ietf-rtgwg-remote-lfa*:

1. Compute 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 that the *draft-ietf-rtgwg-remote-lfa* initially introduced the concept of P space which would have excluded node F since from node C perspective it has a couple of ECMP paths, one of which goes via link C-B. Since however, the remote LFA next-hop is activated when link C-B fails, it is OK to relax the rule and include node F 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 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 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 22](#).

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 same CLI above is used to configure the maximum IGP cost from node C for a Q node to be eligible.

3. Select 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 22](#). From 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 23](#).

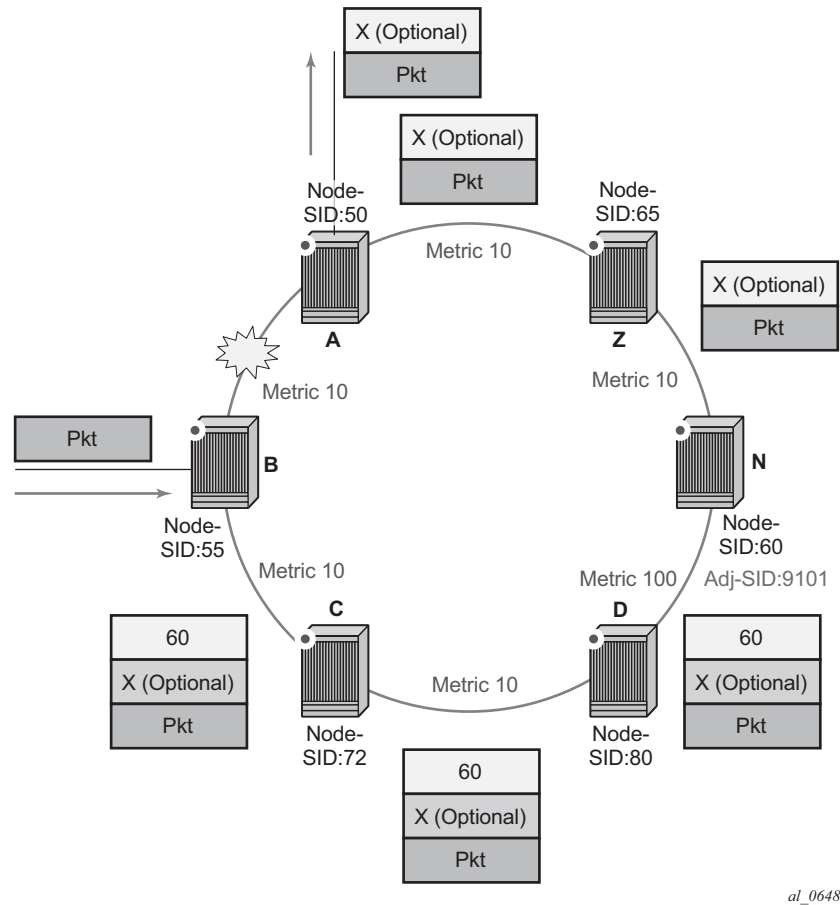


Figure 23: Remote LFA Next-Hop in Segment Routing

Data Path Support

A packet received with a label matching either a node SID or an adjacency SID will be forwarded according to the ILM type and operation, as described in [Table 9](#).

Table 9: Data Path Support

Label type	Operation
Top label is a local node SID	Label is popped and the packet is further processed. If the popped node SID label is the bottom of stack label, the IP packet is looked up and forwarded in the appropriate FIB.
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. 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"> • The BGP label may be popped and the packet looked up in the appropriate FIB. • The BGP label may be swapped to another BGP label. • The BGP label may be stitched to an LDP label.
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 a IP or a service packet over a SR tunnel pushes a maximum of 3 transport labels with a remote LFA nexthop, and 4 transport labels when the P and Q nodes are at most one hop-away from each other (directed LFA if implemented by Node B). This is illustrated in [Figure 24](#).

Page 504



Figure 24: Maximum Pushed Transport Label Stack in Shortest Path Forwarding with Segment Routing

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 a 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 of PE2, followed by the service label for prefix X. When the directed LFA next-hop is activated, node B pushes a couple or more segment routing labels: the node SID for the remote LFA backup node (node D) and the adjacency SID for the next-hop to node N via link D-N.

When node D receives the packet while the directed LFA next-hop is activated, it pops the top segment routing label which corresponds to a local node SID. It then examines the ILM for the next label which corresponds to the adjacency SID. This results in popping this label and forcing the forwarding of the packet over link D-N.

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.

Hash label support

When the hash-label option is enabled in a service context, the insertion of the hash label into the bottom of the label stack of a packet forwarded using segment routing transport tunnel is performed.

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 will be resolved as long as there is a IP Reachability TLV received for same prefix. 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](#).

Segment Routing in Shortest Path Forwarding

I and V flags are set to 1 and 0 respectively when originating the SR capability sub-TLV but 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.

Only IPv4 prefix and adjacency SID sub-TLVs will be originated within MT=0. IPv6 prefix and adjacency SID sub-TLVs can however be received and ignored. The user can get a dump of the octets of the received but not-supported sub-TLVs using the existing show command.

SR OS originates a single prefix SID sub-TLV per IS-IS IP reachability TLV and processes the first one 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 an prefix SID sub-TLV with the following encoding of the flags:

- 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.
- N-flag always set since SROS supports prefix SID of type node SID only.
- P-Flag (no-PHP flag) always set meaning that the label for the prefix SID will be pushed by the PHP router when forwarding to this router. SROS PHP router will process properly a received prefix SID with the P-flag set to zero and will use implicit-null for the outgoing label towards the router which advertised it as long as the P-Flag is also set to 1.
- E-flag (Explicit-Null flag) is always set to zero. An SR OS PHP router will, however, process properly a received prefix SID with the E-flag set to 1 and will use explicit-null for the outgoing label towards the router which advertised it.
- V-flag always set to 0 to indicate an index value for the SID.
- L-flag 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 will still resolve a prefix SID sub-TLV received without the N-flag set but with the prefix length equal to 32. A trap will however be raised by IS-IS.

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.

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 a IP reachability TLV is propagated, along with the prefix SID sub-TLV, by default from L1 to L2 by a L1L2 router. Thus a router in L2 will set up a SR tunnel to the L1L2 router and the L1L2 router will itself set up a SR tunnel to the L1 router which advertised it.

A prefix received in a IP reachability TLV is not propagated, along with the prefix SID sub-TLV, by default from L2 to L1 by a L1L2 router. If the user adds a policy to propagate it, then a router in L1 will set up a SR tunnel to the L1L2 router and the L1L2 router will itself set up a SR tunnel to the L2 router which advertised it.

If a prefix is summarized by an ABR, the prefix SID sub-TLV will not be propagated with the summarized route between levels. If users want to propagate the node SID for a /32 prefix, they must disable route summarization.

SR OS does not propagate into IS-IS the prefix SID sub-TLV of external routes. Thus when the corresponding prefix is redistributed from another protocol such as OSPF, the prefix SID is removed. SROS will however accept the prefix SID sub-TLV of an external route if received from another router and will process it the same way as for an internal route.

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

- F-flag is set to zero to indicate the IPv4 family for the adjacency encapsulation.
- B-Flag not supported and is set to zero.
- V-flag always set to 1.
- L-flag always set to 1.
- 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 will not be processed.
- Weight octet: not supported and is set to all zeros.

SR OS does not originate the SID/Label Binding TLV but can process it properly if received. Note however the following rules and limitations:

- 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 like for the prefix SID sub-TLV, meaning there must be an IP Reachability TLV received for same 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 mapping server Prefix-SID sub-TLV
 - SID received via L2 in a mapping server Prefix-SID sub-TLV
 - 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

- No leaking of this TLV is performed 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 will be used. If that router becomes unreachable, the next reachable one will be used.
- No check is performed if the content of the binding TLVs from different mapping servers are consistent or not.
- Any other sub-TLV, e.g., the SID/Label Sub-TLV, ERO metric and unnumbered interface ID ERO, will be 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.

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:

```
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 will be 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** is specified using the resolution-filter option, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS or OSPF instance.

Refer to the BGP section of the SROS Routing Protocols Guide for more details.

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:


```

configure>router> bgp>next-hop-resolution>
    label-route-transport-tunnel
        [no] family {ipv4, 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 will be 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 ISIS or OSPF instance.

Refer to the BGP section of the SR OS Routing Protocols Guide for more details.

Service Packet Forwarding with Segment Routing

A couple of new SDP sub-types of the MPLS type are added to allow service binding to a 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. The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-isis** and **sr-isis** tunnel types.

The signaling protocol for the service labels for an SDP using a SR tunnel can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

SR tunnels can be used in VPRN with the **auto-bind-tunnel** command:

```

configure>service>vprn>
    auto-bind-tunnel

```

```
resolution {any|disabled|filter}  
resolution-filter  
    [no] gre  
    [no] ldp  
    [no] rsvp  
    sr-isis  
    no sr-isis  
    sr-ospf  
    no sr-ospf  
exit  
exit
```

Both VPN-IPv4 and VPN-IPv6 (6VPE) are supported in a VPRN service using segment routing transport tunnels with the new `auto-bind-tunnel` command.

When the **sr-isis** (**sr-ospf**) value is enabled, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered IS-IS (OSPF) instance.

Refer to the SR OS Layer 3 Services Guide and the BGP section of the SR OS Routing Protocols Guide for more details of the VPRN **auto-bind-tunnel** CLI command.

The following are the service contexts which are supported with SR tunnels:

- VLL, LDP VPLS, IES/VPRN spoke-interface, R-VPLS, BGP EVPN.
- BGP-AD VPLS, BGP-VPLS, BGP VPWS when the `use-provisioned-sdp` option is enabled in the binding to the PW template.
- Intra-AS BGP VPRN for `vpn-ipv4` and `vpn-ipv6` prefixes with both `auto-bind` and `explicit SDP`.
- Multicast over IES/VPRN spoke interface with `spoke-sdp` riding a SR tunnel.

The following service contexts are not supported:

- Inter-AS VPRN.
- Dynamic MS-PW, PW-switching.
- BGP-AD VPLS, BGP-VPLS, BGP VPWS with auto-generation of SDP using a SR tunnel when binding to a PW template.

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.

This also applies to the configuration of the mirror destination for a LI source.

Configuration at mirror source node:

```
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>
```

Notes:

- *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:

```
*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
```

Notes:

- 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 will also be supported with PW redundancy feature when the endpoint spoke-sdp, including the ICB, is using a SR tunnel.

IS-IS Configuration Process Overview

Figure 25 displays the process to provision basic IS-IS parameters.

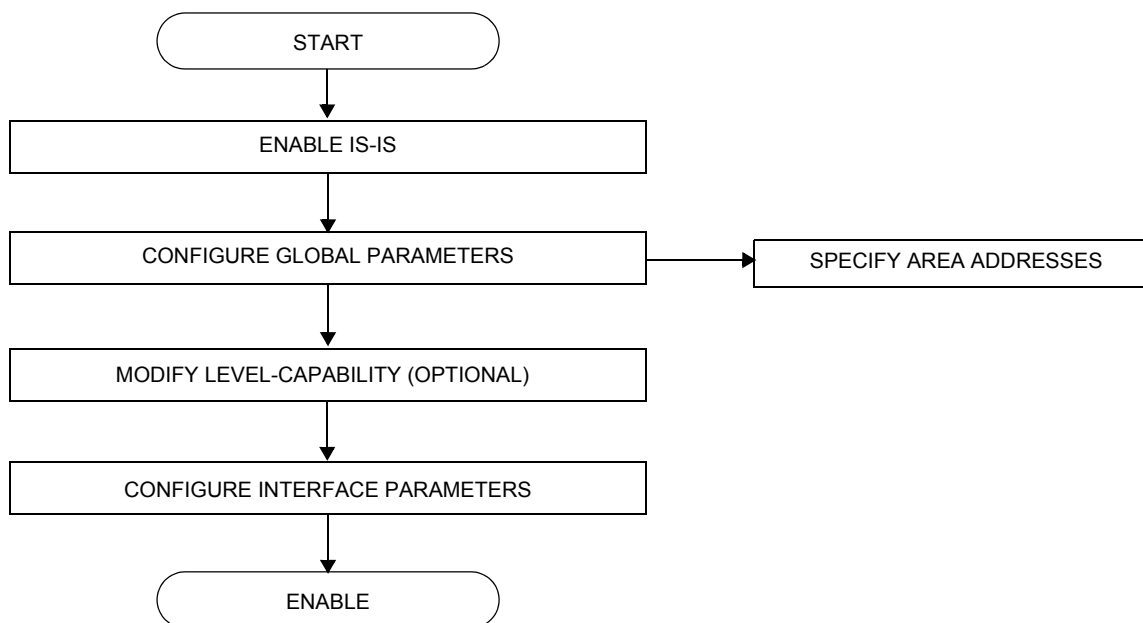


Figure 25: IS-IS Configuration and Implementation Flow

Configuration Notes

This section describes IS-IS configuration caveats.

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.

Configuring IS-IS with CLI

This section provides information to configure intermediate-system-to-intermediate-system (IS-IS) using the command line interface.

Topics in this section include:

- [IS-IS Configuration Overview on page 516](#)
 - [Router Levels on page 516](#)
 - [Area Address Attributes on page 516](#)
 - [Interface Level Capability on page 517](#)
 - [Route Leaking on page 518](#)
- [Basic IS-IS Configuration on page 519](#)
- [Common Configuration Tasks on page 521](#)
 - [Enabling IS-IS on page 522](#)
 - [Modifying Router-Level Parameters on page 522](#)
 - [Configuring ISO Area Addresses on page 524](#)
 - [Configuring Global IS-IS Parameters on page 525](#)
 - [Configuring Interface Parameters on page 526](#)
- [IS-IS Configuration Management Tasks on page 532](#)
 - [Disabling IS-IS on page 532](#)
 - [Modifying Global IS-IS Parameters on page 533](#)
 - [Modifying IS-IS Interface Parameters on page 534](#)
 - [Example: Configuring a Level 1 Area on page 528](#)
 - [Example: Modifying a Router's Level Capability on page 530](#)
 - [Configuring Leaking on page 536](#)
 - [Redistributing External IS-IS Routers on page 540](#)
 - [Specifying MAC Addresses for All IS-IS Routers on page 541](#)

IS-IS Configuration Overview

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.

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.

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.
2. System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.
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

Interface Level Capability

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 10](#) displays configuration combinations and the potential adjacencies that can be formed.

Table 10: 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

Route Leaking

Alcatel-Lucent'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.

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 address(es)
- Configure IS-IS interfaces.

The following output displays IS-IS default values.

```
A:Dut-A>config>router>isis$ info detail
```

```
-----
level-capability level-1/2
no graceful-restart
area-id 01
no authentication-key
no authentication-type
authentication-check
csnp-authentication
lsp-lifetime 1200
no export
hello-authentication
psnp-authentication
traffic-engineering
no reference-bandwidth
no disable-ldp-sync
ipv4-routing
spf-wait 10 1000 1000
no strict-adjacency-check
lsp-wait 5 0 1
level 1
    no authentication-key
    no authentication-type
    csnp-authentication
    external-preference 160
    hello-authentication
    preference 15
    psnp-authentication
    no wide-metrics-only
exit
level 2
    no authentication-key
    no authentication-type
    csnp-authentication
    external-preference 165
    hello-authentication
```

Basic IS-IS Configuration

```
        preference 18
        psnp-authentication
        no wide-metrics-only
    exit
    no shutdown
-----
A:Dut-A>config>router>isis$
```

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 address(es).
- Configure IS-IS interface-specific parameters.

Configuring IS-IS Components

Use the CLI syntax displayed below for:

- [Enabling IS-IS on page 522](#)
 - [Modifying Router-Level Parameters on page 522](#)
 - [Configuring ISO Area Addresses on page 524](#)
 - [Configuring Global IS-IS Parameters on page 525](#)
 - [Configuring Interface Parameters on page 526](#)
 - [Example: Configuring a Level 1 Area on page 528](#)
 - [Example: Modifying a Router's Level Capability on page 530](#)
-

Enabling IS-IS

IS-IS must be enabled in order for the protocol to be active.

NOTE: 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.

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#
```

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#
```


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#
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#
```

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:

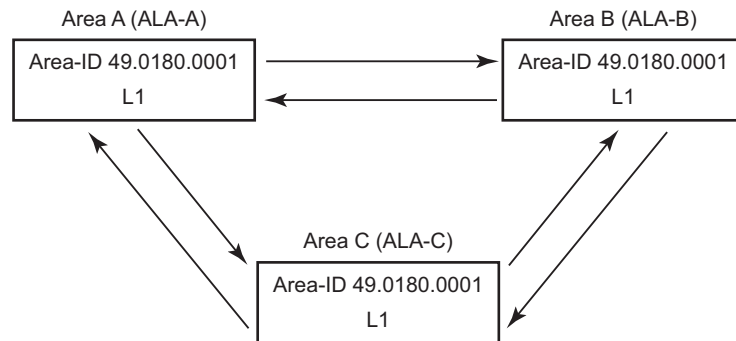
```
config>router# isis
config>router>isis# level 1
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# level 2
config>router>isis>level# wide-metrics-only
config>router>isis>level# exit
config>router>isis# interface ALA-1-2
config>router>isis>if# level-capability level-2
config>router>isis>if# mesh-group 85
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# mesh-group 101
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# mesh-group 85
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# level-capability level-1/2
>router>isis>if# mesh-group 101
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
        mesh-group 85
    exit
    interface "ALA-1-3"
        level-capability level-1
        interface-type point-to-point
        mesh-group 101
    exit
    interface "ALA-1-5"
        level-capability level-1
        interface-type point-to-point
        mesh-group 85
    exit
    interface "to-103"
        mesh-group 101
    exit
-----
A:ALA-A>config>router>isis#
```

Example: Configuring a Level 1 Area

NOTE: Interfaces are configured in the `config>router>interface` context.



OSRG031

Figure 26: Configuring a Level 1 Area

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
-----

```

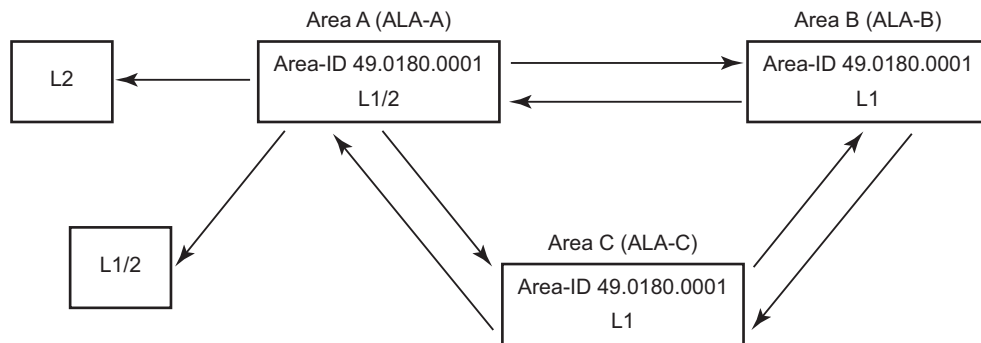
```

A:ALA-C>config>router>isis#

```

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. 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.



OSRG036

Figure 27: Configuring a Level 1/2 Area

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
```

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.

IS-IS Configuration Management Tasks

This section discusses the following IS-IS configuration management tasks:

- [Disabling IS-IS on page 532](#)
 - [Removing IS-IS on page 532](#)
 - [Modifying Global IS-IS Parameters on page 533](#)
 - [Modifying IS-IS Interface Parameters on page 534](#)
 - [Configuring Authentication using Keychains on page 535](#)
 - [Configuring Leaking on page 536](#)
 - [Redistributing External IS-IS Routers on page 540](#)
 - [Specifying MAC Addresses for All IS-IS Routers on page 541](#)
-

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`

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`

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 raiderslost
```

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
    mesh-group 85
exit
interface "ALA-1-3"
    level-capability level-1
    interface-type point-to-point
    mesh-group 101
exit
interface "ALA-1-5"
    level-capability level-1
    interface-type point-to-point
    mesh-group 85
exit
interface "to-103"
    mesh-group 101
exit
interface "A-B"
exit
interface "A-C"
exit
-----
A:ALA-A>config>router>isis#
```

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 on page 528](#) and [Example: Modifying a Router's Level Capability on page 530](#).

Example:

```
config>router# isis
config>router>isis# interface ALA-1-3
config>router>isis>if# mesh-group 85
config>router>isis>if# passive
config>router>isis>if# lsp-pacing-interval 5000
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 49ersrule
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
    mesh-group 85
exit
interface "ALA-1-3"
    level-capability level-1
    interface-type point-to-point
    lsp-pacing-interval 5000
    mesh-group 85
    passive
```

```

exit
interface "ALA-1-5"
    level-capability level-1
    interface-type point-to-point
    mesh-group 85
exit
interface "to-103"
    hello-authentication-key "DvR3l264KQ6vXMTvbAZ1mE" hash
    hello-authentication-type message-digest
    mesh-group 101
exit
interface "A-B"
exit
-----
A:ALA-A>config>router>isis#

```

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
- HMAC-MD5
- HMAC-SHA-1

Configuring Leaking

- HMAC-SHA-256

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.

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
-----
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"
...

```

Configuring Leaking

```
-----  
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 on page 540](#)). 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#
```

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. Refer to the *Route Policy* section of this manual 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#
```


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.

IS-IS Command Reference

Command Hierarchies

Configuration Commands

- [Global Commands on page 543](#)
- [Interface Commands on page 545](#)
- [Show Commands on page 547](#)
- [Clear Commands on page 547](#)
- [Debug Commands on page 547](#)

```

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
— [no] authentication-check
— authentication-key [authentication-key | hash-key] [hash | hash2]
— no authentication-key
— authentication-type {password | message-digest}
— no authentication-type
— auth-keychain keychain-name
— no auth-keychain
— [no] csnp-authentication
— default-route-tag tag
— no default-route-tag
— [no] disable-ldp-sync
— export policy-name [.. policy-name]
— no export
— export-limit number [log percentage]
— no export-limit
— [no] graceful-restart
— [no] helper-disable
— [no] hello-authentication
— [no] hello-padding {adaptive | loose | strict}
— [no] ignore-lsp-errors
— [no] iid-tlv-enable
— [no] interface ip-init-name

```

```

— [no] bfd-enable {ipv4 } [include-bfd-tlv]
— [no] default-instance
— hello-authentication-key [authentication-key | hash-key][hash | hash2]
— no hello-authentication-key
— hello-authentication-type {password | message-digest}
— no hello-authentication-type
— level
    — hello-interval seconds
    — no hello-interval
    — hello-multiplier multiplier
    — no hello-multiplier
    — [no] if-topology mt-id
— interface-type {broadcast | point-to-point}
— no interface-type
— ipv4-node-sid index value
— ipv4-node-sid label value
— no ipv4-node-sid
— ipv4-multicast-routing {native | mt}
— [no] ipv4-multicast-routing
— ipv6-multicast-routing {native | mt}
— [no] ipv6-multicast-routing
— [no] ipv4-routing
— [no] ldp-over-rsvp
— level {1 | 2}
    — authentication-key [authentication-key | hash-key] [hash | hash2]
    — no authentication-key
    — authentication-type {password | message-digest}
    — no authentication-type
    — auth-keychain keychain-name
    — no auth-keychain
    — [no] csnp-authentication
    — default-ipv4-multicast-metric metric
    — default-ipv6-multicast-metric 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
    — [no] loopfree-alternate-exclude
    — 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
    — 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 [0-8]
— no oper-members
— revert-members [0-8]
— no revert-members
— loopfree-alternate [remote-lfa [max-pq-cost value]]
— no loopfree-alternate
— loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy... up to 5]
— lsp-lifetime seconds
— no lsp-lifetime
— lsp-mtu-size size
— no lsp-mtu-size-size
— lsp-refresh-interval seconds
— no lsp-refresh-interval
— [no] lsp-wait lsp-wait [lsp-initial-wait [lsp-second-wait]]
— [no] multi-topology
—   [no] ipv4-multicast
—   [no] ipv6-multicast
—   topology mt-id rtm rtm-id | rtm-name
—   no topology mt-id
— [no] multicast-import
— overload [timeout seconds]
— no overload
— overload-on-boot [timeout seconds]
— no overload-on-boot
— [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
— [no] rsvp-shortcut
— segment-routing
— no segment-routing
—   prefix-sid-range {global | start-label label-value max-index index-value}
—   no prefix-sid-range
—   tunnel-mtu bytes
—   no tunnel-mtu
—   tunnel-table-pref preference
—   no tunnel-table-pref
—   [no] shutdown
— [no] shutdown
— [no] spf-wait spf-wait [spf-initial-wait [spf-second-wait]]
— [no] suppress-default
— summary-address {ip-prefix/mask | ip-prefix [netmask]} level [tag tag]
— no summary-address {ip-prefix/mask | ip-prefix [netmask]}
— [no] traffic-engineering
— [no] unicast-import-disable
— [no] interface-type ip-int-name
—   [no] default-route-tag {ipv4} [include-bfd-tlv]
—   csnp-interval seconds

```

```

— no csnp-interval
— hello-authentication-key [authentication-key | hash-key][hash | hash2]
— no hello-authentication-key
— hello-authentication-type {password | message-digest}
— no hello-authentication-type
— hello-auth-keychain keychain-name
— no hello-auth-keychain
— interface-type {broadcast | point-to-point}
— no interface-type
— no loopfree-alternate-exclude
— [no] metric
— level {1 | 2}
    — hello-authentication-key [authentication-key | hash-key] [hash |
      hash2]
    — no hello-authentication-key
    — hello-authentication-type [password | message-digest]
    — no hello-authentication-type
    — hello-auth-keychain keychain-name
    — no hello-auth-keychain
    — hello-interval seconds
    — no hello-interval
    — hello-multiplier multiplier
    — no hello-multiplier
    — if-topology mt-id
    — no if-topology mt-id
    — [no] ldp-over-rsvp
    — metric metric
    — no metric
    — [no] passive
    — priority number
    — no priority
    — sd-offset offset-value
    — no sd-offset
    — sf-offset
    — no sf-offset
— level-capability {level-1 | level-2 | level-1/2}
— loopfree-alternate route-nh-template template-name
— no lfa-policy-map
— [no] loopfree-alternate
— [no] loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy... up
  to 5]
— lsp-pacing-interval milli-seconds
— no lsp-pacing-interval
— mesh-group [value | blocked]
— no mesh-group
— [no] passive
— retransmit-interval seconds
— no retransmit-interval
— [no] shutdown
— tag tag
— no tag

```

Show Commands

```

show
  — router
    — isis [isis-instance]
      — adjacency [ip-address | ip-int-name | nbr-system-id] [detail]
      — capabilities [system-id | lsp-id ] [level level]
      — database [system-id | lsp-id ] [detail] [level level]
      — hostname
      — interface [ip-int-name | ip-address] [detail]
      — lfa-coverage
      — neighbor
      — 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]
      — routes [ipv4-unicast | ipv4-multicast | ipv6-multicast | mt mt-id-number] [ip-prefix[/prefix-length]] [alternative] [exclude-shortcut]
      — spf [detail] [lfa]
      — spf-log [detail]
      — statistics
      — status
      — summary-address [ip-address [/mask] ]
      — topology [ipv4-unicast | ipv6-unicast | ipv4-multicast | ipv6-multicast | mt mt-id-number] [lfa] [detail]

```

Clear Commands

```

clear
  — router
    — isis [isis-instance]
      — adjacency [system-id]
      — database [system-id]
      — export
      — spf-log
      — statistics

```

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]

```

Command Hierarchies

- **rtm** [*ip-address*]
- **no rtm**
- [**no**] **spf** [*level-number*] [*system-id*]

IS-IS Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>router>isis
 config>router>isis>interface *ip-int-name*
 config>router>isis>if>level *level-number*
 config>router>isis>if>segment-routing

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.

Special Cases **IS-IS Global** — In the **config>router>isis** context, the **shutdown** command disables the IS-IS protocol instance. By default, the protocol is enabled, **no shutdown**.

IS-IS Interface — In the **config>router>isis>interface** context, the command disables the IS-IS interface. By default, the IS-IS interface is enabled, **no shutdown**.

IS-IS Interface and Level — In the **config>router>isis>interface *ip-int-name*>level** context, the command disables the IS-IS interface for the level. By default, the IS-IS interface at the level is enabled, **no shutdown**.

Default **no shutdown** — IS-IS entity is administratively enabled.

IS-IS Commands

isis

Syntax	[no] no isis <i>[isis-instance]</i>				
Context	config>router				
Description	<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 no shutdown command in the config>router>isis context. Alternatively, the IS-IS protocol instance is disabled with the shutdown command in the config>router>isis context.</p> <p>The no form of the command deletes the IS-IS protocol instance. Deleting the protocol instance removes all configuration parameters for this IS-IS instance.</p>				
Parameters	<i>isis-instance</i> — Specifies the instance ID for an IS-IS instance. <table> <tr> <td>Values</td><td>1–31</td></tr> <tr> <td>Default</td><td>0</td></tr> </table>	Values	1–31	Default	0
Values	1–31				
Default	0				

tag

Syntax	tag <i>tag</i> no tag
Context	config>router>isis>interface
Description	This command configures a route tag to the specified IP address of an interface.
Parameters	<i>tag</i> — [1..4294967295]

all-l1isis

Syntax	all-l1isis <i>ieee-address</i> no all-l1isis
Context	config>router>isis
Description	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. You should shut/no shut the IS-IS instance to make the change operational.
Default	all-l1isis 01-80-C2-00-01-00

Parameters *ieee-address* — Specifies the destination MAC address for all L1 I-IS neighbors on the link for this ISIS instance.

all-l2isis

Syntax **all-l2isis** *ieee-address*
no all-l2isis

Context config>router>isis

Description This command enables you to specify the MAC address to use for all L2 IS-IS routers. The MAC address should be a multicast address. You should shut/no shut the IS-IS instance to make the change operational.

Default all-l2isis 01-80-C2-00-02-11

Parameters *ieee-address* — Specifies the destination MAC address for all L2 ISIS neighbors on the link for this ISIS instance.

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** — Rejects authentication mismatches.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>router>isis
 config>router>isis>level *level-number*

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.

Default **no authentication-key** — No authentication key is configured.

Parameters *authentication-key* — 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 (“ ”).

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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax **authentication-type {password | message-digest}**
no authentication

Context config>router>isis
config>router>isis>level *level-number*

Description This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.

Both the authentication key and the authentication type on a segment must match. The **authentication-key** statement must also be included.

Configure the authentication type on the global level in the **config>router>isis** context.

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.

Default **no authentication-type** — No authentication type is configured and authentication is disabled.

Parameters **password** — Specifies that simple password (plain text) authentication is required.

message-digest — Specifies that MD5 authentication in accordance with RFC2104 is required.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>router>isis> config>router>isis>level config>service>vpn>isis> config>service>vpn>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	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

hello-auth-keychain

Syntax	hello-auth-keychain <i>name</i>
Context	config>router>isis config>router>isis>level config>service>vpn>isis>interface config>service>vpn>isis>interface>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 hello-auth-keychain
Parameters	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

default-route-tag

Syntax	default-route-tag <i>tag</i> no default-route-tag
Context	config>router>isis
Description	This command configures the route tag for default route.

IS-IS Commands

Parameters	<i>tag</i> — <i>tag</i> — Assigns a default tag
	Values Accepts decimal or hex formats: ISIS: [0x0..0xFFFFFFFF]H
	Values 1 — 4294967295

csnp-authentication

Syntax	[no] csnp-authentication
Context	config>router>isis config>router>isis>level <i>level-number</i>
Description	This command enables authentication of individual ISIS packets of complete sequence number PDUs (CSNP) type. The no form of the command suppresses authentication of CSNP packets.

csnp-interval

Syntax	csnp-interval seconds no csnp-interval
Context	config>router>isis>interface <i>ip-int-name</i>
Description	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. The no form of the command reverts to the default value.
Default	csnp-interval 10 — CSN PDUs are sent every 10 seconds for LAN interfaces. csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.
Parameters	<i>seconds</i> — The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer. Values 1 — 65535

link-group

Syntax	link-group link-group-name no link-group
Context	config>router>isis
Description	This command specifies the ISIS link group associated with this particular level of the interface.
Parameters	<i>link-group-name</i> — Specifies an ISIS link group on the system up to 32 characters in length.

description

Syntax	description <i>string</i> 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.
Parameters	<i>string</i> — Character string to be associated with the associated link-group.

member

Syntax	[no] member <i>interface-name</i>
Context	config>router>isis>link-group
Description	This command adds or removes a links 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	<i>interface-name</i> — Name of the interface to be added or removed from the associated link-group.

oper-members

Syntax	oper-members [0-8] no oper-members
Context	config>router>isis>link-group
Description	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. The no form of the command reverts the oper-members limit to 1.
Default	oper-members 0

revert-members

Syntax	revert-members [0-8] no revert-members
Context	config>router>isis>link-group
Description	<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 no form of the command reverts the revert-members threshold back to the default which is equal to the oper-member threshold value.</p>
Default	revert-members <i>oper-members</i>

ipv4-unicast-metric-offset

Syntax	ipv4-unicast-metric-offset <i>offset-value</i> no ipv4-unicast-metric-offset
Context	config>router>isis>link-group
Description	<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 no form of the command reverts the offset value to 0.</p>
Default	no ipv4-unicast-metric-offset
Parameters	<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>Values 0 — 6777215</p>

ipv6-unicast-metric-offset

Syntax	ipv6-unicast-metric-offset <i>offset-value</i> no ipv6-unicast-metric-offset
Context	config>router>isis>link-group
Description	<p>This command sets the offset value for the IPv6 unicast 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 topology.</p> <p>The no form of the command reverts the offset value to 0.</p>
Default	no ipv6-unicast-metric-offset

Parameters	<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.
Values	0 — 6777215

ipv4-multicast-metric-offset

Syntax	ipv4-multicast-metric-offset <i>offset-value</i> no ipv4-multicast-metric-offset
Context	config>router>isis>link-group
Description	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. The no form of the command reverts the offset value to 0.
Default	no ipv4-multicast-metric-offset
Parameters	<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
Values	0 — 6777215

ipv6-multicast-metric-offset

Syntax	ipv6-multicast-metric-offset <i>offset-value</i> no ipv6-multicast-metric-offset
Context	config>router>isis>link-group
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	<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
Values	0 — 6777215

default-metric

Syntax	default-metric <i>ipv4 metric</i> no default-metric
Context	config>router>isis>level
Description	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.
Default	10 <i>ipv4 metric</i> — Specifies the default metric for IPv4 unicast.
Values	1 — 16777215

default-ipv4-multicast-metric

Syntax	default-ipv4-multicast-metric <i>metric</i> no default-ipv4-multicast-metric
Context	config>router>isis>level
Description	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.
Default	10
Parameters	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT3)
Values	1 — 16777215

default-ipv6-multicast-metric

Syntax	default-ipv6-multicast-metric <i>metric</i> no default-ipv6-multicast-metric
Context	config>router>isis>level
Description	This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4). The no form of this command deletes the specified default metric and reverts to using the system default of 10.
Default	10
Parameters	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4).

1 — 16777215

default-ipv6-unicast-metric

Syntax	default-ipv6-unicast-metric <i>ipv6 metric</i> no default-ipv6-unicast-metric
Context	config>router>isis>level
Description	This command specifies the default metric for IPv6 unicast.
Default	no default-ipv6-unicast-metric
Parameters	<i>ipv6-metric</i> — Specifies the default metric for IPv6 unicast.
Values	1 — 16777215

disable-ldp-sync

Syntax	[no] disable-ldp-sync
Context	config>router>isis
Description	<p>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 advertized cost is different. It will then disable IGP-LDP synchornization 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.</p> <p>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.</p>
Default	no disable-ldp-sync

export

Syntax	[no] export <i>policy-name</i> [<i>policy-name</i> ...up to 5 max]
Context	config>router>isis
Description	<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>

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.

If an **aggregate** command is also configured in the **config>router** context, then the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of the command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Default **no export** — No export policy name is specified.

Parameters *policy-name* — The export policy name. Up to five *policy-name* arguments can be specified.

export-limit

Syntax **export-limit** *number* [**log** *percentage*]
no export-limit

Context config>router>isis
config>service>vpn>isis

Description 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.
The **no** form of the command removes the parameters from the configuration.

Default no export-limit, the export limit for routes or prefixes is disabled.

Parameters *number* — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 — 4294967295

log percentage — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 — 100

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>router>isis>level *level-number*

Description This command configures the external route preference for the IS-IS level.

The **external-preference** 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.

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.

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 **ecmp** in the **config>router** context.

Default Default preferences are listed in the following table:

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
IS-IS Level 2 external	165	Yes
TMS	167	No

*. Internal preferences are changed using the **preference** command in the **config>router>isis>level level-number** context

Parameters *preference* — The preference for external routes at this level as expressed.

Values 1 — 255

graceful-restart

Syntax **[no] graceful-restart**

Context config>router>isis
config>service>vprn>isis

Description This command enables graceful-restart helper support for ISIS. The router will act as a helper to neighbors who are graceful-restart-capable and are restarting.

When the control plane of a graceful-restart-capable router fails, the neighboring routers (graceful-restart helpers) temporarily preserve adjacency information so packets continue to be forwarded through the failed graceful-restart router using the last known routes. If the control plane of the graceful-restart router comes back up within the timer limits, then the routing protocols re-converge to minimize service interruption.

IS-IS Commands

The **no** form of the command disables graceful restart and removes all graceful restart configurations in the ISIS instance.

Default disabled

helper-disable

Syntax [no] helper-disable

Context config>router>isis>graceful-restart
config>service>vpn>isis>graceful-restart

Description This command disables the helper support for graceful restart.

When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the router will not help the neighbors to restart).

The **no helper-disable** command enables helper support and is the default when graceful-restart is enabled.

Default disabled

hello-authentication

Syntax [no] hello-authentication

Context config>router>isis
config>router>isis>level *level-number*
config>service>vpn>isis
config>service>vpn>isis>interface
config>service>vpn>isis>level

Description This command enables authentication of individual ISIS packets of HELLO type.

The **no** form of the command suppresses authentication of HELLO packets.

hello-padding

Syntax [no] hello-padding {adaptive | loose | strict}

Context config>router>isis
config>service>vpn>isis

Description 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 "lsp-mtu-size" command.

The **no** form of the command disables IS-IS hello padding.

Default	no hello-padding — hello padding is not configured
Parameters	<p>adaptive — 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"> 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. 2. broadcast interface—Padding starts until at least one adjacency is up on the interface. <p>loose — 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"> 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. 2. broadcast interface—Padding starts until there is at least one adjacency (broadcast only has up/down) is up on the interface. <p>strict — 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"> 1. point-to-point interface—Padding is done for all adjacency states, and is continuous. 2. broadcast interface—Padding is done for all adjacency states, and is continuous.

ignore-lsp-errors

Syntax	[no] ignore-lsp-errors
Context	<pre>config>router>isis config>service>vpn>isis</pre>
Description	<p>This command specifies that ISIS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged.</p> <p>The no form of the command specifies that ISIS will not ignore LSP errors.</p>

iid-tlv-enable

Syntax	[no] iid-tlv-enable
Context	<pre>config>router>isis</pre>
Description	<p>This command specifies whether Instance Identifier (IID) TLV has been enabled or disabled for this ISIS instance.</p> <p>When enabled, each I-IS instance marks its packets with the IID TLV containing its unique 16-bit IID for the routing domain. You should shut/no shut the isis instance to make the change operational.</p>

IS-IS Commands

Default no iid-tlv-enable

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.

Default **no interface** — No IS-IS interfaces are defined.

Parameters *ip-int-name* — Identify the IP interface name created in the **config>router>interface** context. The IP interface name must already exist.

bfd-enable

Syntax [no] **bfd-enable** {**ipv4** } [*include-bfd-tlv*]

Context config>router>isis>interface

Description 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.

The **no** form of this command removes BFD from the associated adjacency.

Default no bfd-enable ipv4

Parameters *include-bfd-tlv* — 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.

default-instance

Syntax	[no] default-instance
Context	config>router>isis>interface <i>ip-int-name</i> config>service>vpn>isis>interface <i>ip-int-name</i>
Description	<p>This command enables a non-MI capable router to establish an adjacency and operate with an SR OS 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>The no form of this command disables the functionality so that the router can only establish adjacencies in the standard instance 0.</p>
Default	no default-instance

hello-authentication-key

Syntax	hello-authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no hello-authentication-key
Context	config>router>isis>interface <i>ip-int-name</i> config>router>isis>if>level <i>level-number</i> config>service>vpn>isis>interface config>service>vpn>isis>level
Description	<p>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.</p> <p>To configure the hello authentication key in the interface context use the hello-authentication-key in the config>router>isis>interface context.</p> <p>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.</p> <p>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.</p> <p>When the hello authentication key is configured in the config>router>isis>interface context, it applies to all levels configured for the interface.</p> <p>The no form of the command removes the authentication-key from the configuration.</p>
Default	no hello-authentication-key — No hello authentication key is configured.

- 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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.
 - hash2** — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

hello-authentication-type

- Syntax** **hello-authentication-type {password | message-digest}**
no hello-authentication-type
- Context** config>router>isis>interface *ip-int-name*
config>router>isis>if>level *level-number*
config>service>vpn>isis>interface
config>service>vpn>isis>level
- Description** 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 **authentication-key** statement must also be included.

To configure the hello authentication type at the interface context, use **hello-authentication-type** in the **config>router>isis>interface** context.

To configure or override the hello authentication setting for a given level, configure the **hello-authentication-type** in the **config>router>isis>interface>level** context.

The **no** form of the command disables hello authentication.
- Default** **no hello-authentication-type** — Hello authentication is disabled.
- Parameters** **password** — Specifies simple password (plain text) authentication is required.

message-digest — Specifies MD5 authentication in accordance with RFC2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

hello-interval

Syntax	hello-interval <i>seconds</i> no hello-interval
Context	config>router>isis>if>level <i>level-number</i>
Description	This command configures the interval in seconds between hello messages issued on this interface at this level. The no form of the command to reverts to the default value.
Default	3 — Hello interval default for the designated intersystem. 9 — Hello interval default for non-designated intersystems.
Parameters	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer. Values 1 — 20000

hello-multiplier

Syntax	hello-multiplier <i>multiplier</i> no hello-multiplier
Context	config>router>isis>if>level <i>level-number</i>
Description	This command configures the number of missing hello PDUs from a neighbor after the router declares the adjacency down. The no form of the command reverts to the default value.
Default	3 — The router can miss up to 3 hello messages before declaring the adjacency down.
Parameters	<i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer. Values 2 — 100

if-topology

Syntax	if-topology <i>mt-id</i> no if-topology <i>mt-id</i>
Context	config>router>is-is>interface>level
Description	This command links the associated interface with the specified IS-IS topology.

IS-IS Commands

By default all IS-IS interfaces should be associated the respective unicast topology. To exclude an interface from the respective unicast topology use the command `no if-topology <0|2>`.

The **no** form of this command deletes the specified IS-IS topology.

Default No default

Parameters *mt-id* — Specify the topology to be created.

Values 3 or 4

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>router>isis>interface *ip-int-name*

Description This command configures the IS-IS interface type as either broadcast or point-to-point.

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.

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.

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

Special Cases **SONET** — Interfaces on SONET channels default to the point-to-point type.

Ethernet or Unknown — Physical interfaces that are Ethernet or unknown default to the broadcast type.

Default **point-to-point** — For IP interfaces on SONET channels.

broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.

Parameters **broadcast** — Configures the interface to maintain this link as a broadcast network.

point-to-point — Configures the interface to maintain this link as a point-to-point link.

ipv4-node-sid

Syntax **ipv4-node-sid index** *value*
ipv4-node-sid label *value*
no ipv4-node-sid

Context config>router>isis>interface

Description 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.

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. Also, assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.

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.

Parameters	index <i>value</i> — integer.
	Values 0 — 4294967295
	Default none
	label <i>value</i> — integer.
	Values 0 — 4294967295
	Default none

ipv4-multicast-routing

Syntax	ipv4-multicast-routing {native mt} [no] ipv4-multicast-routing
Context	config>router>isis
Description	<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 no 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>
Default	ipv4-multicast-routing native
Parameters	<p>native — Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.</p> <p>mt — Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.</p>

ipv6-multicast-routing

Syntax	ipv6-multicast-routing {native mt} [no] ipv6-multicast-routing
Context	config>router>isis
Description	<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>

IS-IS Commands

The **no** 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.

Default ipv6-multicast-routing native

Parameters **native** — Causes IPv6 routes from the MT0 topology to be added to the multicast RTM for RPF checks.
mt — Causes IPv6 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

ipv4-routing

Syntax [**no**] **ipv4-routing**

Context config>router>isis

Description This command specifies whether this IS-IS instance supports IPv4.
The **no** form of the command disables IPv4 on the IS-IS instance.

Default ipv4-routing

ldp-over-rsvp

Syntax [**no**] **ldp-over-rsvp**

Context config>router>isis

Description This command allows LDP over RSVP processing in IS-IS.
The **no** form of the command disables LDP over RSVP processing.

Default no ldp-over-rsvp

level

Syntax **level {1 | 2}**

Context config>router>isis
config>router>isis>interface *ip-int-name*

Description 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:

```
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
```

Special Cases **Global IS-IS Level** — The **config>router>isis** context configures default global parameters for both Level 1 and Level 2 interfaces.

IS-IS Interface Level — The **config>router>isis>interface** 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.

Default level 1 or level 2

Parameters 1 — Specifies the ISIS operational characteristics of the interface at level 1.
2 — Specifies the ISIS operational characteristics of the interface at level 2.

level-capability

Syntax **level-capability {level-1 | level-2 | level-1/2}**
no level-capability

Context config>router>isis
config>router>isis>interface *ip-int-name*

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 11 displays configuration combinations and the potential adjacencies that can be formed.

Table 11: 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

Table 11: Potential Adjacency (Continued)

Global Level	Interface Level	Potential Adjacency
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.

Special Cases **IS-IS Router** — In the **config>router>isis** context, changing the **level-capability** performs a restart on the IS-IS protocol instance.

IS-IS Interface — In the **config>router>isis>interface** context, changing the **level-capability** performs a restart of IS-IS on the interface.

Default **level-1/2**

Parameters **level-1** — Specifies the router/interface can operate at Level 1 only.

level-2 — Specifies the router/interface can operate at Level 2 only.

level-1/2 — Specifies the router/interface can operate at both Level 1 and Level 2.

loopfree-alternate

Syntax **loopfree-alternate [remote-lfa [max-pq-cost *value*]]**
no loopfree-alternate

Context **config>router>isis**
config>service>vpn>isis

Description This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol or under the IS-IS VPRN.

When this command is enabled, it instructs 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 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 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 no form of this command disables the LFA computation by IGP SPF.

Default no loopfree-alternate

Parameters **max-pq-cost** *value* — integer used to limit the search of candidate P and Q nodes in remote LFA by setting the maximum IGP cost from the router performing remote LFA calculation to the candidate P or Q node.

Values 1 — 4294967295

Default none

loopfree-alternate-exclude

Syntax **loopfree-alternate-exclude prefix-policy** *prefix-policy* [*prefix-policy...* up to 5]
no loopfree-alternate-exclude

Context config>router>isis
config>service>vpn>isis

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 that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **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 that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form deletes the exclude prefix policy.

Parameters **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

lfa-policy-map

Syntax **lfa-policy-map route-nh-template** *template-name*
no lfa-policy-map

IS-IS Commands

Context	config>router>isis>interface
Description	<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 loopfree-alternate-exclude, 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 no form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.</p>
Parameters	<i>template-name</i> — Specifies the name of the template, up to 32 characters.

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate
Context	configure>router>isis>level configure>router>isis>interface configure>service>vpn>isis>level configure>service>vpn>isis>interface
Description	<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 no form of this command re-instates the default value for this command.</p>
Default	no loopfree-alternate-exclude

lsp-pacing-interval

Syntax	lsp-pacing-interval <i>milliseconds</i> no lsp-pacing-interval
Context	config>router>isis>interface
Description	This command configures the interval between LSP PDUs sent from this interface.

To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSPs). If a value of zero is configured, no LSPs are sent from the interface.

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

Default **100** — LSPs are sent in 100 millisecond intervals.

Parameters *milliseconds* — The interval in milliseconds that IS-IS LSPs can be sent from the interface expressed as a decimal integer.

Values 0 — 65535

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
no lsp-lifetime

Context config>router>isis

Description This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** 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.

The LSP refresh timer is derived from this formula: $\text{lsp-lifetime}/2$

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

Default **1200** — LSPs originated by the router should be valid for 1200 seconds (20 minutes).

Parameters *seconds* — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 — 65535

lsp-mtu-size

Syntax **lsp-mtu-size** *size*
no lsp-mtu-size

Context config>router>isis

Description This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then ISIS 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 bounce IS-IS for the change to take effect.

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

Default 1492

IS-IS Commands

Parameters *size* — Specifies the LSP MTU size.

Values 490 — 9190

lsp-refresh-interval

Syntax **lsp-refresh-interval** *seconds*
no lsp-refresh-interval

Context config>router>isis

Description This command configures the IS-IS LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for **lsp-lifetime** must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The no form of the 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.

Default 600 seconds

Parameters *seconds* — Specifies the refresh interval.

Values 150 — 65535

lsp-wait

Syntax **lsp-wait** *lsp-wait* [*lsp-initial-wait* [*lsp-second-wait*]]

Context config>router>isis

Description This command is used to customize the throttling of IS-IS LSP-generation. 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 **lsp-wait** timer until a maximum value is reached.

Parameters *lsp-max-wait* — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.

Values 1 — 120

Default 5

lsp-initial-wait — Specifies the initial LSP generation delay in seconds.

Values 0 — 100

Default 0

lsp-second-wait — Specifies the hold time in seconds between the first and second LSP generation.

Values 1 — 100

Default 1

ipv4-multicast

Syntax	[no] ipv4-multicast
Context	config>router>is-is>multi-topology
Description	This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance. The no form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.
Default	no ipv4-multicast

ipv6-multicast

Syntax	[no] ipv6-multicast
Context	config>router>is-is>multi-topology
Description	This command enables support for the IPv6 topology (MT4) within the associate IS-IS instance. The no form of this command disables support for the IPv6 topology (MT4) within the associated IS-IS instance.
Default	no ipv6-multicast

multi-topology

Syntax	[no] multi-topology
Context	config>router>isis
Description	This command enables IS-IS multi-topology support.
Default	disabled

topology

Syntax	topology <i>mt-id</i> rtm <i>rtm-id</i> <i>rtm-name</i> no topology <i>mt-id</i>
Context	config>router>is-is>multi-topology
Description	This command creates a new topology within the associate IS-IS instance. In addition, it associates the IS-IS topology with the specified RTM instance. Routes generated from the topology SPF calculation are in turn added to this associate RTM instance. The no form of this command deletes the specified IS-IS topology.
Default	No default

Parameters *mt-id* — Specify the topology to be created (Note: in Release 11.0 this parameters is limited to 3 or 4.
rtm-id — RTM Instance ID that is to be associated with the new IS-IS topology.

Values integer: 3 — 32

rtm-name — string name given to the RTM instance.

multicast-import

Syntax **[no] multicast-import**

Context config>router>isis

Description This command enables the submission of routes into the multicast Route Table Manager (RTM) by IS-IS. The **no** form of the command disables the submission of routes into the multicast RTM.

Default no multicast-import

mesh-group

Syntax **mesh-group {value | blocked}**
no mesh-group

Context config>router>isis>interface *ip-int-name*

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 created 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** — The interface does not belong to a 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 — 2000000000

blocked — Prevents an interface from flooding LSPs.

metric

Syntax	metric <i>metric</i> no metric
Context	config>router>isis>if>level <i>level-number</i>
Description	<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 no form of the command reverts to the default value.</p>
Default	10 — A metric of 10 for the level on the interface is used.
Parameters	<i>metric</i> — The metric assigned for this level on this interface.
Values	1 — 16777215

advertise-passive-only

Syntax	[no] advertise-passive-only
Context	config>router>isis
Description	This command enables and disables IS-IS to advertise only prefixes that belong to passive interfaces.

advertise-router-capability

Syntax	advertise-router-capability {area as} no advertise-router-capability
Context	config>router>isis
Description	<p>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.</p> <p>The parameters (area and as) control the scope of the capability advertisements.</p> <p>The no form of this command, disables this capability.</p>
Default	no advertise-router-capability
Parameters	<p>area — are only advertised within the area of origin.</p> <p>as — are only advertised throughout the entire autonomous system</p>

area-id

Syntax [no] **area-id** *area-address*

Context config>router>isis

Description This command was previously named the **net** *network-entity-title* command. The **area-id** 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.

A maximum of 3 area addresses can be configured.

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.

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.

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.

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.

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.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first area address.

The **no** form of the command removes the area address.

Default **none** — No area address is assigned.

Parameters *area-address* — 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.

overload

Syntax	overload [timeout seconds] no overload				
Context	config>router>isis				
Description	<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> <p>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.</p> <p>The no form of the command causes the router to exit the overload state.</p>				
Default	no overload				
Parameters	<p><i>seconds</i> — The time, in seconds, that this router must operate in overload state.</p> <table> <tr> <td>Default</td><td>infinity (overload state maintained indefinitely)</td></tr> <tr> <td>Values</td><td>60 — 1800</td></tr> </table>	Default	infinity (overload state maintained indefinitely)	Values	60 — 1800
Default	infinity (overload state maintained indefinitely)				
Values	60 — 1800				

overload-on-boot

Syntax	overload-on-boot [timeoutseconds] no overload-on-boot
Context	config>router>isis
Description	<p>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:</p> <ol style="list-style-type: none"> 1. The timeout timer expires. 2. A manual override of the current overload state is entered with the config>router>isis>no overload command. <p>The no overload command does not affect the overload-on-boot function.</p> <p>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:</p> <p>L1 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p>L2 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p>This state can be cleared with the config>router>isis>no overload command.</p>

IS-IS Commands

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.

Default no overload-on-boot

Use **show router ospf status** and/or **show router isis status** commands to display the administrative and operational state as well as all timers.

Parameters **timeout** *seconds* — Configure the timeout timer for overload-on-boot in seconds.

Values 60 — 1800

passive

Syntax [**no**] **passive**

Context config>router>isis>if
config>router>isis>if>level

Description 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.

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.

The **no** form of the command removes the passive attribute.

Special Cases **Service Interfaces** — Service interfaces (defined using the service-prefix command in **config>router**) are passive by default.

All other Interfaces — All other interfaces are not passive by default.

Default **passive** — Service interfaces are passive.
no passive — All other interfaces are not passive.

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 Default preferences are listed in the following table:

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*
IS-IS level 2 external	165	Yes*

*. External preferences are changed using the **external-preference** command in the **config>router>isis>level level-number** context.

Parameters *preference* — The preference for external routes at this level expressed as a decimal integer.

Values 1 — 255

priority

Syntax **priority** *number*
no priority

Context **config>router>isis>if>level** *level-number*

Description This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

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.

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

Default 64

IS-IS Commands

Parameters *number* — The priority for this interface at this level.

Values 0 — 127

sd-offset

Syntax **sd-offset** *offset-value*
no sd-offset

Context config>router>isis>if>level

Description If the pre-FEC error rate of the associated DWDM port crosses the configured **sd-threshold**, 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 **sd-threshold** value is configured under that port.
The **no** form of the command reverts the offset value to 0.

Default no sd-offset

Parameters *offset-value* — Specifies the amount the interface metric is increased by if the **sd-threshold** is crossed.

Values 0 — 16777215

sf-offset

Syntax **sf-offset** *offset-value*
no sf-offset

Context config>router>isis>if>level

Description If the pre-FEC error rate of the associated DWDM port crosses the configured **sf-threshold**, 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 **sf-threshold** value is configured under that port.
The **no** form of the command reverts the offset value to 0.

Default no sf-offset

Parameters *offset-value* — Specifies the amount the interface metric is increased by if the **sf-threshold** is crossed.

Values 0 — 16777215

psnp-authentication

Syntax	[no] psnp-authentication
Context	config>router>isis config>router>isis>level
Description	This command enables authentication of individual ISIS packets of partial sequence number PDU (PSNP) type. The no form of the command suppresses authentication of PSNP packets.

reference-bandwidth

Syntax	reference-bandwidth <i>bandwidth-in-kbps</i> reference-bandwidth [tbps <i>Tera-bps</i>] [gbps <i>Giga-bps</i>] [mbps <i>Mega-bps</i>] [kbps <i>Kilo-bps</i>] no reference-bandwidth
Context	
Description- config>router>isis	This command configures the reference bandwidth that provides the basis of bandwidth relative costing. 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: $\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$ If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 M/bps 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 config>router>isis context.) If the reference bandwidth is not configured, then all interfaces have a default metric of 10. The no form of the command reverts to the default value.
Default	no reference-bandwidth — No reference bandwidth is defined. All interfaces have a metric of 10.
Parameters	<i>bandwidth-in-kbps</i> — The reference bandwidth in kilobits per second expressed as a decimal integer. Values 1 — 1000000000 <i>tbps Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer. Values 1 — 4 <i>gbps Giga-bps</i> — The reference bandwidth in gigabits per second expressed as a decimal integer. Values 1 — 999 <i>mbps Mega-bps</i> — The reference bandwidth in megabits per second expressed as a decimal integer. Values 1 — 999

kbps *Kilo-bps* — reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 999

rsvp-shortcut

Syntax [no] **rsvp-shortcut**

Context config>router>isis

Description This command enables the use of an RSVP-TE shortcut for resolving IGP routes by IS-IS or OSPF routing protocols.

This command instructs IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. 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 which 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 LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When rsvp-shortcut is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **configure>router>mpls>lsp>to**, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP 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 OSPF or IS-IS will only use RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as end-points for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

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 next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
 - select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
 - exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
 - prefer tunnel next-hop over ip next-hop
 - within tunnel next-hops:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index

- within ip next-hops:
 - i. select lowest downstream router-id
 - ii. if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies 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.

The ingress IOM will spray 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 **rsvp-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.

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.

The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

Default **no rsvp-shortcut**

segment-routing

Syntax **segment-routing**
no segment-routing

Context config>router>isis

Description 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 ISIS 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 advertize 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.

prefix-sid-range

Syntax `prefix-sid-range {global | start-label label-value max-index index-value}
no prefix-sid-range`

Context config>router>isis>segment-routing

Description 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) 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. Note that all resulting net label values (start-label + 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.

Parameters	start-label <i>label-value</i> — the label offset for the SR label range of this IGP instance.
	Values 0 — 524287
	Default none
	max-index <i>index-value</i> — the maximum value of the prefix SID index range for this IGP instance.
	Values 0 — 524287
	Default none

tunnel-mtu

Syntax	tunnel-mtu <i>bytes</i> no tunnel-mtu
Context	config>router>isis>segment-routing
Description	<p>This command configures the MTU of all SR tunnels within each IGP instance.</p> <p>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 and directed LFA can add at least two more labels to the tunnel for a total of three. There is no default value for this new command. If the user does not configure a SR tunnel MTU, the MTU will be fully determined by IGP as explained below.</p> <p>The MTU of the SR tunnel is then determined as follows:</p> $SR_Tunnel_MTU = MIN \{Cf_g_SR_MTU, IGP_Tunnel_MTU - 3 \text{ labels}\}$ <p>Where,</p> <p><i>Cfg_SR_MTU</i> 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 fully determined by the IGP interface calculation explained next.</p> <p><i>IGP_Tunnel_MTU</i> 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.</p> <p>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.</p>
Parameters	<i>bytes</i> — the size of the Maximum Transmission Unit (MTU) in bytes.
	Values 512— 9198
	Default none

tunnel-table-pref

Syntax	tunnel-table-pref <i>preference</i> no tunnel-table-pref
Context	config>router>isis>no segment-routing
Description	<p>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.</p> <p>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 .</p> <p>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-OSPF and SR-ISIS).</p> <p>The global default TTM preference for the tunnel types is as follows:</p> <ul style="list-style-type: none"> • 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 <p>The default value for SR-OSPF (SR-ISIS) is the same regardless if one or more OSPF (IS-IS) 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.</p>
Parameters	<p><i>preference</i> — integer value to represent the preference of IS-IS SR tunnels in TTM.</p> <p>Values 1— 255</p> <p>Default 11</p>

advertise-tunnel-link

Syntax	[no] advertise-tunnel-link
Context	config>router>isis
Description	<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 rsvp-shortcut and advertise-tunnel-link options are enabled for a given IGP instance, then the advertise-tunnel-link 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>Note that the igp-shortcut option under the LSP name governs the use of the LSP with both the rsvp-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.</p> <p>The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.</p> <p>The no form of this command disables forwarding adjacency and hence disables the advertisement of RSVP LSP into IGP.</p>
Default	no advertise-tunnel-link

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> no retransmit-interval
Context	config>router>isis>interface <i>ip-int-name</i>
Description	<p>This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.</p> <p>The no form of the command reverts to the default value.</p>
Default	100
Parameters	<i>seconds</i> — The interval in seconds that IS-IS LSPs can be sent on the interface.
Values	1 — 65535

spf-wait

Syntax	[no] spf-wait <i>spf-wait</i> [<i>spf-initial-wait</i> [<i>spf-second-wait</i>]]
Context	config>router>isis
Description	This command defines the maximum interval between two consecutive SPF calculations in seconds. 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 <i>spf-second-wait</i> interval. For example, if the <i>spf-second-wait</i> 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 <i>spf-wait</i> value. The SPF interval will stay at <i>spf-wait</i> 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 <i>spf-initial-wait</i> .
Default	no spf-wait
Parameters	<p><i>spf-wait</i> — Specifies the maximum interval in seconds between two consecutive spf calculations.</p> <p>Values 1 — 120</p> <p>Default 10</p> <p><i>spf-initial-wait</i> — Specifies the initial SPF calculation delay in milliseconds after a topology change.</p> <p>Values 10 — 100000</p> <p>Default 1000</p> <p><i>spf-second-wait</i> — Specifies the hold time in milliseconds between the first and second SPF calculation.</p> <p>Values 1 — 100000</p> <p>Default 1000</p>

summary-address

Syntax	summary-address { <i>ip-prefix/mask</i> <i>ip-prefix</i> [<i>netmask</i>]} <i>level</i> [tag <i>tag</i>] no summary-address { <i>ip-prefix/mask</i> <i>ip-prefix</i> [<i>netmask</i>]}
Context	config>router>isis
Description	This command creates summary-addresses.
Default	none
Parameters	<p><i>ip-prefix/mask</i> — Specifies information for the specified IP prefix and mask length.</p> <p>Values <i>ipv4-prefix</i>: a.b.c.d (host bits must be 0) <i>ipv4-prefix-length</i>: 0 — 32</p> <p><i>netmask</i> — The subnet mask in dotted decimal notation.</p> <p>Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)</p>

level — Specifies IS-IS level area attributes.

Values level-1, level-2, level-1/2

tag *tag* — Assigns an OSPF, RIP or ISIS tag to routes matching the entry.

Values Accepts decimal or hex formats:
 OSPF and ISIS: [0x0..0xFFFFFFFF]H
 RIP: [0x0..0xFFFF]H

suppress-default

Syntax [no] **suppress-default**

Context config>router>isis

Description This command enables or disables IS-IS to suppress the installation of default routes.

traffic-engineering

Syntax [no] **traffic-engineering**

Context config>router>isis

Description This command configures traffic-engineering and determines if IGP shortcuts are required.

Default disabled

unicast-import-disable

Syntax [no] **unicast-import-disable**

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 disabled

wide-metrics-only

Syntax [no] **wide-metrics-only**

Context config>router>isis>level *level-number*

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.

Show Commands

isis

Syntax	isis [<i>isis-instance</i>]
Context	show>router
Description	This command displays information for a specified IS-IS instance.
Parameters	<i>isis-instance</i> — Specifies the instance ID for an IS-IS instance.
Values	1–31
Default	0

adjacency

Syntax	adjacency [<i>ip-address</i> <i>ip-int-name</i> <i>nbr-system-id</i>] [detail]
Context	show>router>isis
Description	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.
Parameters	<p><i>ip-address</i> — When specified, only adjacencies with that interface is displayed.</p> <p>Values ipv4-address: a.b.c.d (host bits must be 0)</p> <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>detail — All output is displayed in the detailed format.</p>
Output	Standard and Detailed IS-IS Adjacency Output — The following table describes the standard and detailed command output fields for an IS-IS adjacency.

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.

Label	Description (Continued)
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.

Sample Output

```
*A:Dut-C# show router isis adjacency
=====
Router 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
=====
Router Base ISIS Instance 0 Adjacency
=====
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 Supressed : Disabled
Number of Restarts: 0
Last Restart at : Never
=====
*A:Dut-C#
```



```

*A:ALA-A# show router isis adjacency 180.0.7.12
=====
ISIS 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
=====
ISIS 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
=====
ISIS Adjacency
=====
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

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

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

```

Show Commands

```
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 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#
```

```
A:Dut-A# show router isis status
```

```
=====
ISIS Status
```

```
=====
System Id           : 0100.2000.1001
Admin State         : Up
Ipv4 Routing        : Enabled
Last Enabled        : 08/28/2006 10:22:17
Level Capability     : L2
Authentication Check : True
Authentication Type  : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Enabled
Graceful Restart     : Disabled
GR Helper Mode       : Disabled
LSP Lifetime        : 1200
LSP Wait             : 1 sec (Max)    1 sec (Initial)    1 sec (Second)
Adjacency Check      : loose
L1 Auth Type        : none
L2 Auth Type        : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference        : 15
L2 Preference        : 18
L1 Ext. Preference   : 160
L2 Ext. Preference   : 165
L1 Wide Metrics      : Disabled
L2 Wide Metrics      : Enabled
L1 LSDB Overload     : Disabled
L2 LSDB Overload     : Disabled
L1 LSPs              : 0
L2 LSPs              : 15
Last SPF             : 08/28/2006 10:22:25
SPF Wait             : 1 sec (Max)    10 ms (Initial)    10 ms (Second)
Export Policies      : None
Area Addresses       : 49.0001
```

```
=====
* indicates that the corresponding row element may have been truncated.
A:Dut-A#
```

capabilities

Syntax	capabilities [<i>system-id</i> <i>/sp-id</i>] [<i>level level</i>]
Context	show>router>isis
Description	This command displays the IS-IS capability information.

Sample Output

```
*A:Dut-C# show router isis capabilities
=====
Router Base ISIS Instance 0 Capabilities
=====
Displaying Level 1 capabilities
-----
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 , SRGB Base:20000, Range:10001
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
Level (1) Capability Count : 4
Displaying Level 2 capabilities
-----
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
LSP ID : Dut-B.00-01
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
```

Show Commands

```
LSP ID : Dut-C.00-01
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
LSP ID : Dut-D.00-01
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 , SRGB Base:20000, Range:10001
SR Alg: metric based SPF
LSP ID : Dut-E.00-01
Level (2) Capability Count : 8
=====
*A:Dut-C#
```

database

- Syntax** **database** [*system-id* | *lsp-id*] [*detail*] [*level level*]
- Context** show>router>isis
- Description** This command displays the entries in the IS-IS link state database.
- Parameters**

system-id — Only the LSPs related to that *system-id* are listed. If no *system-id* or *lsp-id* are specified, all database entries are listed.

lsp-id — Only the specified LSP (hostname) is listed. If no *system-id* or *lsp-id* are specified, all database entries are listed.

Sample Output

```
*A:ALA-A# show router isis database
=====
ISIS 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
=====
```

```

ISIS Database
=====
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

```

Show Commands

```
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
=====
*A:Dut-B#
```

hostname

Syntax **hostname**

Context show>router>isis

Description This command displays the hostname database. There are no options or parameters.

Output **IS-IS Hostname Output** — The following table describes output fields for IS-IS hostname output.

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
=====
Hosts
=====
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

Syntax `interface [ip-int-name | ip-address] [detail]`

Context `show>router>isis`

Description This command shows IS-IS interface information. When no *ip-addr* or the *ip-int-name* is specified, all interfaces are listed.

Parameters *ip-address* — Only displays the interface information associated with the specified IP address.

Values *ipv4-address* a.b.c.d (host bits must be 0)

ip-int-name — Only displays the interface information associated with the specified IP interface name.

detail — All output is given in the detailed format.

Output **IS-IS Interface Output** — The following table describes 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. 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
=====
ISIS 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:cses-V93# show router isis interface "system" detail
=====
```

Show Commands

```
Router Base ISIS Instance 0 Interfaces
=====
-----
Interface      : system                      Level Capability: L1L2
Oper State     : Down                      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: Disabled                 LFA              :

Level          : 1                        Adjacencies      : 0
Auth Keychain  : Disabled
Auth Type      : None                      Metric           : 0
Hello Timer    : 9                        IPv6-Ucast-Met   : 0
Priority       : 64                        IPv6-Mcast-Met   : 0
Passive        : No                       IPv4-Mcast-Met   : 0
SD-Offset     : 0                        SF-Offset        : 0
Hello Mult.    : 3

Level          : 2                        Adjacencies      : 0
Auth Keychain  : Disabled
Auth Type      : None                      Metric           : 0
Hello Timer    : 9                        IPv6-Ucast-Met   : 0
Priority       : 64                        IPv6-Mcast-Met   : 0
Passive        : No                       IPv4-Mcast-Met   : 0
SD-Offset     : 0                        SF-Offset        : 0
Hello Mult.    : 3
=====
```

lfa-coverage

- Syntax** lfa-coverage
- Context** show>router>isis
- Description** This command displays IS-IS LFA coverage information.

Sample Output

```
*A:SR# show router isis lfa-coverage
=====
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#

```

neighbor

Syntax neighbor

Context show>router>isis

Description This command displays neighboring route information in the IS-IS route table.

Sample Output

```

A:linus-212# show router isis neighbor
=====
Topology Table
=====
Node Interface Nexthop
-----
IS-IS IP paths (MT-ID 0), Level 1
-----
linus-211.00 ab linus-211
linus-216.00 ab linus-211
linus-216.01 ab linus-211
-----
IS-IS IP paths (MT-ID 0), Level 2
-----
linus-211.00 ab linus-211
linus-216.00 ab linus-211
linus-216.01 ab linus-211
=====

```

Show Commands

prefix-sids

Syntax **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*]

Context show>router>isis

Description This command displays IS-IS prefix SIDs.

Sample Output

```
*A:Dut-C# show router isis prefix-sids
```

```
=====
Router Base ISIS Instance 0 Prefix/SID Table
=====
```

Prefix	SID	Lvl/Typ	SRMS MT	AdvRtr Flags
4.0.0.1/32	1	2/Int.	N	Dut-B
			0	RNnP
4.0.0.1/32	1	2/Int.	N	Dut-C
			0	RNnP
4.0.0.1/32	1	1/Int.	N	Dut-D
			0	NnP
4.0.0.1/32	1	2/Int.	N	Dut-D
			0	NnP
4.0.0.1/32	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
-----
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

Syntax **routes** [ipv4-unicast | ipv4-multicast | ipv6-multicast | mt *mt-id-number*] [*ip-prefix[/prefix-length]*] [alternative] [exclude-shortcut]

Context show>router>isis

Description This command displays the routes in the IS-IS route table.

Parameters **ipv4-unicast** — Displays IPv4 unicast parameters.
mt *mt-id-number* — Displays multi-topology parameters.

Values 0, 2

alternative — Displays LFA details.

Output **IS-IS Route Output** — The following table describes 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.

Label	Description (Continued)
Hostname	Hostname for the specific <i>system-id</i> .

Sample Output

```

*A:Dut-C# show router isis routes exclude-shortcut alternative
=====
Route Table
=====
Prefix[Flags]           Metric    Lvl/Typ    Ver.    SysID/Hostname
NextHop                MT         AdminTag
-----
1.1.2.0/24              20        1/Int.     0       Dut-A
  1.1.3.1                0         0
1.1.2.0/24              20        1/Int.     0       Dut-B
  1.2.3.2                0         0
1.1.3.0/24              10        1/Int.     0       Dut-C
  0.0.0.0                0         0
1.2.3.0/24              10        1/Int.     0       Dut-C
  0.0.0.0                0         0
1.2.4.0/24              20        1/Int.     0       Dut-B
  1.2.3.2                0         0
1.3.5.0/24              10        1/Int.     0       Dut-C
  0.0.0.0                0         0
1.4.5.0/24              20        1/Int.     0       Dut-E
  1.3.5.5                0         0
1.4.6.0/24              30        1/Int.     0       Dut-B
  1.2.3.2                0         0
1.4.6.0/24              30        1/Int.     0       Dut-E
  1.3.5.5                0         0
10.20.1.1/32            10        1/Int.     0       Dut-A
  1.1.3.1                0         0
10.20.1.2/32            10        1/Int.     0       Dut-B
  1.2.3.2                0         0
10.20.1.3/32            0         1/Int.     0       Dut-C
  0.0.0.0                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
Flags: L = LFA nexthop available
=====*A:Dut-C#

*A:Dut-B# show router isis routes
=====
Route Table
=====

```

```

=====
Prefix [Flags]      Metric  Lvl/Typ  Ver.  SysID/Hostname
NextHop            MT      AdminTag
-----
10.20.1.2/32        0        1/Int.   3      Dut-B
  0.0.0.0            0
10.20.1.3/32 [L]    10        2/Int.   2      Dut-C
  10.20.3.3          0
10.20.1.4/32        10        2/Int.   3      Dut-D
  10.20.4.4          0
10.20.1.5/32        20        2/Int.   3      Dut-C
  10.20.3.3          0
10.20.1.6/32        20        2/Int.   3      Dut-D
  10.20.4.4          0
10.20.3.0/24        10        1/Int.   3      Dut-B
  0.0.0.0            0
10.20.4.0/24        10        1/Int.   3      Dut-B
  0.0.0.0            0
10.20.5.0/24        20        2/Int.   2      Dut-C
  10.20.3.3          0
10.20.6.0/24        20        2/Int.   4      Dut-D
  10.20.4.4          0
10.20.9.0/24        20        2/Int.   3      Dut-D
  10.20.4.4          0
10.20.10.0/24       30        2/Int.   3      Dut-C
  10.20.3.3          0
-----
Routes : 11
Flags: L = LFA nexthop available
=====
*A:Dut-B#

*A:Dut-B# show router isis routes alternative
=====
Route Table
=====
Prefix [Flags]      Metric  Lvl/Typ  Ver.  SysID/Hostname
NextHop            MT      AdminTag
Alt-Nexthop        Alt-Metric
-----
0 1/Int. 3 Dut-B -10.20.1.2/32
  0.0.0.0 0
10.20.1.3/32 10 2/Int. 2 Dut-C
  10.20.3.3 0
  10.20.3.3 (lfa) 15
10.20.1.4/32 10 2/Int. 3 Dut-D
  10.20.4.4 0
10.20.1.5/32 20 2/Int. 3 Dut-C
  10.20.3.3 0
10.20.1.6/32 20 2/Int. 3 Dut-D
  10.20.4.4 0
10.20.3.0/24 10 1/Int. 3 Dut-B
  0.0.0.0 0
10.20.4.0/24 10 1/Int. 3 Dut-B
  0.0.0.0 0
10.20.5.0/24 20 2/Int. 2 Dut-C
  10.20.3.3 0
10.20.6.0/24 20 2/Int. 4 Dut-D
  10.20.4.4 0

```

Show Commands

```

10.20.9.0/24                20          2/Int.    3      Dut-D
   10.20.4.4                0           0
10.20.10.0/24              30          2/Int.    3      Dut-C
   10.20.3.3                0           0
-----Routes : 11
Flags: LFA = Loop-Free Alternate nexthop
=====
*A:Dut-B#

*A:Dut-A# show router isis routes
=====
Route Table
=====
Prefix                               Metric    Lvl/Typ Ver.   SysID/Hostname
NextHop                             MT
-----
10.10.1.0/24                        10        1/Int.   5      Dut-A
   0.0.0.0                          0
10.10.3.0/24                        20        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.4.0/24                        20        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.5.0/24                        30        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.9.0/24                        60        1/Int.   52    Dut-F
   10.10.21.6                      0
10.10.10.0/24                      70        1/Int.   52    Dut-F
   10.10.21.6                      0
10.10.12.0/24                      20        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.13.0/24                     10        1/Int.   7      Dut-A
   0.0.0.0                          0
10.10.14.0/24                     20        1/Int.   52    Dut-F
   10.10.21.6                      0
10.10.15.0/24                     30        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.16.0/24                     30        1/Int.  137    Dut-B
   10.10.1.2                        0
10.10.21.0/24                     10        1/Int.  48    Dut-A
   0.0.0.0                          0
10.10.22.0/24                     30        1/Int.  137    Dut-B
   10.10.1.2                        0
10.20.1.1/32                       0         1/Int.  10     Dut-A
   0.0.0.0                          0
10.20.1.2/32                       10        1/Int.  137    Dut-B
   10.10.1.2                        0
10.20.1.3/32                       20        1/Int.  137    Dut-B
   10.10.1.2                        0
10.20.1.4/32                       20        1/Int.  137    Dut-B
   10.10.1.2                        0
10.20.1.5/32                       30        1/Int.  137    Dut-B
   10.10.1.2                        0
10.20.1.6/32                       10        1/Int.   52    Dut-F
   10.10.21.6                      0
3FFE::A0A:100/120                 10        1/Int.   5      Dut-A
   ::                               0
10.10.1.0/24                      10        1/Int.  65     Dut-A
   0.0.0.0                          2

```

```

10.10.13.0/24          10          1/Int.  65      Dut-A
    0.0.0.0            2
10.10.21.0/24          10          1/Int.  65      Dut-A
    0.0.0.0            2
10.20.1.1/32           0          1/Int.  65      Dut-A
    0.0.0.0            2
3FFE::A0A:100/120      10          1/Int.  65      Dut-A
    ::                2
3FFE::A0A:300/120      20          1/Int.  116     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:400/120      20          1/Int.  116     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:500/120      30          1/Int.  130     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:900/120      60          1/Int.  71      Dut-F
    FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:A00/120      70          1/Int.  71      Dut-F
    FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:C00/120      20          1/Int.  116     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:D00/120      10          1/Int.  65      Dut-A
    ::                2
3FFE::A0A:E00/120      20          1/Int.  71      Dut-F
    FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
3FFE::A0A:F00/120      30          1/Int.  130     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:1000/120     30          1/Int.  130     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A0A:1500/120     10          1/Int.  65      Dut-A
    ::                2
3FFE::A0A:1600/120     30          1/Int.  127     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:101/128      0          1/Int.  65      Dut-A
    ::                2
3FFE::A14:102/128      10          1/Int.  116     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:103/128      20          1/Int.  130     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:104/128      20          1/Int.  127     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:105/128      30          1/Int.  130     Dut-B
    FE80::2281:1FF:FE01:1-"ip-3FFE::A0A:101" 2
3FFE::A14:106/128      10          1/Int.  71      Dut-F
    FE80::2285:FFFF:FE00:0-"ies-1-3FFE::A0A:1501" 2
-----
Routes : 43
=====
*A:Dut-A#

*A:SRR# show router isis routes 1.1.1.0/24
=====
Route Table
=====
Prefix[Flags]          Metric    Lvl/Typ   Ver.   SysID/Hostname
NextHop               MT        AdminTag
-----
1.1.1.0/24 [L]         7540      1/Int.    6109   SRL
    60.60.1.1          0         0
-----

```

Show Commands

```
No. of Routes: 1
Flags: L = LFA nexthop available
=====
*A:SRR# show router isis routes 1.1.1.0/24 alternative
=====
Route Table
=====
Prefix[Flags]           Metric      Lvl/Typ    Ver.   SysID/Hostname
  NextHop              MT          AdminTag
Alt-Nexthop            Alt-Metric  Alt-Type
-----
1.1.1.0/24              7550        1/Int.     6114   SRL
  60.60.1.1              0           0
  11.22.12.4 (LFA)      16784764    linkProtection
-----

No. of Routes: 1
Flags: LFA = Loop-Free Alternate nexthop
=====
*A:SRR#

*A:SRR# show router isis routes 1.1.1.0/24
=====
Route Table
=====
Prefix[Flags]           Metric      Lvl/Typ    Ver.   SysID/Hostname
  NextHop              MT          AdminTag
-----
1.1.1.0/24 [L]          7540        1/Int.     6109   SRL
  60.60.1.1              0           0
-----

No. of Routes: 1
Flags: L = LFA nexthop available
=====
*A:SRR# show router isis routes 1.1.1.0/24 alternative
=====
Route Table
=====
Prefix[Flags]           Metric      Lvl/Typ    Ver.   SysID/Hostname
  NextHop              MT          AdminTag
Alt-Nexthop            Alt-Metric  Alt-Type
-----
1.1.1.0/24              7550        1/Int.     6114   SRL
  60.60.1.1              0           0
  11.22.12.4 (LFA)      16784764    linkProtection
-----

No. of Routes: 1
Flags: LFA = Loop-Free Alternate nexthop
=====
*A:SRR#

*A:Dut-C# show router isis routes ipv4-unicast
=====
Router Base ISIS Instance 0 Route Table
=====
Prefix[Flags]           Metric      Lvl/Typ    Ver.   SysID/Hostname
  NextHop              MT          AdminTag
                               SID[Flags]
```


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

Flags: L = LFA nexthop available

*A:Dut-C# show router isis routes ipv4-unicast alternative

Router Base ISIS Instance 0 Route Table

Prefix[Flags]	Metric	Lvl/Typ	Ver.	SysID/Hostname
NextHop			MT	AdminTag
Alt-Nexthop			Alt-Metric	Alt-Type
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

Show Commands

```

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
Flags: L = Loop-Free Alternate nexthop
Legend: LP = linkProtection, NP = nodeProtection
=====
*A:Dut-C#
```

spf

Syntax	spf [detail] [lfa]
Context	show>router>isis
Description	This command displays information regarding SPF calculation.
Output	Router ISIS Output — The following table describes the output fields for ISIS SPF.

Label	Description
Node	The route node and mask.
Interface	The outgoing interface name for the route.
Metric	The route's metric.
Nexthop	The system ID of nexthop or hostname.
SNPA	The Subnetwork Points of Attachment (SNPA) where a router is physically attached to a subnetwork.

Sample Output

```

linus-212>show>router>isis# show router isis spf
=====
Path Table
=====
Node Interface Nexthop
-----
linus-211.00 ab linus-211
x linux-222 *
linus-216.00 ab linus-211
linus-216.00 ab linus-211
linus-216.01 ab linus-211
=====

A:ALA-A# show router isis spf
=====
Path Table
=====
Node Interface Nexthop
-----
abr_sjc.00 if2/2 dist_oak
abr_sjc.00 if2/3 dist_nj
dist_oak.00 if2/2 dist_oak
dist_nj.00 if2/3 dist_nj
acc_nj.00 if2/3 dist_nj
acc_ri.00 if2/3 dist_nj
core_west.00 if2/8 core_west
core_east.00 lag-1 core_east
asbr_west.00 if2/8 core_west
asbr_east.00 if2/5 asbr_east
abr_sjc.00 lag-1 core_east
abr_sjc.00 if2/8 core_west
abr_lax.00 lag-1 core_east
abr_lax.00 if2/8 core_west
abr_dfw.00 if2/5 asbr_east
abr_dfw.00 lag-1 core_east
abr_dfw.00 if2/8 core_west
dist_arl.00 if2/5 asbr_east
dist_arl.00 lag-1 core_east
dist_arl.00 if2/8 core_west
dist_msq.00 if2/5 asbr_east
dist_msq.00 lag-1 core_east
dist_msq.00 if2/8 core_west
acc_arl.00 if2/5 asbr_east
acc_arl.00 lag-1 core_east
acc_arl.00 if2/8 core_west
acc_msq.00 if2/5 asbr_east
acc_msq.00 lag-1 core_east
acc_msq.00 if2/8 core_west
acc_msq.03 if2/5 asbr_east
acc_msq.03 lag-1 core_east
acc_msq.03 if2/8 core_west
acc_msq.04 if2/5 asbr_east
acc_msq.04 lag-1 core_east
acc_msq.04 if2/8 core_west
=====
A:ALA-A#

```

Show Commands

```
A:ALA-A# show router isis spf detail
=====
Path Table
=====
Node       : abr_sjc.00                      Metric : 20
Interface  : if2/2                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_oak

Node       : abr_sjc.00                      Metric : 20
Interface  : if2/3                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : dist_oak.00                    Metric : 10
Interface  : if2/2                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_oak

Node       : dist_nj.00                    Metric : 10
Interface  : if2/3                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : acc_nj.00                      Metric : 20
Interface  : if2/3                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : acc_ri.00                      Metric : 20
Interface  : if2/3                          SNPA   : 00:00:00:00:00:00
Nexthop    : dist_nj

Node       : core_west.00                  Metric : 10
Interface  : if2/8                          SNPA   : 00:00:00:00:00:00
Nexthop    : core_west

...
=====
A:ALA-A#

*A:SRR>config>router>isis# show router isis spf lfa detail
=====
Path Table
=====
Level 1
-----
Node       : SRL.00                      Metric : 1000
Interface  : aps-10                      SNPA   : n/a
Nexthop    : SRL

LFA intf   : sru4-3                      LFA Metric : 16778214
LFA nh     : SRU4                        LFA type   : linkProtection

Node       : SRU4.00                      Metric : 1000
Interface  : sru4-3                      SNPA   : n/a
Nexthop    : SRU4

LFA intf   : aps-11                      LFA Metric : 16777214
LFA nh     : SRU4                        LFA type   : linkProtection

Node       : ESS-7.00                    Metric : 16779214
Interface  : aps-10                      SNPA   : n/a
Nexthop    : SRL
```

LFA intf : sru4-3	LFA Metric : 33555428
LFA nh : SRU4	LFA type : nodeProtection
Node : ESS-7.09	Metric : 16778214
Interface : sru4-3	SNPA : n/a
Nexthop : SRU4	
Node : ESS-7.0A	Metric : 16778214
Interface : sru4-3	SNPA : n/a
Nexthop : SRU4	
Node : ESS-7.0B	Metric : 16778214
Interface : sru4-3	SNPA : n/a
Nexthop : SRU4	
Node : SRC.00	Metric : 1000
Interface : src-asap-2	SNPA : n/a
Nexthop : SRC	
Node : SRC.01	Metric : 2000
Interface : aps-10	SNPA : n/a
Nexthop : SRL	
Node : SRC.02	Metric : 16778214
Interface : src-asap-2	SNPA : n/a
Nexthop : SRC	
Node : SRC.03	Metric : 1063
Interface : aps-10	SNPA : n/a
Nexthop : SRL	

Level 2

Node : SRL.00	Metric : 1000
Interface : aps-1	SNPA : n/a
Nexthop : SRL	
Node : SRL.00	Metric : 1000
Interface : if-30.30.1.2	SNPA : n/a
Nexthop : SRL	
Node : SRL.00	Metric : 1000
Interface : if-40.40.1.2	SNPA : n/a
Nexthop : SRL	
Node : SRL.09	Metric : 2000
Interface : aps-1	SNPA : n/a
Nexthop : SRL	
Node : SRL.09	Metric : 2000
Interface : if-30.30.1.2	SNPA : n/a
Nexthop : SRL	
Node : SRL.09	Metric : 2000
Interface : if-40.40.1.2	SNPA : n/a
Nexthop : SRL	
Node : SRL.0A	Metric : 2000

Show Commands

```

Interface : aps-1                      SNPA      : n/a
Nexthop   : SRL

Node      : SRL.0A                      Metric    : 2000
Interface : if-30.30.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : SRL.0A                      Metric    : 2000
Interface : if-40.40.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : SRU4.00                     Metric    : 2000
Interface : aps-1                       SNPA      : n/a
Nexthop   : SRL

Node      : SRU4.00                     Metric    : 2000
Interface : if-30.30.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : SRU4.00                     Metric    : 2000
Interface : if-40.40.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : SRU4.14                     Metric    : 16778214
Interface : germ-1                      SNPA      : n/a
Nexthop   : SROne

Node      : ESS-7.00                    Metric    : 1000
Interface : ess-7-1                     SNPA      : 00:16:4d:c1:25:e4
Nexthop   : ESS-7

LFA intf  : ess-7-2                      LFA Metric : 16777214
LFA nh    : ESS-7                        LFA type   : linkProtection

Node      : ESS-7.05                     Metric    : 16778214
Interface : aps-1                       SNPA      : n/a
Nexthop   : SRL

Node      : ESS-7.05                     Metric    : 16778214
Interface : if-30.30.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : ESS-7.05                     Metric    : 16778214
Interface : if-40.40.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : ESS-7.05                     Metric    : 16778214
Interface : ess-7-1                     SNPA      : 00:16:4d:c1:25:e4
Nexthop   : ESS-7

Node      : ESS-7.06                     Metric    : 16778214
Interface : aps-1                       SNPA      : n/a
Nexthop   : SRL

Node      : ESS-7.06                     Metric    : 16778214
Interface : if-30.30.1.2                SNPA      : n/a
Nexthop   : SRL

Node      : ESS-7.06                     Metric    : 16778214

```

Interface : if-40.40.1.2	SNPA	: n/a
Nexthop : SRL		
Node : ESS-7.06	Metric	: 16778214
Interface : ess-7-1	SNPA	: 00:16:4d:c1:25:e4
Nexthop : ESS-7		
Node : ESS-7.0C	Metric	: 16778214
Interface : ess-7-1	SNPA	: 00:16:4d:c1:25:e4
Nexthop : ESS-7		
Node : ESS-7.0D	Metric	: 16778214
Interface : ess-7-1	SNPA	: 00:16:4d:c1:25:e4
Nexthop : ESS-7		
Node : ESS-7.0E	Metric	: 16778214
Interface : ess-7-1	SNPA	: 00:16:4d:c1:25:e4
Nexthop : ESS-7		
Node : ESS-7.12	Metric	: 16778214
Interface : ess-7-1	SNPA	: 00:16:4d:c1:25:e4
Nexthop : ESS-7		
Node : ESS-7.12	Metric	: 16778214
Interface : germ-1	SNPA	: n/a
Nexthop : SROne		
Node : SRC.00	Metric	: 1000
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
LFA intf : aps-1	LFA Metric	: 2063
LFA nh : SRL	LFA type	: linkProtection
Node : SRC.05	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.06	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.07	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.08	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.09	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.0A	Metric	: 16778214
Interface : src-asap-1	SNPA	: n/a
Nexthop : SRC		
Node : SRC.0B	Metric	: 16778214

Show Commands

```
Interface : src-asap-1          SNPA      : n/a
Nexthop   : SRC

Node      : SRC.0C              Metric    : 16778214
Interface : src-asap-1          SNPA      : n/a
Nexthop   : SRC

Node      : SRC.0D              Metric    : 16778214
Interface : src-asap-1          SNPA      : n/a
Nexthop   : SRC

Node      : SRC.0E              Metric    : 16778214
Interface : src-asap-1          SNPA      : n/a
Nexthop   : SRC

Node      : SROne.00            Metric    : 1000
Interface : germ-1              SNPA      : n/a
Nexthop   : SROne

LFA intf  : aps-1               LFA Metric : 2000
LFA nh    : SRL                 LFA type   : linkProtection

=====
*A:SRR>config>router>isis#
```

spf-log

Syntax	spf-log [detail]
Context	show>router>isis
Description	This command displays IS-IS SPF log information.
Parameters	detail — Displays detailed logged information.
Output	Router ISIS SFP Log Output — The following table describes the ISIS 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.
Reason	Displays the reason(s) for the SPF run.

Label	Description (Continued)
	NEWADJ: An adjacency changed.
	NEWLSP: A new LSP was received.
	NEWAREA: An area changed.
	NEWREACH: A prefix changed.
	ECMPCHANGED: An ECMP path changed.
	NEWMETRIC: A prefix metric changed.
	RESTART: The graceful restart exited.
	LSPEXPIRED: An LSP expired.
	DBCHANGED: The LSP database was cleared by an administrator.
	LSPCONTENT: The content of an LSP changed.
	NEWPREF: The external route preference changed.
	NEWNLPID: The routed protocols (IPv4 or IPv6) changed.
	MANUALREQ: An SPF calculation was requested by an administrator.
	ADMINTAGCHANGED: An administrative tag changed.
	TUNNELCHANGED: An MPLS tunnel changed.

Sample Output

```
*A:Dut-C# show router isis spf-log
```

```
=====
Router 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# show router isis spf-log
```

Show Commands

```
A:SetupCLI# show router isis spf-log detail
=====
ISIS 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#

A:ALA-48# show router isis spf-log
=====
ISIS SPF Log
=====
When      Duration      L1 Nodes  L2 Nodes  Event Count
-----
01/30/2007 11:01:54    <0.01s    1         1         3
```

```
-----
Log Entries : 1
=====
```

```
A:ALA-48#
```

statistics

Syntax **statistics**

Context show>router>isis

Description This command displays information regarding IS-IS traffic statistics.

Output **IS-IS Statistics Output** — This table describes 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.

Sample Output

```
*A:Dut-C# show router isis statistics
```

```
=====
Router Base ISIS Instance 0 Statistics
=====
```

Show Commands

```
ISIS Instance      : 0
Purge Initiated    : 0
Sid SRGB err       : 0
LSP Regens.        : 17
Sid dupl err       : 0

CSPF Statistics
Requests           : 0
Paths Found        : 0
Request Drops      : 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#
```

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.

Label	Description
System-id	Neighbor system ID.
Admin State	Up — IS-IS is administratively up. Down — IS-IS is administratively down.

Label	Description (Continued)
Ipv4 Routing	Enabled — IPv4 routing is enabled. Disabled — IPv4 routing is disabled.
Last Enabled	The date/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. 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. 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. 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	When enabled, excludes a prefix entry defined in the specified LFA policy from LFA calculation.

Sample Output

```
A:SetupCLI# show router isis status
=====
ISIS Status
=====
System Id           : 0100.2003.0040
Admin State         : Up
Ipv4 Routing        : Enabled
Ipv6 Routing        : Disabled
Last Enabled        : 10/01/2011 04:11:47
Level Capability     : L1L2
Authentication Check : True
Authentication Type  : MD5
```

Show Commands

```
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Enabled
Graceful Restart : Enabled
GR Helper Mode : Enabled
LSP Lifetime : 2400
LSP Wait : 120 sec (Max) 0 sec (Initial) 1 sec (Second)
LSP MTU Size : 9190 (Config) 9190 (Oper)
Adjacency Check : strict
L1 Auth Type : password
L2 Auth Type : md5
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference : 30
L2 Preference : 22
L1 Ext. Preference : 40
L2 Ext. Preference : 34
L1 Wide Metrics : Enabled
L2 Wide Metrics : Enabled
L1 LSDB Overload : Disabled
L2 LSDB Overload : Disabled
L1 LSPs : 1
L2 LSPs : 1
L1 Default Metric : 10
L2 Default Metric : 10
L1 IPv6 Def Metric : 10
L2 IPv6 Def Metric : 10
Last SPF : 10/01/2011 04:11:49
SPF Wait : 120 sec (Max) 1000 ms (Initial) 1000 ms (Second)
Export Policies : None
Multicast Import : None
Multi-topology : Disabled
Advertise-Passive-On*: Enabled
Suppress Default : Enabled
Default Route Tag : 1
Area Addresses : 49.0001
: 47.1234.1234.1234.1234.ffff
Ldp Sync Admin State : Down
LDP-over-RSVP : Enabled
RSVP-Shortcut : Enabled
Advertise-Tunnel-Link: Disabled
Export Limit : 0
Exp Lmt Log Percent : 0
Total Exp Routes(L1) : 0
Total Exp Routes(L2) : 0
IID TLV : Enabled
All-L1-MacAddr : 01:80:c2:00:00:14
All-L2-MacAddr : 01:80:c2:00:00:15
Loopfree-Alternate : Enabled
L1 LFA : Included
L2 LFA : Included
```

=====

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

Task 113366: IPFRR

```

-----

*A:SRR>config>router>isis# show router isis status
=====
ISIS Status
=====
System Id           : 1100.2000.1002
Admin State         : Up
Ipv4 Routing        : Enabled
Ipv6 Routing        : Enabled, Native
Last Enabled        : 07/25/2011 18:11:34
Level Capability     : L1L2
Authentication Check : True
Authentication Type  : None
CSNP-Authentication : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication : Enabled
Traffic Engineering : Enabled
Graceful Restart     : Disabled
GR Helper Mode       : Disabled
LSP Lifetime        : 350
LSP Wait            : 5 sec (Max) 0 sec (Initial) 1 sec (Second)
LSP MTU Size        : 1492 (Config) 1492 (Oper)
Adjacency Check     : loose
L1 Auth Type        : none
L2 Auth Type        : none
L1 CSNP-Authenticati* : Enabled
L1 HELLO-Authenticat* : Enabled
L1 PSNP-Authenticati* : Enabled
L1 Preference       : 15
L2 Preference       : 18
L1 Ext. Preference  : 160
L2 Ext. Preference  : 165
L1 Wide Metrics     : Enabled
L2 Wide Metrics     : Enabled
L1 LSDB Overload    : Disabled
L2 LSDB Overload    : Disabled
L1 LSPs             : 86
L2 LSPs             : 113
L1 Default Metric   : 10
L2 Default Metric   : 10
L1 IPv6 Def Metric   : 10
L2 IPv6 Def Metric   : 10
Last SPF            : 07/26/2011 15:27:09
SPF Wait            : 10 sec (Max) 1000 ms (Initial) 1000 ms (Second)
Export Policies     : isis_from_ALL_to_0
Multicast Import     : both
Multi-topology      : Enabled
Advertise-Passive-On* : Disabled
Suppress Default     : Disabled
Default Route Tag    : None
Area Addresses       : 49.0001
Ldp Sync Admin State : Up
LDP-over-RSVP       : Enabled
RSVP-Shortcut        : Disabled
Advertise-Tunnel-Link : Disabled
Export Limit         : 0
Exp Lmt Log Percent  : 0

```

Show Commands

```
Total Exp Routes(L1) : 0
Total Exp Routes(L2) : 0
IID TLV                : Disabled
All-L1-MacAddr         : 01:80:c2:00:00:14
All-L2-MacAddr         : 01:80:c2:00:00:15
Loopfree-Alternate     : Enabled
=====
* indicates that the corresponding row element may have been truncated.
*A:SRR>config>router>isis#

*A:Dut-C# show router isis status

=====
Router Base ISIS Instance 0 Status
=====
System Id              : 0100.2000.1003
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           : Enabled, Native
Mcast Ipv4 Routing     : Enabled, Native
Mcast Ipv6 Routing     : Enabled, Native
Last Enabled           : 01/26/2015 11:22:13
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
LSP Lifetime           : 1200
LSP Refresh Interval   : 600
LSP Wait               : 5 sec (Max) 0 sec (Initial) 1 sec (Second)
LSP MTU Size           : 1300 (Config) 1300 (Oper)
Adjacency Check        : loose
L1 Auth Keychain       : Disabled
L1 Auth Type           : none
L1 CSNP-Authenticati* : Enabled
L1 HELLO-Authenticat* : Enabled
L1 PSNP-Authenticati* : Enabled
L1 Preference          : 15
L1 Ext. Preference    : 160
L1 Wide Metrics        : Enabled
L1 LSDB Overload       : Disabled
L1 LSPs                : 4
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                : 01/26/2015 11:22:51
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)    : 0
IID 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-Authenticat*   : Enabled
L2 PSNP-Authenticati*   : Enabled
L2 Preference           : 18
L2 Ext. Preference      : 165
L2 Wide Metrics         : Enabled
L2 LSDB Overload        : Disabled
L2 LSPs                 : 8
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
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
RSVP-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
Remote-LFA              : Enabled
L1 LFA                  : Included
L2 LFA                  : Included
Advertise Router Cap     : Area
Hello Padding           : Disabled
Ignore Lsp Errors       : Disabled
Reference Bandwidth     : 0
Ucast Import Disable    : None
Segment Routing          : Up
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C#

```

Show Commands

The following two examples illustrate sample output and detail sample output with LFA policies configured in the configure router ISIS context.

```
*A:SRR# show router isis status
```

```
=====
ISIS Status
=====
System Id           : 0100.2000.1003
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-Authenticat*: 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
```

```

All-L1-MacAddr      : 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     : 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
RSVP-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

=====
ISIS Interfaces
=====
-----
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

```

Show Commands

```
Admin Groups      : None
Ldp Sync          : outOfService
Ldp Timer State   : Disabled
Route Tag         : None
Ldp Sync Wait     : Disabled
Ldp Tm Left       : 0
LFA               : Included

Level            : 1
Desg. IS         : Dut-C
Auth Keychain     : Disabled
Auth Type         : None
Hello Timer       : 9
Priority          : 64
Passive           : No
SD-Offset         : 0
Hello Mult.       : 3
Adjacencies       : 0
Metric           : 10
IPv6-Ucast-Met    : 10
IPv6-Mcast-Met    : 10
IPv4-Mcast-Met    : 10
SF-Offset         : 0

Level            : 2
Desg. IS         : Dut-C
Auth Keychain     : Disabled
Auth Type         : None
Hello Timer       : 9
Priority          : 64
Passive           : No
SD-Offset         : 0
Hello Mult.       : 3
Adjacencies       : 0
Metric           : 10
IPv6-Ucast-Met    : 10
IPv6-Mcast-Met    : 10
IPv4-Mcast-Met    : 10
SF-Offset         : 0
```

=====

*A:SRR#

summary-address

Syntax **summary-address** [*ip-address* [/*mask*]]

Context show>router>isis

Description Displays ISIS summary addresses.

Output **Router ISIS Summary Address Output** — The following table describes the ISIS 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
=====
ISIS 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

Syntax **topology** [**ipv4-unicast** | **ipv6-unicast** | **ipv4-multicast** | **ipv6-multicast**] **mt** *mt-id-number* [**lfa**] [**detail**]

Context show>router>isis

Description This command shows IS-IS topology information.

Parameters **ipv4-unicast** — Displays IPv4 unicast parameters.

ipv4-multicast — Displays IPv4 multicast parameters.

multicastmulticast

mt *mt-id-number* — Displays multi-topology parameters.

Values 0, 2, 3, 4

lfa — Displays LFA (loop free alternative) information.

detail — Displays detailed topology information.

Output **Router ISIS Topology Output** — The following table describes the ISIS 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
=====
Topology Table

```

Show Commands

```

=====
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 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
=====
*A:Dut-A#

```

Clear Commands

isis

Syntax	isis [<i>isis-instance</i>]
Context	clear>router>isis
Description	This command enables the context to clear and reset ISIS protocol entities.
Parameters	<i>isis-instance</i> — Specifies the instance ID for the IS-IS instance. Values 1 — 31

adjacency

Syntax	adjacency [<i>system-id</i>]
Context	clear>router>isis
Description	This command clears and resets the entries from the IS-IS adjacency database.
Parameters	<i>system-id</i> — When the system ID is entered, only the specified entries are removed from the IS-IS adjacency database.

database

Syntax	database [<i>system-id</i>]
Context	clear>router>isis
Description	This command removes the entries from the IS-IS link-state database which contains information about PDUs.
Parameters	<i>system-id</i> — When the system ID is entered, only the specified entries are removed from the IS-IS link-state database.

Clear Commands

export

Syntax	export
Context	clear>router>isis
Description	This command re-evaluates route policies participating in the export mechanism, either as importers or exporters of routes.

spf-log

Syntax	spf-log
Context	clear>router>isis
Description	This command clears the SPF log.

statistics

Syntax	statistics
Context	clear>router>isis
Description	This command clears and resets IS-IS statistics.

Debug Commands

isis

Syntax	isis [<i>isis-instance</i>]
Context	debug>router
Description	This command enables the context to debug IS-IS protocol entities.
Parameters	<i>isis-instance</i> — Specifies the IS-IS instance.
Values	1 — 31

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. The no form of the command disables debugging.

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

Syntax	[no] graceful-restart
Context	debug>router>isis
Description	This command enables debugging for IS-IS graceful-restart. The no form of the command disables debugging.

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-address</i>] no interface
Context	debug>router>isis
Description	This command enables debugging for IS-IS interface. The no form of the command disables debugging.

leak

Syntax	leak [<i>ip-address</i>] no leak
Context	debug>router>isis
Description	This command enables debugging for IS-IS leaks. The no form of the command disables debugging.

lsdb

Syntax	[no] lsdb [<i>level-number</i>] [<i>system-id</i> <i>lsp-id</i>]
Context	debug>router>isis
Description	This command enables debugging for Link State DataBase (LSDB). The no form of the command disables debugging.

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 [<i>packet-type</i>] [<i>ip-int-name</i> <i>ip-address</i>] [detail]
Context	debug>router>isis
Description	This command enables debugging for IS-IS packets. The no form of the command disables debugging.

rtm

Syntax	rtm [<i>ip-address</i>] 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.

spf

Syntax	[no] spf [<i>level-number</i>] [<i>system-id</i>]
Context	debug>router>isis
Description	This command enables debugging for IS-IS SFP. The no form of the command disables debugging.

In This Chapter

This chapter provides information about the Border Gateway Protocol (BGP) and its implementation in SR-OS.

Topics in this chapter include:

- [BGP Overview on page 642](#)
- [BGP Sessions on page 643](#)
- [BGP Design Concepts on page 651](#)
- [BGP Messages on page 656](#)
- [BGP Path Attributes on page 661](#)
- [BGP Routing Information Base \(RIB\) on page 680](#)
- [BGP Applications on page 701](#)
- [BGP Configuration Process Overview on page 715](#)
- [Configuration Notes on page 716](#)

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.

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 28](#) 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 28](#) the router ALA-A forms an EBGP session with ALA-D.



OSRG053

Figure 28: BGP Sessions

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 on page 654](#).

In SR-OS a BGP session is 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 7x50 is the **active** side of TCP connections to remote neighbors, meaning that as soon as a session leaves the *Idle* state 7x50 attempts to setup 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 BGP session can be configured for **passive** mode so that the 7x50 router only listens for an incoming

connection and does not attempt to setup the outgoing connection. The source IP address used to setup the TCP connection to the 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
-

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 a **shutdown** command). 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.
-

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 on page 658](#) 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.

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.

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.

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.

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 7x50 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 7x50 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. Note that 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.

BGP Graceful Restart

Some BGP routers do not have redundant control plane processor modules or else do not support BGP HA with the same quality or coverage as 7x50 routers. When dealing with such routers or certain error conditions BGP graceful restart is a good option for minimizing the network disruption caused by a control plane reset. BGP graceful restart assumes that the router restarting its BGP sessions has the ability/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). In order for BGP graceful restart to be used on a session both routers must advertise the BGP graceful restart capability during the OPEN message exchange; see the section titled [BGP Advertisement on page 657](#) for more details.

On 7x50 routers BGP graceful restart is enabled on one or more BGP sessions by configuring the **graceful-restart** command in the global, group or neighbor context. The command causes the GR capability to be advertised and enables helper mode support for IPv4 (AFI1, SAFI1), IPv6 (AFI 2,

SAFI), VPN-IPv4 and VPN-IPv6 routes. Note that the GR capability advertised by a 7x50 router does not list the supported AFI/SAFI unless **enable-notification** is configured.

On a 7x50 router 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 7x50 router marks the received IPv4, IPv6, VPN-IPv4 and VPN-IPv6 routes from the peer as 'stale' and starts a restart timer. (As noted above the 'stale' state is not factored into the BGP decision process and not made visible to other routers in the network.) The restart timer derives its initial value from the Restart Time carried in the peer's last GR capability. (The default Restart Time advertised by 7x50 routers is 300 seconds but this can be changed using the **restart-time** command.) When the restart timer expires helping stops if the session has 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 forwarding state was preserved then helping continues and the peers exchange routes per the normal procedure. When each router has advertised all its routes for a particular 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/SAFI. When the helping router receives an EOR it deletes all remaining stale routes of the AFI/SAFI that were not refreshed in the most recent set of UPDATE messages; there is an upper limit on the amount of time that routes can remain stale (before being deleted if they were not refreshed) and this is configurable using the **stale-routes-time**.



NOTE: 7x50 routers always abort the GR helper process, regardless of the failure trigger, if there is a second reset before GR has successfully completed.

BGP Session Security

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.

TTL Security Mechanism

The TTL security mechanism 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.

On 7x50 routers 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 multi-hop EBGP and IBGP sessions as well.



NOTE: When a 7x50 router sends IP packets to an IBGP peer they are originated with an IP TTL value of 64. When a 7x50 router sends IP packets to an EBGP peer they are originated with an IP TTL value of 1, except if **multihop** is configured, and in that case the TTL value is taken from the **multihop** command.

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.


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 29](#). 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 on page 652](#). 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 on page 654](#).



al_0138

Figure 29: Fully Meshed BGP Configuration

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 subdivided 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 30 depicts the same network as Figure 29 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.

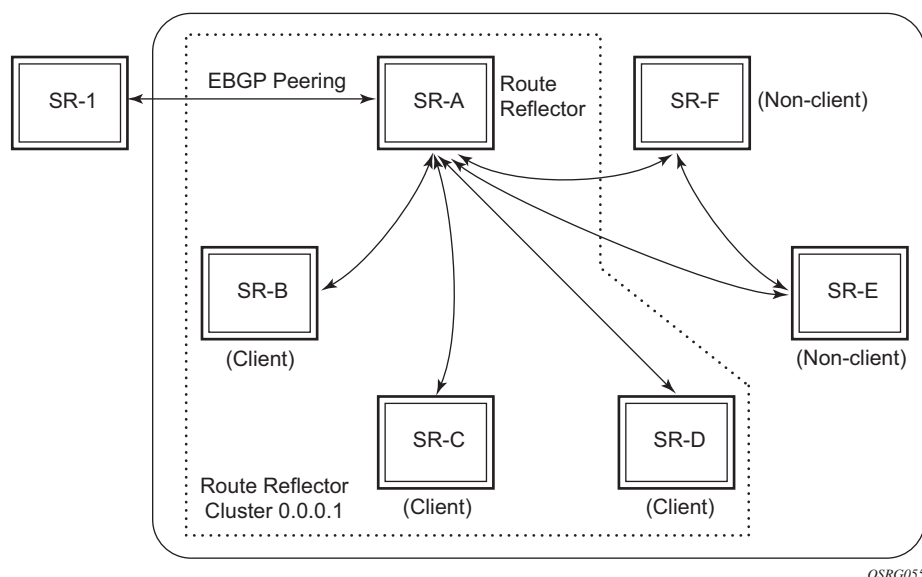


Figure 30: BGP Configuration with Route Reflectors

A 7x50 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 on a 7x50 router (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.

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. Note that 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.

BGP Messages

BGP protocol operation relies on the exchange of BGP messages between peers. 7x50 and most other routers support the following 5 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 7x50 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.

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 the section titled [BGP Session States on page 644](#) 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 7x50 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 on page 664](#). More details about 4-byte AS numbers are described in the section titled [4-Octet Autonomous System Numbers on page 665](#).
- 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 that changes to the configured **hold-time** trigger a session reset.
- BGP Identifier — The router ID of the BGP speaker. In Open messages sent by 7x50, 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 that 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, which is described in the next section. When a BGP router receives an Open message with an unsupported optional parameter type it terminates the session. A 7x50 router always sends a optional parameter in its Open message unless **disable-capability-negotiation** is configured.

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.

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.

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. SROS 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)
-

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 a 7x50 router packs as many NLRI into a single Update message as possible. This requires identifying all the routes that share the same path attribute values.

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. On 7x50 routers 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 on page 644](#).

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 graceful restart is enabled for this scenario; see the section titled [BGP Graceful Restart on page 647](#) 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.
-

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 a reset.

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 on page 677](#)).

A 7x50 router only sends Route Refresh messages for AFI/SAFI associated with VPN routes that carry Route Target extended communities - i.e. 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 that 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.

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_AGG (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)

Origin

- CONNECTOR (optional transitive)
 - PMSI_TUNNEL (optional transitive)
 - AIGP (optional non-transitive)
-

Origin

The ORIGIN path attribute indicates the origin of the path information. There are 3 supported values:

- IGP (0)
- EGP (1)
- Incomplete (2)

When a 7x50 router originates a VPN-IP prefix (from a non-BGP route) it sets the value of the Origin attribute to IGP. When a 7x50 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.

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 7x50 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. (Note that 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 on page 654](#).
 - 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. Note that 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.

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.

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.

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 7x50 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

Next-Hop

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 7x50 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 as 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 7x50 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.

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 and the NEXT_HOP attribute is not included in the Update message; see the section titled [Multi-Protocol BGP Attributes on page 677](#) for more details.

In IPv4 and IPv6 routes advertised by a 7x50 router the BGP next-hop address is set as follows:

- When a route is advertised to an EBGP peer the BGP next-hop is always changed to the local-address used with the EBGP peer and this behavior cannot be overridden, even with a BGP export policy. (See the section titled [BGP Sessions on page 643](#) for an explanation

of how the local-address is determined.) The one exception to this rule occurs when the third-party-nexthop command is applied:

- When a route is received from one EBGp peer and is advertised to another EBGp that is in the same IP subnet and has been configured with the third-party-nexthop command (at the BGP instance, group or neighbor level), the BGP next-hop in the advertised route remains unchanged.
- When a route is advertised to an IBGP or confederation-EBGP peer and the route is not locally originated the advertising router does not modify the next-hop by default, however:
 - If the **next-hop-self** command is applied to a confederation-EBGP peer this changes the next-hop to the local-address used with that peer.
 - If the **next-hop-self** command is applied to an IBGP peer this changes the next-hop to the local-address used with that peer, but only if the route came from a confed-EBGP or EBGp peer.
 - A BGP export policy applied to an IBGP or confederation-EBGP session can change the next-hop to any IPv4 address, regardless of the route source (IBGP, EBGp, confed-EBGP).
- When a 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 exported into BGP, with certain exceptions (e.g. black-hole next-hop).

In VPN-IPv4 routes advertised by a 7x50 router the BGP next-hop address is set as follows:

- When a route is advertised to an EBGp peer the BGP next-hop is changed to the local-address used with the EBGp peer if **enable-inter-as-vpn** is configured; otherwise there is no change to the next-hop.
- When a route is received from an EBGp peer and advertised to an IBGP or confederation-EBGP peer the BGP next-hop is changed to the local-address used with the IBGP or confederation-EBGP peer if **enable-inter-as-vpn** is configured. If **enable-inter-as-vpn** is not configured the next-hop may be changed with the **next-hop-self** command but this is not recommended because it can result in a change of the next-hop without a change in the VPN label.
- When a 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 changes the next-hop to the local-address used with the peer.

In Label-IPv4 routes advertised by a 7x50 router the BGP next-hop address is set as follows:

- When a route is advertised to an EBGp peer the BGP next-hop is always changed to the local-address used with the EBGp peer and this behavior cannot be overridden.

- When a route is received from an EBGp peer and advertised to an IBGP or confederation-EBGP peer next-hop-self is applied automatically (i.e. the next-hop is modified to the local-address used with the peer), however:
 - A BGP export policy applied to the IBGP or confederation-EBGP session can change the next-hop to any IPv4 address
 - If the **next-hop-unchanged label-ipv4** command is applied to the receiving IBGP or confederation-EBGP peer this overrides the automatic next-hop-self and causes no modification to the BGP next-hop
 - *At the current time SR-OS does not support next-hop-self for label-IPv4 routes advertised to a confed-EBGP peer.*
- When a route is received from an IBGP peer and reflected to another IBGP peer the next-hop is not modified by default, however:
 - If the **next-hop-self** command is applied to the receiving IBGP peer this changes the next-hop to the local-address used with that peer.
 - A BGP export policy applied to the IBGP session can change the next-hop to any IPv4 address.

In 6PE routes advertised by a 7x50 router the BGP next-hop address is set as follows:

- When a 6PE route is locally-originated and advertised to any BGP peer the BGP next-hop is an IPv4-mapped IPv6 address allocated from the ::FFFF/96 range. The bottom 32 bits of the IPv6 address is the IPv4 local-address used with the peer.
 - *At the current time SR-OS does not support sending and receiving 6PE routes with EBGp peers.*
- When a route is received from an IBGP peer and reflected to another IBGP peer the next-hop is not modified by default, however:
 - A BGP export policy applied to the IBGP session can apply **next-hop-self** or change the next-hop to any IPv4-mapped IPv6 address. Note that the **next-hop-self** command at the group/neighbor configuration level has no effect in this case.
- When a route is advertised to a confederation-EBGP peer the next-hop is not modified by default, however:
 - If the **next-hop-self** command is applied to the session this changes the next-hop to the IPv4-mapped IPv6 address corresponding to the IPv4 local-address used with the peer.
 - A BGP export policy applied to the IBGP session can change the next-hop to any IPv4-mapped IPv6 address.

Next-Hop IPv4 Address Family over IPv6

For IBGP sessions, next-hop information is taken from the system interface. If the system interface does not have an IPv4 address configured, no next-hop will be populated without a routing policy

applied to the BGP session, and BGP NLRI messages is not sent for the IPv4 address family. The use of an export policy allows the operator to configure next-hop information explicitly.

For EBGp sessions, the next-hop information must be taken from an export routing policy that explicitly sets the next-hop based on operator configuration. If the export policy is not set, the BGP NLRI messages are not sent for the IPv4 address family due to no `next_hop`.

Next-Hop VPN-IPv4 Address Family over IPv6

For IBGP and EBGp sessions, next-hop information is specified as the system IP address.

Next-Hop VPN-IPv6 Address Family over IPv6

For IBGP sessions, the next-hop information is specified as the system IP address encoded as an IPv4-mapped-IPv6 address.

For EBGp sessions, the next-hop information is specified as the system IP address encoded as an IPv4-mapped-IPv6 address, by the way of an export policy configured by the user.

Next-Hop Resolution

For a BGP router to use a BGP route for forwarding it must know how to reach the BGP next-hop of the route. The process of determining the local interface or tunnel that should be 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 SR-OS details are explained below:

- Next-hop resolution is always done for IPv4 routes. Note the following:
 - BGP routes are eligible to resolve a BGP next-hop only if the **use-bgp-routes** command is configured.
 - 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.
 - If the LPM route (accepted by the policy) is a BGP route then the BGP next-hop of that route is looked up and this time other BGP routes are not eligible to be resolving routes, whether or not **use-bgp-routes** is configured. In other words 7x50 routers support BGP routes resolving BGP routes with one level of recursion.
 - BGP shortcuts are discussed further in the section titled [Next-hop Resolution Using Tunnels on page 701](#).
- Next-hop resolution is always done for IPv6 routes. The 7x50 router looks for an eligible IPv6 route that matches the BGP next-hop address in the route table. Note the following:
 - BGP routes are eligible to resolve a BGP next-hop only if the **use-bgp-routes** command is configured.
 - 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.
 - If the LPM route (accepted by the policy) is a BGP route then the BGP next-hop of that route is looked up and this time other BGP routes are not eligible to be resolving routes, whether or not **use-bgp-routes** is configured. In other words 7x50 routers support BGP routes resolving BGP routes with one level of recursion.
- SR-OS attempts to resolve the next-hop of a VPN-IPv4 or VPN-IPv6 route only if it is imported into one or more VPRNs or if it is advertised with a new BGP next-hop. If the

next-hop is part of a local subnet the next-hop is automatically resolved by the direct route. If the next-hop is remote (more than one hop away):

- The 7x50 router looks for a tunnel in the tunnel-table with a destination that matches the BGP next-hop address. If the route is imported into VPRNs the tunnel types eligible to resolve the BGP next-hop are controlled by the **auto-bind** configurations of the VPRNs. If the route is advertised with a new BGP next-hop the eligible tunnel types are controlled by the **transport-tunnel** command.
- If there is no matching tunnel-table entry then the BGP next-hop is unresolved and the VPN-IP route is effectively *invalid* despite displaying as *valid* and *best*. A VPN-IP route that is *invalid* due to an unresolved next-hop can be advertised to any type of peer, but only if the next-hop is not changed.
- SR-OS always attempts to resolve the next-hop of a label-IPv4 route. If the next-hop is part of a local subnet the next-hop is automatically resolved by the direct route. If the next-hop is remote (more than one hop away):
 - The 7x50 router looks for a tunnel in the tunnel-table with a destination that matches the BGP next-hop address and a type allowed by the **transport-tunnel** command. If there are multiple matches the tunnel with the lowest preference is used (RSVP is preferred over LDP).
 - If there is no matching and eligible entry in the tunnel table but there is a /32 static route with a black-hole next-hop that matches the BGP next-hop address this static route automatically resolves the BGP next-hop.
 - If there is no matching and eligible entry in the tunnel table and no /32 static black-hole route then the BGP next-hop is unresolved and the label-IPv4 route is considered *invalid*. However note that a label-IPv4 route that is *invalid* due to an unresolved next-hop can still be reflected to an IBGP peer, whether or not **next-hop-self** is applied to the route.
- SR-OS always attempts to resolve the next-hop of a 6PE route.
 - The 7x50 router looks for an LDP tunnel in the tunnel-table with a destination that matches the IPv4 address contained in the IPv4-mapped IPv6 BGP next-hop address.
 - If there is no matching LDP entry in the tunnel table but there is a /128 static route with a black-hole next-hop that matches the IPv4-mapped IPv6 BGP next-hop address this static route automatically resolves the BGP next-hop.
 - If there is no matching LDP entry in the tunnel table and no /128 static black-hole route then the BGP next-hop is unresolved and the 6PE route is considered *invalid*. However note that a 6PE route that is *invalid* due to an unresolved next-hop can still be reflected to an IBGP peer, whether or not **next-hop-self** is applied to the route.
- SR-OS does not check for next-hop reachability in Flow-spec and RTC routes.

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.

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-hop(s) 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.

Using Multiple Address Families over IPv6 BGP Sessions

To ease transition to IPv6 and the deployment of IPv6 into service provider environments, SR-OS permits the transport of the following address families over an IPv6 transported BGP session (a BGP session where both neighbors are configured and transported over IPv6):

- IPv4
- VPN-IPv4
- IPv6
- VPN-IPv6

As the IPv4, VPN-IPv4 and VPN-IPv6 address families require an IPv4 NEXT_HOP address to be present in the BGP NLRI messaging, the following approaches are taken in SR-OS:

- For iBGP sessions, SR-OS will use the configured System Interface IPv4 address as the NEXT_HOP address; unless specifically overwritten by a routing export policy.
- For eBGP sessions, SR-OS requires the use of a routing export policy to set the NEXT_HOP to an appropriate address, such as the IPv4 address configured on the interface between eBGP neighbors.

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.

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.

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.

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:

Route Aggregation Path Attributes

- 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 7x50 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 7x50 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 on page 662](#).
 - 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 on page 662](#).
 - 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 on page 662](#).
- The COMMUNITY attribute contains all the communities from all the contributing routes.
- No MED attribute is included by default. Note that SR-OS does not require all the contributing routes to have the same MED value.

Community and Extended Community Attributes

A BGP route can be associated with one or more standard communities and one or more extended communities. All the standard communities are carried in a single COMMUNITIES attribute and all the extended communities currently supported by SR-OS are carried in a single EXTENDED_COMMUNITIES attribute.

Each standard community is 4 bytes; the first 2 bytes encode the AS number of the administrative entity that assigned the value in the last 2 bytes. In SR-OS a standard community member is input as *AS:value* to reflect this format. There are several well-known standard communities that 7x50 and most other BGP routers recognize:

- NO_EXPORT: When a route carries this community it must not be advertised outside a confederation boundary (i.e. 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 (i.e. to confed-EBGP peers or EBGp peers).

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 an associated community, this community is automatically

added to the BGP route. (Note that 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.)

If it is necessary to remove all the standard communities from all routes advertised to a BGP peer SR-OS supports the **disable-communities standard** command.

Extended communities provide more flexibility than standard communities. Each extended community is 8 bytes. The first 1 or 2 bytes identifies the type/sub-type and the remaining 6 or 7 bytes is a value. As of release 12.0R1 SR-OS supports the following types of extended communities:

- 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)
- 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 4-octet AS-specific
 - Route target (type 0x0202)
 - Route origin (type 0x0203)
 - OSPF domain ID (type 0x0205)
 - Source AS (type 0x0209)
- Transitive opaque
 - OSPF route type (type 0x0306)
- Transitive experimental
 - OSPF domain ID (type 0x8005)
 - Flow-spec traffic rate (type 0x8006)
 - Flow-spec traffic action (type 0x8007)
 - Flow-spec redirect (type 0x8008)
 - Layer 2 info (type 0x800A)

- EVPN
 - MAC mobility (type 0x0600)

Route target and route origin extended communities can be added to or removed from BGP routes using BGP import and export policies. Other types of extended communities are added automatically to the relevant types of routes.

If it is necessary to remove all the extended communities from all routes advertised to a BGP peer SR-OS supports the **disable-communities extended** command.

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 on page 652](#).

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 of a 7x50 router is controlled primarily by the **family** command. The AFI/SAFI supported by SR-OS as of Release 12.0R1 and the method of configuring the AFI/SAFI support is summarized in [Table 12](#).

Table 12: Multi-Protocol BGP support in SR-OS

Name	AFI	SAFI	Configuration Commands
IPv4 unicast	1	1	family ipv4
IPv4 multicast	1	2	family mcast-ipv4
IPv4 labeled unicast	1	4	family ipv4 advertise-label ipv4
NG-MVPN IPv4	1	5	family mvpn-ipv4
MDT-SAFI	1	66	family mdt-safi
VPN-IPv4	1	128	family vpn-ipv4
VPN-IPv4 multicast	1	129	family mcast-vpn-ipv4
RT constrain	1	132	family route-target
IPv4 flow-spec	1	133	family flow-ipv4
IPv6 unicast	2	1	family ipv6
IPv6 multicast	2	2	family mcast-ipv6
6PE	2	4	family ipv6 advertise-label ipv6
NG-MVPN IPv6	2	5	family mvpn-ipv6
VPN-IPv6	2	128	family vpn-ipv6
IPv6 flow-spec	2	133	family flow-ipv6
Multi-segment PW	25	6	family ms-pw
L2 VPN	25	65	family l2-vpn
EVPN	25	70	family evpn

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. Note that 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.

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 on page 665](#).

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 on page 681](#))
- 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 7x50 implementation AIGP is supported only in the base router BGP instance and only for the following types of routes: IPv4, label-IPv4, IPv6 and 6PE. 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 7x50 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).

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 7x50 BGP features in the context of the RIB architecture outlined above.

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 on page 659](#).
 - BGP import policies
-

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.

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.

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.

BGP Decision Process

When a BGP router has multiple routes in its LOC-RIB for the same NLRI its BGP decision process is responsible for deciding which one is the best. The best path can be used by the local router (e.g. for its own forwarding) and advertised to other BGP peers.

On 7x50 routers the BGP decision process orders *valid* LOC-RIB routes based on the following sequence of comparisons (if there multiple routes tied at step N then proceed to step N+1):

1. Select the route with the best (numerically lowest) route preference.
2. Select the route with the highest Local Preference (LOCAL_PREF).

3. From all routes with an AIGP metric (if any) select the route with the lowest sum of:
 - a. AIGP metric value stored with the LOC-RIB copy of the route.
 - b. The route table (or tunnel table) cost between the calculating router and the BGP NEXT_HOP in the received route.
 4. Select the route with the shortest AS Path. Note that AS numbers in AS_CONFED_SEQ and AS_CONFED_SET elements do not count towards the AS path length. Skip this step if **as-path-ignore** is configured for the address family.
 5. Select the route with the lowest Origin (IGP=0, EGP=1, Incomplete=2).
 6. Select the route with the lowest MED. Only compare MED in routes from the same neighbor AS by default. A missing MED attribute is considered equivalent to a MED value of 0 by default. Defaults can be changed with the **always-compare-med** command.
 7. Prefer routes learned from EBGp peers over routes learned from IBGP and confed-EBGP peers.
 8. Select the route with the lowest route or tunnel table cost to the NEXT_HOP. If **ignore-nh-metric** is configured skip this step.
 9. Select the route with lowest next-hop type (resolved in route-table = 0, resolved in tunnel-table = 1). If **ignore-nh-metric** is configured skip this step.
 10. Select the route received by the peer with the lowest Router ID; this comes from the ORIGINATOR_ID attribute (if present) or else the BGP identifier of the peer (received in its OPEN message). If **ignore-router-id** is configured and two EBGp routes are being compared keep the current best path and skip steps 11 and 12.
 11. Select the route with the shortest CLUSTER_LIST length.
 12. Select the route received from the peer with the lowest IP address.
-

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; 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 13](#) shows how the MED comparison of two paths is influenced by different forms of the **always-compare-med** command.

Table 13: MED Comparison with always-compare-med

Command	MED comparison step in decision process
no always-compare-med 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 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.

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*).

BGP Route Installation in the Route Table

If the best BGP path for an IPv4 or IPv6 prefix is the most preferred route to the destination it is installed in the IP route table unless **disable-route-table-install** is configured. The best BGP path is the most preferred route if has the numerically lowest route preference among all routes, of all protocols, to the destination. The default preference value for BGP routes is 170 but this can be changed using the **preference** command in the BGP or policy configuration.



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 the best path can be installed in the route table and there are other BGP paths (LOC-RIB routes) for the same IPv4 or IPv6 prefix that are nearly as good as the best path the additional paths can also be installed in the route table. This is called *BGP-Multipath* and it must be explicitly enabled using the **multipath** command. The **multipath** command specifies the maximum number of BGP paths (up to 32), including the overall best path, that BGP can install in the route table for any particular IPv4 or IPv6 prefix; in this scenario each BGP path is effectively one ECMP next-hop of the IP route and 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.



NOTE: In order for BGP to install a route with N ECMP next-hops in the route-table the associated routing instance must have the **ecmp** command in its configuration and the max number of ECMP next-hops specified as part of that command must have a value greater than or equal to N .

In SR-OS 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 both of the following criteria:

1. It is the overall best BGP path or else it is tied with the overall best path up to and including step 9 of the decision process as summarized in the section titled [BGP Decision Process on page 681](#).
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 EBGP peers 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.

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 EBGp peers.

It is important to note that *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, 6PE). 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.

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.

In SR OS, a Link Bandwidth Extended Community can be added to an IPv4, IPv6, VPN-IPv4 or VPN-IPv6 route using either route policies or the **ebgp-link-bandwidth** command. The **ebgp-link-bandwidth** command is supported in BGP group and neighbor configuration contexts and automatically adds (on import) a Link Bandwidth Extended Community to received routes from single-hop (directly connected) EBGp peers. Note that when a route is advertised to an EBGp peer, the Link Bandwidth Extended Community, if present, is always removed. The Link Bandwidth Extended Community associated with a BGP route can be displayed using the **show router bgp routes** commands; for the bandwidth value, the system automatically converts the binary value in the extended community to a decimal number in units of Mbps (1000000 bit/s).

7x50 routers automatically performs weighted ECMP for an IP BGP route when the ingress card is FP2 or better and all 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.

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 (requires chassis mode D)

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 on page 670](#) 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 on page 683](#)) also applies to this case.

For BGP tunnels there is no support for the equivalent of *IBGP-Multipath*. That is, if a BGP next-hop of the label-IPv4 route in the tunnel table is resolved by an LDP tunnel with multiple ECMP next-hops load-sharing is not supported across the LDP ECMP next-hops; only the first next-hop carries traffic towards the BGP next-hop.

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, labeled-IPv4, IPv6, 6PE, VPN-IPv4 and VPN-IPv6 routes. The scenarios supported by the base router BGP context are outlined in [Table 14](#).

Note that BGP fast reroute information specific to IP VPNs is described in the BGP Fast Reroute in a VPRN section of the 7x50 SR OS Services Guide.

Table 14: 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 an LDP or RSVP shortcut tunnel	IPv4 route with next-hop B resolved by an IPv4 route or an LDP or RSVP shortcut tunnel	Yes
IPv6	IPv6 route with next-hop A resolved by an IPv6 route OR 6PE route with next-hop A resolved by an LDP tunnel	IPv6 route with next-hop B resolved by an IPv6 route OR 6PE route with next-hop B resolved by an LDP tunnel	Yes, but if the 6PE routes are label-per-prefix the ingress card must be IOM3 or better for PIC
IPv4	Lbl-IPv4 route with next-hop A resolved by an LDP or RSVP tunnel	Lbl-IPv4 route with next-hop B resolved by an LDP or RSVP tunnel	Yes, if ingress card is IOM3 or better
IPv4	Lbl-IPv4 route with next-hop A resolved to an interface	Lbl-IPv4 route with next-hop B resolved to an interface	Yes, if ingress card is IOM3 or better
MPLS or Service	Lbl-IPv4 route with next-hop A resolved by an LDP or RSVP tunnel	Lbl-IPv4 route with next-hop B resolved by an LDP or RSVP tunnel	Yes
MPLS or Service	Lbl-IPv4 route with next-hop A resolved to an interface	Lbl-IPv4 route with next-hop B resolved to an interface	Yes

Calculating Backup Paths

In SR-OS BGP fast reroute is optional and must be enabled using either the **backup-path** or **install-backup-path** command.

The **backup-path** command is used in the base router context to control fast reroute on a per-routing instance and per-family (IPv4 and IPv6) basis. The command supports options to enable fast reroute for IPv4 prefixes only, for IPv6 prefixes only, or for all IPv4 and IPv6 prefixes.

The **install-backup-path** command is used to designate a specific set of IPv4 or IPv6 prefixes that are eligible for BGP fast reroute protection. The command enables a BGP import policy to restrict the set of routes that are programmed with a backup path.

When BGP fast reroute is enabled the control plane attempts to find an eligible backup path for every received IPv4 and/or IPv6 prefix, depending on configuration. In general the backup path is the single best path remaining after the primary ECMP paths and any paths with the same BGP next-hops as these paths have been removed. However the following points should be noted:

- A backup path is not calculated for a prefix if the best path is a labeled-IPv4 route and it has been programmed with multiple ECMP next-hops through different BGP next-hops.
- For labeled-IPv4 prefixes that are re-advertised with a new BGP next-hop the programmed backup path is the same for all prefixes that have the same best path and received label, even if the calculated backup path is different for some of the prefixes.

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.

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.

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*. 7x50 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 (1-25), 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.

Route Flap Damping (RFD)

Route flap damping is a mechanism supported by 7x50 and 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. Note that 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.

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.

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.

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. Note that 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. Note that 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. Note that 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). Also note that 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. Note that 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.

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. 7x50 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.

The 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 7x50 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.

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 the 7x50 router 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. Note that 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), the 7x50 router delays 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.

On 7x50 routers received RTC routes have no effect on the advertisement on MVPN-IPv4, MVPN-IPv6 and L2-VPN routes.

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, called the *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). When all RIB-OUT routes have been sent to a peer the MRAI timer associated with that session is started and when it expires the RIB-OUT changes that have accumulated while the timer was running trigger the sending of a new set of UPDATE messages to the peer.

It may be important 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 peer-level **min-route-advertisement** time and applies the minimum setting to routes belonging to specific address families; routes of other address families continue to be advertised according to the session-level MRAI setting. The address families that can be configured with **rapid-update** support are:

- L2-VPN
- MVPN-IPv4

- MVPN-IPv6
- MDT-SAFI
- EVPN

In many cases the default MRAI is appropriate for all address families (or at least those not included in the above list) when it applies to UPDATE messages that advertise reachable NLRI but it is less than ideal 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 and this is impeded by the effect of the MRAI timer at each router hop. SR-OS provides a solution for this problem by supporting a configuration command called **rapid-withdrawal**. 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, and this timer remains governed by the **min-route-advertisement** time or the **rapid-update** command, if it applies. When **rapid-withdrawal** is enabled it applies to all address families.

Advertise-Inactive

Standard BGP rules do not allow a BGP route to be advertised to peers unless it is the best path and it is ‘used’ locally. An IPv4 or IPv6 BGP route is considered ‘used’ if it is the *active* route to the destination in the route table. If there are multiple routes from different protocols for the same IP destination the BGP route is ‘used’ only if it has the numerically lowest route preference among all these routes; for further details refer to the section titled [BGP Route Installation in the Route Table on page 683](#).

In some cases it may be useful to advertise the best BGP path to peers despite the fact that it is *inactive* —i.e. because there are one or more lower-preference 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*.

As a global BGP configuration option the **advertise-inactive** command applies to all IPv4 and IPv6 routes and all sessions that advertise these routes. When the command is configured and the best BGP path is inactive it is automatically advertised to every peer unless rejected by a BGP export policy.

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 and IPv6 as options. *Best-External* for IPv4 applies to both regular IPv4 unicast routes as well as labeled-IPv4 (SAFI4) routes. Similarly, *Best-External* for IPv6 applies to both regular IPv6 unicast routes as well as 6PE (SAFI4) routes.

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 7x50 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.

Add-Paths

Add-Paths is a BGP enhancement that allows a BGP router to advertise multiple distinct paths for the same prefix/NLRI. This provides a number of potential benefits, including reduced routing churn, faster convergence, and better loadsharing.

In order 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 31](#).

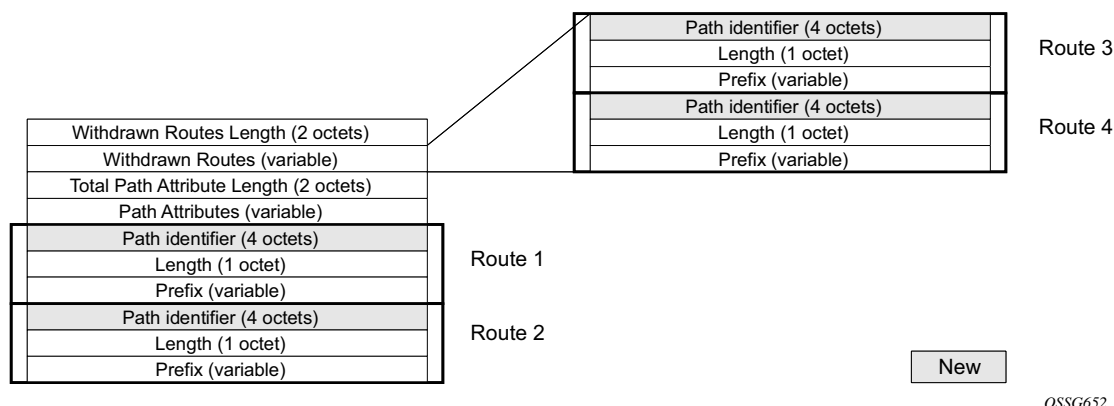


Figure 31: BGP Update Message with Path Identifier for IPv4 NLRI

Add-Paths is only supported by the base router BGP instance and the EBGp and IBGP sessions it forms with other *Add-Paths* capable peers. The ability to send and receive multiple paths per prefix is configurable per family, with the supported options being:

- IPv4 (including labeled IPv4 routes)
- VPN-IPv4
- IPv6 (including labeled IPv6 routes)
- VPN-IPv6

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 SROS 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.

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. 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) EBGPs 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 EBGPs peer back to different EBGPs peer in the same neighbor AS.

BGP Applications

SR OS implements the following BGP applications:

- [Next-hop Resolution Using Tunnels on page 701](#)
- [BGP Flow-Spec on page 705](#)

Next-hop Resolution Using Tunnels

The user enables the resolution of IPv4 prefixes using tunnels to BGP next-hops in TTM with the following command:

```
configure>router> bgp>next-hop-resolution
    shortcut-tunnel
        [no] family {ipv4}
            resolution {any|disabled|filter}
            resolution-filter
                [no] bgp
                [no] ldp
                [no] rsvp
            [no] disallow-igp
            exit
        exit
    exit
```

The **shortcut-tunnel** and **family** nodes are simply contexts to configure the binding of BGP unlabelled routes to tunnels.

The default resolution of a BGP unlabelled route is performed in RTM. The user must configure the **resolution** option to enable resolution to tunnels in TTM. If the **resolution** option is explicitly set to **disabled**, the binding to tunnel is removed and resolution resumes in RTM to IP next-hops.

If **resolution** is set to **any**, any supported tunnel type in BGP shortcut context will be selected following TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

The following tunnel types are supported in a BGP shortcut context and in order of preference: RSVP, LDP, and BGP.

- The **rsvp** value instructs BGP to search for 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 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.

Next-hop Resolution Using Tunnels

- The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.
- The **bgp** value instructs BGP to search for a BGP LSP with a RFC 107 label route prefix matching the address of the BGP next-hop.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

If **disallow-igp** is enabled, the BGP route will not be activated using IP next-hops in RTM if no tunnel next-hops are found in TTM.

The user enables the resolution of RFC 3107 BGP label route prefixes using tunnels to BGP next-hops in TTM with the following command:

```
configure>router> bgp>next-hop-resolution>
    label-route-transport-tunnel
        [no] family {ipv4, vpn}
            resolution {any|disabled|filter}
            resolution-filter
                [no] ldp
                [no] rsvp
            exit
        exit
    exit
```

The **label-route-transport-tunnel** and **family** nodes are simply contexts to configure the binding of IPv4 or IPv6 BGP labeled routes to tunnels.

This command provides a separate control for the different families of RFC 3107 BGP label routes: core IPv4 routes and inter-AS option B vpn-ipv4 and vpn-ipv6 routes at ASBR.

By default, core IPv4 routes and inter-AS option B VPN label routes resolve to LDP without the user needing to enter this command. IPv6 BGP labeled routes are currently resolving to IPv4 LDP tunnel only with the 6PE feature and do not require this command.

If 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 will be selected following TTM preference.

The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, and LDP.

- The **rsvp** value instructs BGP to search for 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 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** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding 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 will be selected again following the TTM preference.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

The user enables the resolution of vpn-ipv4 and vpn-ipv6 prefixes using tunnels to MP-BGP peers with the following command:

```
configure>service>vprn>
    auto-bind-tunnel
        resolution {any|disabled|filter}
        resolution-filter
        [no] gre
        [no] ldp
        [no] rsvp
        [no] sr-isis
        exit
    exit
```

The **auto-bind-tunnel** node is simply a context to configure the binding of VPRN routes to tunnels. The user must configure the **resolution** option to enable auto-bind resolution to tunnels in TTM. If the **resolution** option is explicitly set to disabled, the auto-binding to tunnel is removed.

If **resolution** is set to **any**, any supported tunnel type in VPRN context will be selected following TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.

- The **rsvp** value instructs BGP to search for 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 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** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

When the **sr-isis** value is enabled, a SR tunnel to the BGP next-hop is selected in the TTM from the lowest numbered IS-IS instance.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

When an explicit SDP to a BGP next-hop is configured in a VPRN service (**configure>service>vprn>spoke-sdp**), it overrides the **auto-bind-tunnel** selection for that BGP next-hop only. There is no support for reverting automatically to the **auto-bind-tunnel** selection if the explicit SDP goes down. The user must delete the explicit spoke-SDP in the VPRN service context to resume using the **auto-bind-tunnel** selection for the BGP next-hop.

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 15](#).

Table 15: Subcomponents of IPv4 Flow Route NLRI

Subcomponent Name [Type]	Value Encoding	SROS 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]	One or more (operator, value) pairs	No
Destination Port [5]	One or more (operator, value) pairs	Partial. No support for multiple ranges.
Source Port [6]	One or more (operator, value) pairs	Partial. No support for multiple ranges.
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	No
DSCP [11]	One or more (operator, value) pairs	Partial. Only a single value is supported.

Table 15: Subcomponents of IPv4 Flow Route NLRI (Continued)

Subcomponent Name [Type]	Value Encoding	SROS Support
Fragment [12]	One or more (operator, bitmask) pairs	Partial. No support for matching DF bit, first-fragment or last-fragment.

The NLRI of an IPv6 flow route can contain one or more of the subcomponents shown in [Table 16](#).

Table 16: Subcomponents of IPv6 Flow Route NLRI

Subcomponent Name [Type]	Value Encoding	SROS 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]	One or more (operator, value) pairs	No
Destination Port [5]	One or more (operator, value) pairs	Partial. No support for multiple ranges.
Source Port [6]	One or more (operator, value) pairs	Partial. No support for multiple ranges.
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	No
Traffic Class [11]	One or more (operator, value) pairs	Partial. Only a single value is supported.
Flow Label[13]	One or more (operator, value) pairs	No

Validating Received Flow Routes

Table 17 summarizes the actions that may be associated with an IPv4 or IPv6 flow route and how each type of action is encoded.

Table 17: IPv4 Flowspec Actions

Action	Encoding	SROS Support
Rate Limit	Extended community type 0x8006	Partial. Only rate=0 is supported.
Sample/Log	Extended community type 0x8007. S-bit	Yes
Next Entry	Extended community type 0x8007. T-bit	No
Redirect to VRF	Extended community type 0x8008.	Yes
Mark Traffic Class	Extended community type 0x8009.	No

IPv4 and IPv6 flow routes received from a BGP peer must be validated before they can be installed as filter entries. A flow route is considered invalid if:

1. The flow route is received from an EBGp peer and the left most AS number in the AS_PATH attribute does not equal the peer's AS number (from the **group/neighbor** configuration).
2. The **flowspec-validate** command is enabled, the flow route has a destination prefix subcomponent D, and the flow route was received from a peer that did not advertise the best route to D and all more-specific prefixes.

After received flow routes are validated they are processed by the relevant import policies, if applicable.



NOTE: A flow route never matches a prefix entry in a prefix-list, even if the destination IPv4 (or IPv6) prefix subcomponent or the source IPv4 (or IPv6) prefix subcomponent of the NLRI is a match.

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 and best the system attempts to construct an IPv4 or IPv6 filter entry from the NLRI contents and the action(s) encoded in the UPDATE message. If successful, the filter entry is added to the system-created 'fSpec-0' IPv4 or 'fSpec-0' IPv6 filter policy. The 'fSpec-0' IPv4 filter policy is applied to the following:

Configuration of TTL Propagation for BGP Label Routes

- Ingress IPv4 traffic on a network interface, if its configuration includes the **flowspec** command.
- Ingress IPv4 traffic on an IES SAP interface, if its configuration includes the **flowspec** command.
- Ingress IPv4 traffic on an IES spoke SDP interface, if its configuration includes the **flowspec** command.

Similarly the 'fSpec-0' IPv6 filter policy is applied to the following:

- Ingress IPv6 traffic on a network interface, if its configuration includes the **flowspec-ipv6** command.
- Ingress IPv6 traffic on an IES SAP interface, if its configuration includes the **flowspec-ipv6** command.
- Ingress IPv6 traffic on an IES spoke SDP interface, if its configuration includes the **flowspec-ipv6** command.

A user-defined filter policy can be applied to a base router interface that has flow-spec enabled. When an interface has both a user-defined filter policy and the system-created 'fSpec-0' filter policy, the filter rules are installed in the following order:

1. User-defined filter entries
2. Flow-spec entries (in order, determined by comparison of the NLRI described in RFC 5575).
3. User-defined filter default action.

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.

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 6PE, 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 that 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.

Note that 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.

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]`

Note that 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. Note 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 (6PE). As such the CLI command above has no impact on prefix packets forwarded in this context.

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.

7x50 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; at the current time, TCP-MD5 and other forms of session security are not supported. A 7x50 router can setup an RPKI-RTR session using the base routing table or the management router.

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 the prefix bits in the route match the prefix bits in the VRP entry (up to its min prefix length), AND the route prefix length is greater than or equal to the VRP entry min prefix length, AND the route prefix length is less than or equal to the VRP entry max prefix length, AND the origin AS of the route matches the origin AS of the VRP entry.
- A route is **covered** by a VRP entry if the prefix bits in the route match the prefix bits in the VRP entry (up to its min prefix length), AND the route prefix length is greater than or equal to the VRP entry min prefix length, AND 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**.

BGP Route Leaking

It is possible to leak a copy of a BGP route (including all its path attributes) from one routing instance to another in the same 7x50 system. This BGP route leaking capability applies to IPv4 and IPv6 routes (without labels). 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 or IPv6 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 or IPv6 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. There are separate leak-import policies for IPv4 and IPv6 routes and 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 note that 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. Note that 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 that a leaked route cannot be exported as a VPN-IP route and then re-imported into another local VPRN.

BGP Configuration Process Overview

Figure 32 displays the process to provision basic BGP parameters.

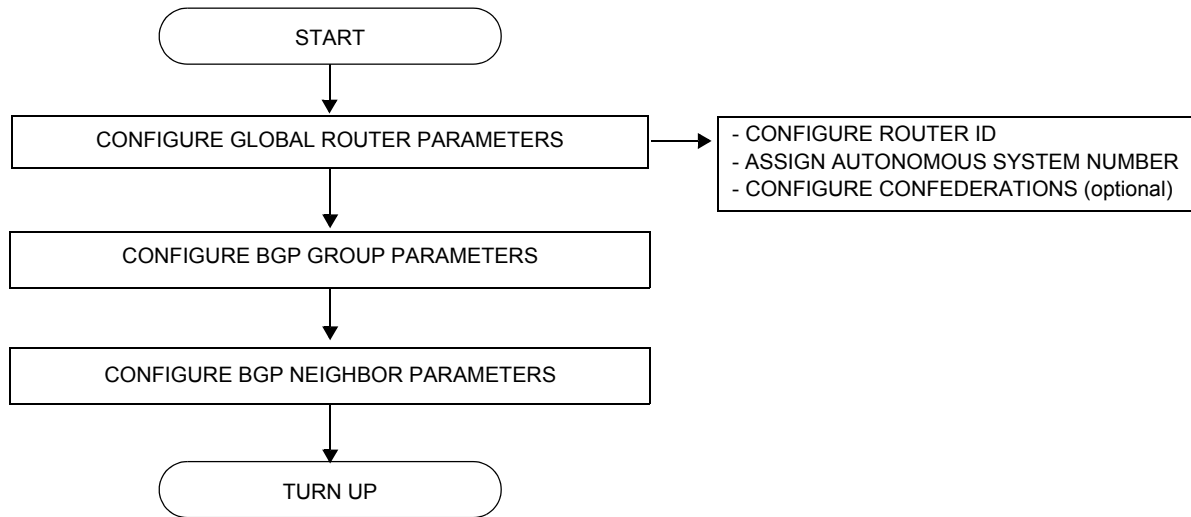


Figure 32: BGP Configuration and Implementation Flow

Configuration Notes

This section describes BGP configuration caveats.

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.
-

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 on page 717](#))
- 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.

BGP MIB Notes

The router implementation of the RFC 1657 MIB variables listed in [Table 18](#) differs from the IETF MIB specification.

Table 18: SR OS and IETF MIB Variations

MIB Variable	Description	RFC 1657 Allowed Values	Allowed Values
bgpPeerMinRouteAdvertisementInterval	Time interval in seconds for the MinRouteAdvertisementInterval timer. The suggested value for this timer is 30.	1 — 65535	^a 1 — 255

a. 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 19](#), there are three possible results:

Table 19: 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 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 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) 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 values as specified in [Table 18](#) and [Table 19](#), 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
```

Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

Topics in this section include:

- [BGP Configuration Overview on page 720](#)
 - [Preconfiguration Requirements on page 720](#)
 - [BGP Hierarchy on page 720](#)
 - [Internal and External BGP Configurations on page 720](#)
 - [BGP Confederations on page 728](#)
 - [BGP Route Reflectors on page 730](#)
- [Basic BGP Configuration on page 722](#)
- [Common Configuration Tasks on page 724](#)
 - [Creating an Autonomous System on page 725](#)
 - [Configuring a Router ID on page 726](#)
 - [BGP Components on page 732](#)
 - [Configuring Group Attributes on page 732](#)
 - [Configuring Neighbor Attributes on page 733](#)
 - [Configuring Route Reflection on page 734](#)
 - [Configuring a Confederation on page 735](#)
- [BGP Configuration Management Tasks on page 736](#)
 - [Modifying an AS Number on page 736](#)
 - [Modifying the BGP Router ID on page 737](#)
 - [Deleting a Neighbor on page 739](#)
 - [Deleting Groups on page 740](#)

BGP Configuration Overview

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.
-

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.

Internal and External BGP Configurations

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.

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 - Note that 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#

```

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:

1. Prepare a plan detailing the autonomous system(s), the router belonging to each group, group names, and peering connections.
2. Associate each router with an autonomous system number.
3. Configure each router with a router ID.
4. Associate each router with a peer group name.
5. Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
6. Specify neighbors.
7. Specify the autonomous system number associated with each neighbor.

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#
```

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.

Note that it is possible to configure an SR OS node 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.

Note that 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#
```

BGP Confederations

Follow these steps to configure a confederation:

1. Configure the autonomous system number of the confederation using the `confederation` command in the `config>router` context..
2. Configure the BGP confederation members using the `confederation` command in the `config>router` context.
3. Configure IBGP peering within the (local) sub-confederation.
4. Configure one or more confed-EBGP peerings to peers in other neighboring sub-confederations.

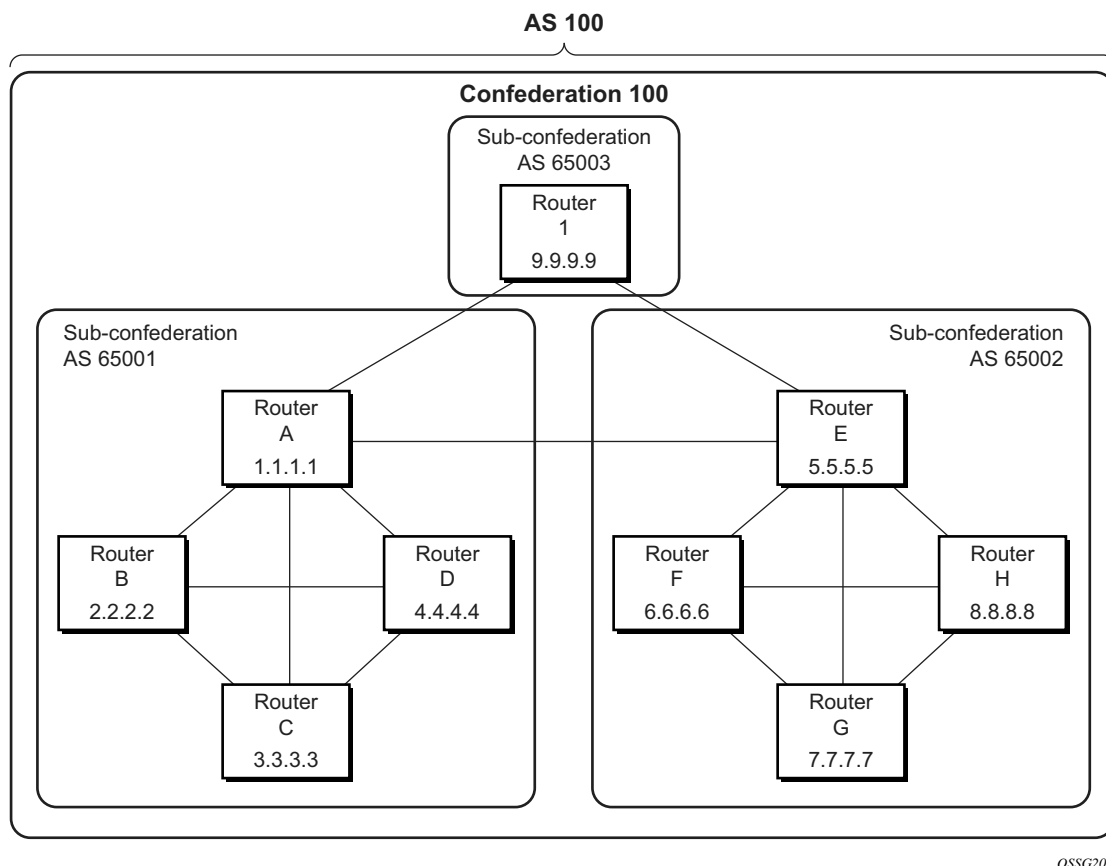


Figure 33: Confederation Network Diagram Example

The following configuration displays the minimum BGP configuration for routers in sub-confederation AS 65001 outlined in [Figure 34](#).

```

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

```

BGP Route 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.

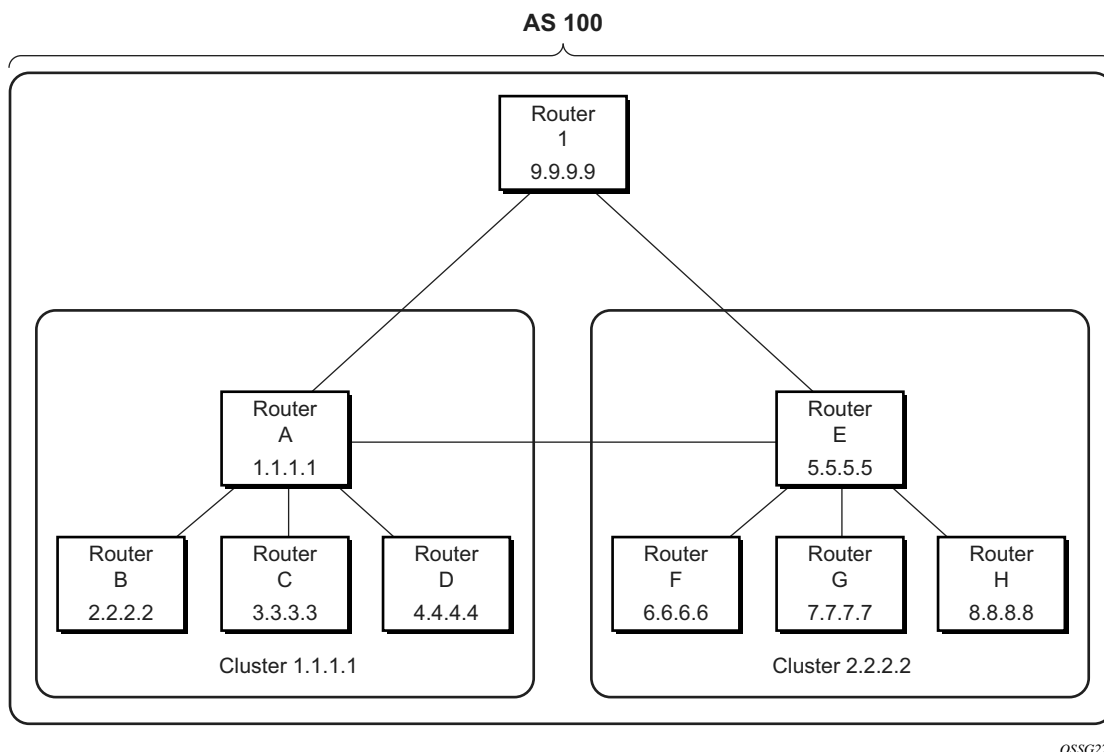


Figure 34: Route Reflection Network Diagram Example

The following configuration displays the minimum BGP configuration for routers in Cluster 1.1.1.1 outlined in [Figure 34](#).

```
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
```

BGP Components

Use the CLI syntax displayed below to configure the following BGP attributes:

- [BGP Components on page 732](#)
 - [Configuring Group Attributes on page 732](#)
 - [Configuring Neighbor Attributes on page 733](#)
 - [Configuring Route Reflection on page 734](#)
 - [Configuring a Confederation on page 735](#)
-

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#
```


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
          min-as-origination 15
          local-preference 170
          peer-as 10701
        exit
        neighbor 17.5.1.2
          hold-time 90
          keepalive 30
          min-as-origination 15
          local-preference 100
          min-route-advertisement 30
          preference 170
          peer-as 10702
        exit
      exit
    ...
-----
ALA-B>config>router>bgp#
```

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#
```

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.

NOTE: 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#
```

BGP Configuration Management Tasks

This section discusses the following BGP configuration management tasks:

- [Modifying an AS Number on page 736](#)
 - [Modifying a Confederation Number on page 737](#)
 - [Modifying the BGP Router ID on page 737](#)
 - [Modifying the Router-Level Router ID on page 738](#)
 - [Deleting a Neighbor on page 739](#)
 - [Deleting Groups on page 740](#)
-

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 system(s), 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`

Modifying a Confederation Number

Modifying a confederation number will cause BGP to restart automatically. Changes immediately take effect.

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#
```

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#
```

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#
```

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.
```

BGP Command Reference

Command Hierarchies

Configuration Commands

- [Global BGP Commands on page 741](#)
- [Group BGP Commands on page 744](#)
- [Neighbor BGP Commands on page 746](#)
- [Show Commands on page 749](#)
- [Clear Commands on page 750](#)
- [Debug Commands on page 750](#)

```

config
— router [router-name]
— confederation confed-as-num members as-number [as-number... (up to 15 max)]
— no confederation [confed-as-num members as-number [as-number... (up to 15 max)]]
— [no] mh-primary-interface interface-name
— [no] address {ip-address/mask | ip-address netmask}
— [no] description description-string
— [no] shutdown
— [no] mh-secondary-interface interface-name
— [no] address {ip-address/mask | ip-address netmask}
— [no] description description-string
— [no] shutdown
— [no] hold-time holdover-time
— [no] mh-secondary-interface
— router-id id-address
— no router-id
— [no] bgp
— [no] add-paths
— [no] ipv4 send send-limit receive [none]
— [no] ipv4 send send-limit
— [no] no ipv4
— [no] ipv6 send send-limit receive [none]
— [no] ipv6 end send-limit
— [no] no ipv6
— [no] vpn-ipv4 [send send-limit receive [none]]
— [no] vpn-ipv4 end send-limit
— [no] no vpn-ipv4
— [no] vpn-ipv6 send send-limit receive [none]
— [no] vpn-ipv6 end send-limit
— [no] no vpn-ipv6
— [no] advertise-external [ipv4] [ipv6]
— [no] advertise-inactive
— [no] aggregator-id-zero

```

```

— auth-keychain name
— authentication-key [authentication-key | hash-key] [hash | hash2]
— no authentication-key
— [no] backup-path [ipv4] [ipv6]
— best-path-selection
    — always-compare-med {zero | infinity}
    — always-compare-med strict-as {zero | infinity}
    — no always-compare-med
    — as-path-ignore [ipv4] [vpn-ipv4] [mcast-ipv4] [mvpn-ipv4] [l2-vpn]
    — no as-path-ignore
    — [no] compare-origin-validation-state
    — [no] deterministic-med
    — ignore-nh-metric
    — ignore-router-id
    — [no] origin-invalid-unusable
— [no] bfd-enable
— cluster cluster-id
— no cluster
— connect-retry seconds
— no connect-retry
— damp-peer-oscillations [idle-hold-time initial-wait second-wait max-wait] [error-inter-  
val minutes]
— [no] damping
— description description-string
— no description
— [no] disable-4byte-asn
— [no] disable-client-reflect
— disable-communities [standard] [extended]
— no disable-communities
— [no] disable-fast-external-failover
— [no] disable-route-table-install
— [no] enable-inter-as-vpn
— [no] enable-peer-tracking
— [no] enable-rr-vpn-forwarding
— error-handling
    — [no] update-fault-tolerance
— export policy-name [policy-name ... (up to 5 max)]
— no export
— family [ipv4] [vpn-ipv4] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mdt-safi] [ms-pw] [flow-  
ipv4] [route-target] [mcast-vpn-ipv4] [evpn] [mcast-ipv6]
— no family
— [no] flowspec-validate
— [no] graceful-restart
    — enable-notification
    — restart-time seconds
    — stale-routes-time time
    — no stale-routes-time
— hold-time seconds [min seconds2]
— no hold-time
— [no] ibgp-multipath
— import policy-name [policy-name ... (up to 5 max)]
— no import
— keepalive seconds
— no keepalive

```

```

— 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
— multihop tll-value
— no multihop
— multipath max-paths
— no multipath
— [no] mvpn-vrf-import-subtype-new
— next-hop-resolution
  — label-route-transport-tunnel
    — family {ipv4}
      — resolution {any | filter | disabled}
      — resolution-filter [ldp] [rsvp] [sr-sis]
  — policy policy-name
  — no policy
  — shortcut-tunnel
    — family {ipv4}
      — [no] disallow-igp
      — resolution {any | filter | disabled}
      — resolution-filter [bgp] [ldp] [rsvp] [sr-sis]
    — [no] use-bgp-routes
— [no] outbound-route-filtering
  — [no] extended-community
    — [no] accept-orf
    — send-orf [comm-id...(up to 32 max)]
    — no send-orf comm-id
— [no] path-mtu-discovery
— peer-tracking-policy policy-name
— preference preference
— no preference
— purge-timer minutes
— no purge-timer
— [no] rapid-update {[l2-vpn] [mvpn-ipv4] [mdt-safi][evpn]}
— [no] rapid-withdrawal
— [no] remove-private [limited] [skip-peer-as]
— route-target-list comm-id [comm-id...(up to 15 max)]
— no route-target-list [comm-id]
— router-id ip-address
— no router-id
— [no] shutdown
— [no] split-horizon
— transport-tunnel ldp | rsvp-te | mpls
— [no] third-party-nexthop
— [no] third-party-nexthop
— [no] vpn-apply-import

```

```

config
— router [router-name]
— [no] bgp
— [no] group name
— [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
— 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
— [no] aggregator-id-zero
— [no] aigp
— authentication-key [authentication-key | hash-key] [hash | hash2]
— no authentication-key
— auth-keychain name
— [no] bfd-enable
— 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
— description description-string
— no description
— [no] disable-4byte-asn
— [no] disable-client-reflect
— disable-communities [standard] [extended]
— no disable-communities
— [no] disable-fast-external-failover
— enable-origin-validation [ipv4] [ipv6]
— ebgp-link-bandwidth [ipv4] [ipv6] [vpn-ipv4] [vpn-ipv6]
— no enable-origin-validation
— [no] enable-peer-tracking
— error-handling
— [no] update-fault-tolerance
— export policy-name [policy-name...(up to 5 max)]
— no export
— family [ipv4] [vpn-ipv4] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [flow-ipv4]
[mdt-safi] [mcast-vpn-ipv4] [evpn] [mcast-ipv6]
— no family
— [no] flowspec-validate
— [no] relax-redirect-as-check
— [no] graceful-restart
— enable-notification

```

```

— restart-time seconds
— stale-routes-time time
— no stale-routes-time
— hold-time seconds [strict]
— no hold-time
— import policy-name [policy-name ...(up to 5 max)]
— no import
— keepalive seconds
— no keepalive
— local-address ip-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
— multihop ttl-value
— no multihop
— [no] next-hop-self {[ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4] [l2-vpn]} [multi-
homing primary-anycast secondary-anycast]
— [no] outbound-route-filtering
— [no] extended-community
— [no] accept-orf
— send-orf [comm-id...(up to 32 max)]
— 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] third-party-nexthop
— [no] vpn-apply-import

```

```

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
— 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-label [ipv4]
— [no] advertise-label
— [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
— 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
— description description-string
— no description
— [no] disable-4byte-asn
— [no] disable-client-reflect
— disable-communities [standard] [extended]
— no disable-communities
— [no] disable-fast-external-failover
— ebgp-link-bandwidth [ipv4] [ipv6] [vpn-ipv4] [vpn-ipv6]
— enable-origin-validation [ipv4] [ipv6]
— no enable-origin-validation
— [no] enable-peer-tracking
— error-handling
— [no] update-fault-tolerance
— export policy-name [policy-name...(up to 5 max)]
— no export
— family [ipv4] [vpn-ipv4] [mcast-ipv4] [l2-vpn] [mvpn-ipv4] [mvpn-
ipv6] [flow-ipv4] [mdt-safi] [route-target] [mcast-vpn-ipv4] [evpn]
[mcast-ipv6]
— no family

```

```

— [no] flowspec-validate
    — [no] relax-redirect-as-check
— [no] graceful-restart
    — enable-notification
    — restart-time seconds
    — stale-routes-time time
    — no stale-routes-time
— hold-time seconds [strict]
— no hold-time
— import policy-name [policy-name ... (up to 5 max)]
— no import
— keepalive seconds
— no keepalive
— local-address ip-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
— multihop ttl-value
— no multihop
— [no] next-hop-self
— [no] outbound-route-filtering
    — [no] extended-community
        — [no] accept-orf
            — send-orf [comm-id... (up to 32 max)]
            — 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] third-party-nexthop
— ttl-security min-ttl-value
— no ttl-security
— type {internal | external}
— no type
— [no] third-party-nexthop
— [no] vpn-apply-import

```

Command Hierarchies

Other BGP-Related Commands

config

- **router** *[router-name]*
 - **autonomous-system** *as-number*
 - **no autonomous-system**
 - **router-id** *ip-address*
 - **no router-id**

Show Commands

```

show
  — router [router-instance]
    — bgp
      — auth-keychain keychain-name
      — damping [damp-type] [detail]
      — damping [ip-prefix | prefix-length] [detail]
      — group [name] [detail]
      — neighbor [ip-address] [detail]
      — neighbor [as-number] [detail]
      — neighbor ip-address [family [type mvpn-type]] filter1 [brief]
      — neighbor ip-address [family] filter2
      — neighbor as-number [family] filter2
      — neighbor ip-address orf [filter3]
      — neighbor ip-address graceful-restart
      — next-hop [family] [ip-address] [detail]
      — paths
      — route-target
      — 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 mdt-safi [rd rd] [grp-address mcast-grp-address] [brief]
      — routes ms-pw [rd rd] [aii-type2 aii-type2] [brief]
      — routes flow-ipv4
        — policy-test policy-name family family prefix prefix/pfxlen [longer] [neighbor
          neighbor] [display-rejects] [detail]
      — summary [all]
      — summary [family family] [neighbor ip-address]
    — fib
    — fib
    — route-table

```

Command Hierarchies

Clear Commands

```
clear
  — router
    — bgp
      — damping [{prefix/ip-prefix-length] [neighbor ip-address]} | {group name}]
      — flap-statistics [{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]
      — neighbor {ip-address | as as-number | external | all} statistics
      — neighbor ip-address end-of-rib
      — protocol
```

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
```

Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

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.

description

Syntax	description <i>description-string</i> no description
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
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 is associated with the configuration context.
Parameters	<i>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, etc.), the entire string must be enclosed within double quotes.

BGP Commands

bgp

Syntax	[no] bgp
Context	config>router
Description	<p>This command creates the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation.</p> <p>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.</p>

add-paths

Syntax	[no] add-paths
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>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.</p>
Default	no add-paths

ipv4

Syntax	ipv4 send <i>send-limit</i> receive [none] ipv4 send <i>send-limit</i> no ipv4
Context	config>router>bgp>add-paths config>router>bgp>group>add-paths config>router>bgp>group>neighbor>add-paths
Description	<p>This command is used to configure the add-paths capability for IPv4 routes (including labeled IPv4 routes). By default, add-paths is not enabled for IPv4 routes.</p> <p>The maximum number of paths per IPv4 prefix 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. Entering the command without optional parameters negotiates the ability to both send and receive multiple paths per IPv4 prefix with each peer and configures the router to send the two best paths per prefix to each peer using the default Add-N, N=2 path selection algorithm.</p> <p>The no form of the command disables add-paths support for IPv4 routes, causing sessions established using add-paths for IPv4 to go down and come back up without the add-paths capability.</p>
Default	no ipv4
Parameters	<p>send <i>send-limit</i> — The maximum number of paths per IPv4 prefix 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).</p> <p>Values 1 — 16, none</p> <p>receive — The router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers</p> <p>none — The router does not negotiate the Add-Paths receive capability for VPN-IPv4 routes with its peers.</p>

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 is used to configure the add-paths capability for IPv6 routes (including 6PE routes). By default, add-paths is not enabled for IPv6 routes.

The maximum number of paths per IPv6 prefix 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 IPv6 routes, causing sessions established using add-paths for IPv6 to go down and come back up without the add-paths capability.

Default no ipv6

Parameters **send** *send-limit* — The maximum number of paths per IPv6 prefix 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 — 16, none

receive — The router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers

none — The router does not negotiate the Add-Paths receive capability for VPN-IPv6 routes with its peers.

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 is used to configure 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** — 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 — 16, none

receive — The router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers

none — The router does not negotiate the Add-Paths receive capability for VPN-IPv64 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	<p>This command is used to configure the add-paths capability for VPN-IPv6 routes. By default, add-paths is not enabled for VPN-IPv6 routes.</p> <p>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.</p> <p>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.</p>
Default	no vpn-ipv6
Parameters	<p><i>send-limit</i> — 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).</p> <p>Values 1 — 16, none</p> <p>receive — The router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers</p> <p>none — The router does not negotiate the add-paths receive capability for VPN-IPv6 routes with its peers.</p>

advertise-external

Syntax	[no] advertise-external [ipv4] [ipv6]
Context	config>router>bgp
Description	<p>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.</p> <p>The no form of the command disables Advertise Best External for the BGP family.</p>
Default	no advertise-external
Parameters	<p>ipv4 — Enable/disable best-external advertisement for all IPv4 (unicast and labeled-unicast) routes.</p> <p>ipv6 — Enable/disable best-external advertisement for all IPv6 (unicast and labeled-unicast) routes.</p>

advertise-inactive

Syntax	[no] advertise-inactive
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>The no form of the command disables the advertising of inactive BGP routers to other BGP peers.</p>
Default	no advertise-inactive

advertise-label

Syntax	advertise-label [ipv4] no advertise-label
Context	config>router>bgp>group>neighbor
Description	<p>This command configures the IPv4 transport peers to exchange prefixes using 6PE, as well as RFC 3107-labeled IPv4 routes.</p> <p>If IPv4 is enabled all IPv4 routes advertised to the remote BGP peer will be sent with an RFC 3107-formatted label for the destination route.</p> <p>The no form of the command disables any or all configured options.</p> <p>The command must include one or more of the options above.</p>
Default	no advertise-label
Parameters	ipv4 — Specifies the advertisement label address family for core IPv4 routes. This keyword can be specified only for an IPv4 peer.

aggregator-id-zero

Syntax	[no] aggregator-id-zero
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command is used to set 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.</p>

BGP Commands

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** — BGP adds the AS number and router ID to the aggregator path attribute.

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

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>vpn>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] [vpn-ipv4] [mcast-ipv4] [mvpn-ipv4] [l2-vpn]**
no as-path-ignore

Context config>router>bgp>best-path-selection
 config>service>vpn>bgp>best-path-selection

Description This command determines whether the AS path is used to determine the best BGP route. If this option is present, the AS paths of incoming routes are not used in the route selection process. 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 IPv4 routes.

vpn-ipv4 — Specifies that the length AS-path will be ignored for all IPv4 VPRN routes.

mcast-ipv4 — Specifies that the AS-path length will be ignored for all IPv4 multicast routes.

mvpn-ipv4 — Specifies that the AS-path length will be ignored for all mVPN IPv4 multicast 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

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 authentication <i>key</i> can be any combination of ASCII characters up to 255 characters long. The no form of the command reverts to the default value.
Default	MD5 Authentication is disabled by default.

- Parameters**
- authentication-key* — 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 (“ ”).
 - 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** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.
 - hash2** — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup-path

- Syntax** **[no] backup-path [ipv4]**
- Context** config>router
config>service>vprn
- 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 or a particular VPRN. 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 path(s) 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 path(s), 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. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs.
By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.
- Default** no backup-path
- Parameters** **ipv4** — enable the use of a backup path for BGP-learned IPv4 prefixes

best-path-selection

- Syntax** **best-path-selection**
- Context** config>router>bgp
- Description** This command enables path selection configuration.

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<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 no form of this command removes BFD from the associated IGP/BGP protocol adjacency.</p>
Default	no bfd-enable

cluster

Syntax	cluster <i>cluster-id</i> no cluster
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command configures the cluster ID for a route reflector server.</p> <p>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.</p> <p>When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.</p> <p>For redundancy, a cluster can have multiple route reflectors.</p> <p>Confederations can also be used to remove the full IBGP mesh requirement within an AS.</p> <p>The no form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.</p>
Default	no cluster — No cluster ID is defined.
Parameters	<i>cluster-id</i> — The route reflector cluster ID is expressed in dot decimal notation.
Values	Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

confederation

Syntax	confederation <i>confed-as-num</i> members <i>member-as-num</i> no confederation <i>confed-as-num</i> [members <i>member-as-num</i>]
Context	config>router
Description	<p>This command creates confederation autonomous systems within an AS.</p> <p>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.</p> <p>The no form of the command deletes the specified member AS from the confederation.</p> <p>When members are not specified in the no statement, the entire list is removed and confederations is disabled.</p> <p>When the last member of the list is removed, confederations is disabled.</p>
Default	no confederation — No confederations are defined.
Parameters	<p><i>confed-as-num</i> — The confederation AS number expressed as a decimal integer.</p> <p>Values 1 — 65535</p> <p>members <i>member-as-num</i> — 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>.</p>

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	<p>This command configures the BGP connect retry timer value in seconds.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to the default value.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	120 <i>seconds</i>
Parameters	<p><i>seconds</i> — The BGP Connect Retry timer value in seconds expressed as a decimal integer.</p> <p>Values 1 — 65535</p>

damp-peer-oscillations

Syntax	damp-peer-oscillations [<i>idle-hold-time initial-wait second-wait max-wait</i>] [<i>error-interval minutes</i>]
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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>.</p> <p>The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.</p>
Default	<i>no damp-peer-oscillations</i>
Parameters	<p><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.</p> <p>Values 0 — 2048</p> <p>Default 0</p> <p><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.</p> <p>Values 1 — 2048</p> <p>Default 5</p> <p><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.</p> <p>Values 1 — 2048</p> <p>Default 60</p> <p><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.</p> <p>Values 0 — 2048</p> <p>Default 30</p>

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** — Learned route damping is disabled.

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 peer(s).

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-client-reflect

Syntax	[no] disable-client-reflect
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command disables the reflection of routes by the route reflector to the clients in a specific group or neighbor.</p> <p>This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.</p> <p>The no form re-enables client reflection of routes.</p>
Default	no disable-client-reflect — Client routes are reflected to all client peers.

disable-communities

Syntax	disable-communities [standard] [extended] no disable-communities
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command configures BGP to disable sending communities.
Parameters	<p>standard — Specifies standard communities that existed before VPRNs or 2547.</p> <p>extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.</p>

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.

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.

ebgp-link-bandwidth

Syntax	ebgp-link-bandwidth <i>family</i> [<i>family</i> ... (up to 4 max)] no ebgp-link-bandwidth								
Context	config>router>bgp>group config>router>bgp>group>neighbor								
Description	<p>When the ebgp-link-bandwidth command is configured, BGP automatically adds a link-bandwidth extended community to every route (of the selected types) received from directly connected (single-hop) EBGp peers within the scope of the command.</p> <p>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.</p>								
Default	no ebgp-link-bandwidth — No link bandwidth extended community is automatically added to received BGP routes.								
Parameters	<i>family</i> — The BGP address family.								
Values	<table> <tr> <td><i>ipv4</i></td> <td>The command applies to IPv4 and label-IPv4 routes.</td> </tr> <tr> <td><i>ipv6</i></td> <td>The command applies to IPv6 and 6PE routes.</td> </tr> <tr> <td><i>vpn-ipv4</i></td> <td>The command applies to VPN-IPv4 routes.</td> </tr> <tr> <td><i>vpn-ipv6</i></td> <td>The command applies to VPN-IPv6 routes.</td> </tr> </table>	<i>ipv4</i>	The command applies to IPv4 and label-IPv4 routes.	<i>ipv6</i>	The command applies to IPv6 and 6PE routes.	<i>vpn-ipv4</i>	The command applies to VPN-IPv4 routes.	<i>vpn-ipv6</i>	The command applies to VPN-IPv6 routes.
<i>ipv4</i>	The command applies to IPv4 and label-IPv4 routes.								
<i>ipv6</i>	The command applies to IPv6 and 6PE routes.								
<i>vpn-ipv4</i>	The command applies to VPN-IPv4 routes.								
<i>vpn-ipv6</i>	The command applies to VPN-IPv6 routes.								

enable-origin-validation

Syntax	enable-origin-validation [ipv4] [ipv6] no enable-origin-validation
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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:</p> <ul style="list-style-type: none"> Valid (0) Not-Found (1) Invalid (2)

BGP Commands

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 6PE (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 IPv4 and label-IPv4 routes.

ipv6 — Enables origin validation processing for IPv6 and 6PE 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 reconvergence 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	<p>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 transport-tunnel command under config>router>bgp 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:</p> <ul style="list-style-type: none"> • from X (client) to Y (client) • from X (client) to Y (non-client) • from X (non-client) to Y (client) <p>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.</p> <p>Note that it is advised to leave this command disabled (for scaling and convergence reasons).</p>
Default	no enable-rr-vpn-forwarding

export

Syntax	export <i>policy-name</i> [<i>policy-name</i> ...] no export [<i>policy-name</i>]
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command specifies the export route policy used to determine which routes are advertised to peers.</p> <p>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.</p> <p>When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of fifteen (15) policy names can be configured. The first policy that matches is applied.</p> <p>When multiple export commands are issued, the last command entered overrides the previous command.</p> <p>When no export policies are specified, BGP routes are advertised and non-BGP routes are not advertised by default.</p> <p>The no form of the command removes the policy association with the BGP instance. To remove association of all policies, use the no export command without arguments.</p>

BGP Commands

Default **no export** — No export policy is specified. BGP routes are advertised and non-BGP routes are not advertised.

Parameters *policy-name* — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

family

Syntax **family** [ipv4] [vpn-ipv4][mcast-ipv4] [l2-vpn] [mvpn-ipv4] [flow-ipv4] [ms-pw] [route-target] [mcast-vpn-ipv4] [evpn]
no family

Context config>router>bgp
 config>router>bgp>group
 config>router>bgp>group>neighbor

Description This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.

Default **ipv4**

Parameters **evpn** — Exchanges Ethernet VPN routes using AFI 25 and SAFI 70.
ipv4 — Provisions support for IPv4 routing information.
vpn-ipv4 — Exchanges IPv4 VPN routing information.
mcast-ipv4 — Exchanges multicast IPv4 routing information.
l2-vpn — Exchanges Layer 2 VPN information.
mvpn-ipv4 — Exchanges Multicast VPN related information.
flow-ipv4 — Exchanges IPv4 flowspec routes belonging to AFI 1 and SAFI 133.
ms-pw — Exchanges dynamic MS-PW related information.
route-target — Exchanges RT constraint routes for VPN route filtering.
mcast-vpn-ipv4 — –Exchanges Multicast Routes in VPN using SAFI 129.
mcast-ipv6 — –Exchanges multicast IPv6 routing information.

flowspec-validate

Syntax	flowspec-validate no flowspec-validate
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command enables/disables validation of received flowspec routes. A flow route with a destination prefix subcomponent that is received from a particular peer is considered valid if and only if that peer also advertised the best unicast route to the destination prefix and any of its more-specific components. Also, when a flow route is received from an EBGp peer, the left most AS number in the AS_PATH attribute must equal the peer's AS number. If validation is enabled and a flowspec route is not valid, it is not eligible for import into the RIB, it is not used for filtering, a log/trap is generated, and it is not propagated to other flowspec peers.</p> <p>The no form of the command disables the validation procedure.</p>
Default	no flowspec-validate

route-target-list

Syntax	route-target-list <i>comm-id</i> [<i>comm-id</i> ..[up to 15 max]] no route-target-list [<i>comm-id</i>]
Context	config>router>bgp
Description	<p>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.</p> <p>The route-target-list is assigned at the global level and applies to all peers connected to the system.</p> <p>This command is only applicable if the router is a route-reflector server.</p> <p>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.</p>
Default	no route-target-list
Parameters	<i>comm-id</i> — Specifies the route target community in the form <0..65535>:<0..65535>

third-party-nexthop

Syntax	third-party-nexthop no third-party-nexthop
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>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.</p>
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	<p>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.</p> <p>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.</p>
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	<p>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.</p> <p>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.</p>

Default **no vpn-apply-import**

graceful-restart

Syntax **[no] graceful-restart**

Context config>router>bgp
 config>router>bgp>group
 config>router>bgp>group>neighbor

Description The command enables BGP graceful restart helper procedures (the “receiving router” role defined in the standard) for all received IPv4, IPv6, VPN-IPv4, and VPN-IPv6 routes. In order for helper mode to be available for a particular address family, both peers must signal GR support for the address family during capability negotiation.

When a neighbor covered by 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

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>update-error-handling
 config>router>bgp>group> update-error-handling
 config>router>bgp>group>neighbor> update-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 fault-tolerance

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

restart-time

Syntax	restart-time <i>seconds</i> 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 30 seconds.				
Default	no restart time				
Parameters	<i>seconds</i> — The restart-time that is advertised in the router's graceful-restart capability. <table><tr><td>Values</td><td>0 — 4095 seconds</td></tr><tr><td>Default</td><td>30</td></tr></table>	Values	0 — 4095 seconds	Default	30
Values	0 — 4095 seconds				
Default	30				

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 restart time

Parameters *time* — Specify the amount of time that stale routes should be maintained after a graceful restart is initiated.

Values 1 — 3600 seconds

group

Syntax [**no**] **group** *name*

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 peer groups are defined.

Parameters *name* — 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, etc.), the entire string must be enclosed within double quotes.

hold-time

Syntax **hold-time** *seconds* [**min** *seconds2*]
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.

BGP Commands

Default 90 seconds

Parameters *seconds* — 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 — 65535

seconds2 — 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**

ignore-nh-metric

Syntax **ignore-nh-metric**
no ignore-nh-metric

Context config>router>bgp>best-path-selection
config>service>vpn
config>service>vpn>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>vpn 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>vpn>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 no ignore-router-id
Context	config>router>bgp>best-path-selection config>service>vpn>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

origin-invalid-unusable

Syntax	origin-invalid-unusable no origin-invalid-unusable
Context	config>router>bgp>best-path-selection
Description	<p>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.</p> <p>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.</p>
Default	no origin-invalid-unusable

import

Syntax	import <i>policy-name</i> [<i>policy-name</i> ...] no import [<i>policy-name</i>]
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the config>router>policy-options context.</p> <p>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.</p> <p>When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of fifteen (15) policy names can be specified. The first policy that matches is applied.</p>

BGP Commands

When multiple **import** commands are issued, the last command entered will override 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 with the BGP instance. To remove association of all policies, use **no import** without arguments.

Default **no import** — No import policy specified (BGP routes are accepted).

Parameters *policy-name* — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

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 30 seconds

Parameters *seconds* — The keepalive timer in seconds expressed as a decimal integer.

Values 0 — 21845

local-address

Syntax	local-address <i>ip-address</i> no local-address
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>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.</p> <p>The no form of the command removes the configured local-address for BGP.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	<p>no local-address - The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.</p> <p><i>ip-address</i> — 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.</p> <p>Values ipv4-address: a.b.c.d (host bits must be 0)</p>

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	<p>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.</p> <p>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.</p> <p>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.</p>

BGP Commands

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 *as-number* — The virtual autonomous system number expressed as a decimal integer.

Values 1 — 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** *local-preference*
 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 — Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.
Parameters	<i>local-preference</i> — The local preference value to be used as the override value expressed as a decimal integer.
Values	0 — 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	<p>This command configures how the BGP peer session handles loop detection in the AS path.</p> <p>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.</p> <p>Note that dynamic configuration changes of loop-detect are not recognized.</p> <p>The no form of the command used at the global level reverts to default, which is loop-detect ignore-loop.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	loop-detect ignore-loop
Parameters	<p>drop-peer — Sends a notification to the remote peer and drops the session.</p> <p>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.</p> <p>ignore-loop — Ignores routes with loops in the AS path but maintains peering.</p> <p>off — Disables loop detection.</p>

mdt-safi

Syntax	[no] mdt-safi
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	This command enables peer capability to exchange MDT-SAFI address family advertisements.

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	<p>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.</p> <p>The specified value can be overridden by any value set via a route policy.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to default where the MED is not advertised.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	no med-out
Parameters	<p><i>number</i> — The MED path attribute value expressed as a decimal integer.</p> <p>Values 0 — 4294967295</p> <p>igp-cost — The MED is set to the IGP cost of the given IP prefix.</p>

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	<p>This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.</p> <p>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.</p> <p>The no form of the command used at the global level reverts to default.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	30 seconds
Parameters	<p><i>seconds</i> — The minimum route advertising interval, in seconds, expressed as a decimal integer.</p> <p>Values 1 — 255</p>

multihop

Syntax	multihop <i>ttl-value</i> no multihop
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGp peer multiple hops away.</p> <p>The no form of the command is used to convey to the BGP instance that the EBGp peers are directly connected.</p> <p>The no form of the command used at the global level reverts to default.</p> <p>The no form of the command used at the group level reverts to the value defined at the global level.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	<p>1 — EBGp peers are directly connected.</p> <p>64 — IBGP</p>
Parameters	<p><i>ttl-value</i> — The TTL value expressed as a decimal integer.</p> <p>Values 1 — 255</p>

multipath

Syntax	multipath <i>max-paths</i> no multipath
Context	config>router>bgp
Description	<p>This command enables BGP multipath.</p> <p>When multipath is enabled BGP load shares traffic across multiple links. Multipath can be configured to load share traffic across a maximum of 32 routes. If the equal cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.</p> <p>This configuration parameter is set at the global level (applies to all peers).</p> <p>Multipath is effectively disabled if the value is set to one. When multipath is disabled, and multiple equal cost routes are available, the route with the lowest next-hop IP address will be used.</p> <p>The no form of the command used at the global level reverts to default where multipath is disabled.</p>
Default	no multipath
Parameters	<p><i>max-paths</i> — The number of equal cost routes to use for multipath routing. If more equal cost routes exist than the configured value, routes with the lowest next-hop value are chosen. Setting this value to 1 disables multipath.</p> <p>Values 1 — 16</p>

mvpn-vrf-import-subtype-new

Syntax	[no] mvpn-vrf-import-subtype-new
Context	config>router>bgp
Description	<p>When enabled, the type/subtype in advertised routes is encoded as 0x010b.</p> <p>The no form of the command (the default) encodes the type/subtype as 0x010a (to preserve backwards compatibility).</p>

next-hop-resolution

Syntax	next-hop-resolution
Context	config>router>bgp
Description	This command enables the context to configure next-hop resolution parameters.

label-route-transport-tunnel

Syntax	label-route-transport-tunnel
Context	config>router>bgp>next-hop-res
Description	<p>This command enables the context to configure the resolution of RFC 3107 BGP label route prefixes using tunnels to BGP next-hops in TTM.</p> <p>The label-route-transport-tunnel and family nodes are simply contexts to configure the binding of IPv4 or IPv6 BGP labeled routes to tunnels.</p> <p>This command provides a separate control for the different families of RFC 3107 BGP label routes: core IPv4 routes, core IPv6 (6PE), and inter-AS option B vpn-ipv4 and vpn-ipv6 routes at ASBR.</p> <p>By default, core IPv4 routes and inter-AS option B VPN label routes resolve to LDP without the user needing to enter this command. IPv6 BGP labeled routes are currently resolving to IPv4 LDP tunnel only with the 6PE feature and do not require this command.</p> <p>If 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 will be selected following TTM preference.</p> <p>The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).</p> <p>The ldp value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.</p> <p>The rsvp value instructs BGP to search for 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 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.</p>

If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

When the **sr-isis** (**sr-ospf**) value is enabled, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF).

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

family

Syntax	family ipv4
Context	config>router>bgp>next-hop-resolution>label-route-transport-tunnel
Description	This command configures the address family for configuring the resolution of RFC 3107 BGP label routes using tunnels to BGP peers.
Parameters	ipv4 — selects the IPv4 address family for configuring the resolution of BGP label routes using tunnels to BGP peers.

resolution

Syntax	resolution {any filter disabled}
Context	config>router>bgp>next-hop-resolution>label-route-transport-tunnel>family
Description	This command configures the resolution mode in the resolution of BGP label routes using tunnels to BGP peers.
Parameters	<p>any — enables the binding to any supported tunnel type in BGP label route context following TTM preference.</p> <p>filter — enables the binding to the subset of tunnel types configured under resolution-filter.</p> <p>disabled — disables the resolution of BGP label routes using tunnels to BGP peers.</p>

resolution-filter

Syntax	resolution-filter [ldp] [rsvp] [sr-isis]
Context	config>router>bgp>next-hop-resolution>label-route-transport-tunnel>family
Description	<p>This command configures the subset of tunnel types which can be used in the resolution of BGP label routes using tunnels to BGP peers.</p> <p>The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).</p>
Parameters	ldp — selects the LDP tunnel type.

rsvp — selects the RSVP-TE tunnel type.

sr-isis — selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM.

policy

Syntax	policy <i>policy-name</i> no policy
Context	config>router>bgp>next-hop-res
Description	<p>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.</p> <p>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.</p>
Default	no policy
Parameters	<i>policy-name</i> — The 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, etc.), 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-res
Description	<p>This command enables the context to configure the resolution of BGP prefixes using tunnels to BGP next-hops in TTM.</p> <p>The shortcut-tunnel and family nodes are simply contexts to configure the binding of BGP unlabelled routes to tunnels.</p> <p>The default resolution of a BGP unlabelled route is performed in RTM. The user must configure the resolution option to enable resolution to tunnels in TTM. If the resolution option is explicitly set to disabled, the binding to tunnel is removed and resolution resumes in RTM to IP next-hops.</p> <p>If resolution is set to any, any supported tunnel type in BGP shortcut context will be selected following TTM preference. If one or more explicit tunnel types are specified using the resolution-filter option, then only these tunnel types will be selected again following the TTM preference.</p> <p>The following tunnel types are supported in a BGP shortcut context and in order of preference: RSVP, LDP, Segment Routing (SR), and BGP.</p>

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **bgp** value instructs BGP to search for a BGP LSP with a RFC 107 label route prefix matching the address of the BGP next-hop.

The **rsvp** value instructs BGP to search for 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 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.

When the **sr-isis (sr-ospf)** value is enabled, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

If **disallow-igp** is enabled, the BGP route will not be activated using IP next-hops in RTM if no tunnel next-hops are found in TTM.

family

Syntax	family ipv4
Context	config>router>bgp>next-hop-res>shortcut-tunnel
Description	This command configures the address family for configuring the resolution of BGP prefixes using tunnels to BGP peers.
Parameters	ipv4 — selects the IPv4 address family for configuring the resolution of BGP prefixes using tunnels to BGP peers.

resolution

Syntax	resolution {any filter disabled}
Context	config>router>bgp>next-hop-res>shortcut-tunnel>family
Description	This command configures the resolution mode in the resolution of BGP prefixes using tunnels to BGP peers.
Parameters	<p>any — enables the binding to any supported tunnel type in BGP shortcut context following TTM preference.</p> <p>filter — enables the binding to the subset of tunnel types configured under resolution-filter.</p> <p>disabled — disables the resolution of BGP prefixes using tunnels to BGP peers.</p>

resolution-filter

Syntax	resolution-filter [bgp] [ldp] [rsvp] [sr-isis]
Context	config>router>bgp>next-hop-res>shortcut-tunnel>family
Description	<p>This command configures the subset of tunnel types which can be used in the resolution of BGP label routes using tunnels to BGP peers.</p> <p>The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).</p>
Parameters	<p>bgp — selects the BGP label route tunnel type.</p> <p>ldp — selects the LDP tunnel type.</p> <p>rsvp — selects the RSVP-TE tunnel type.</p> <p>sr-isis — selects the Segment Routing (SR) tunnel type programmed by an ISIS instance in TTM.</p>

peer-tracking-policy

Syntax	peer-tracking-policy <i>policy-name</i> no peer-tracking-policy
Context	config>router>bgp config>service>vpn>bgp
Description	<p>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.</p> <p>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.</p>
Default	no peer-tracking-policy
Parameters	<i>policy-name</i> — The 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, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the con-fig>router>pol-icy-options context.

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 resolve BGP nexthop for IPv4 and IPv6 families on this router instance.
Default	no use-bgp-routes

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

extended-community

Syntax	[no] extended-community
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>The extended-community command opens the configuration tree for sending or accepting extended-community based BGP filters.</p> <p>In order for the no version of the command to work, all sub-commands (send-orf, accept-orf) must be removed first.</p>
Default	Community filtering is not enabled by default.

accept-orf

Syntax	[no] accept-orf
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>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.</p>
Default	Accepting ORFs is not enabled by default.

send-orf

Syntax	send-orf [<i>comm-id</i> ...(up to 32 max)] no send-orf [<i>comm-id</i>]
Context	config>router>bgp config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command instructs the router to negotiate the send capability in the BGP outbound route filtering (ORF) negotiation with a peer.</p> <p>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.</p> <p>The no form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer.</p> <p>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.</p> <p>If the <i>comm-id</i> parameter(s) 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> parameter(s) specified contain no route targets, then the router will not send an ORF.</p>
Default	no send-orf — Sending ORF is not enabled by default.
Parameters	<i>comm-id</i> — Any community policy which consists 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.

neighbor

Syntax	[no] neighbor <i>ip-address</i>
Context	config>router>bgp>group
Description	<p>This command creates a BGP peer/neighbor instance within the context of the BGP group.</p> <p>This command can be issued repeatedly to create multiple peers and their associated configuration.</p> <p>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.</p>
Default	No neighbors are defined.
Parameters	<p><i>ip-address</i> — The IP address of the BGP peer router in dotted decimal notation.</p> <p>Values ipv4-address: a.b.c.d (host bits must be 0)</p>

next-hop-self

Syntax	[no] next-hop-self {[ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4] [l2-vpn]} [multihoming <i>primary-anycast secondary-anycast</i>]
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.</p> <p>This is primarily used to avoid third-party route advertisements when connected to a multi-access network.</p> <p>In addition, this command can be used to enable and configure the multi-homing resiliency mechanism replacing the usual BGP nexthop with a configured anycast address.</p> <p>The no form of the command used at the group level allows third-party route advertisements in a multi-access network.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	no next-hop-self — Third-party route advertisements are allowed.
Parameters	<p>ipv4 — Provisions support for IPv4 routing information.</p> <p>vpn-ipv4 — Exchanges IPv4 VPN routing information.</p> <p>ipv6 — Exchanges IPv6 routing information.</p> <p>mcast-ipv4 — Exchanges multicast IPv4 routing information.</p> <p>l2-vpn — Exchanges Layer 2 VPN information.</p> <p><i>primary-anycast</i> — Specifies the anycast address that the local node will use to replace the BGP nexthop address in route updates associated peers.</p>

secondary-address — Specifies the anycast address that the local node is to track.

passive

Syntax	[no] passive
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>Enables/disables passive mode for the BGP group or neighbor.</p> <p>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.</p> <p>The no form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	no passive — BGP will actively try to connect to all the configured peers.

peer-as

Syntax	peer-as <i>as-number</i>
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.</p> <p>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.</p> <p>For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.</p> <p>This is required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.</p>
Default	No AS numbers are defined.
Parameters	<i>as-number</i> — The autonomous system number expressed as a decimal integer.
Values	1 — 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 *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 the configured peer(s). 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 170

Parameters *preference* — The route preference expressed as a decimal integer.

Values 1 — 255

purge-timer

Syntax	purge-timer <i>minutes</i> no purge-timer
Context	config>router>bgp
Description	<p>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.</p> <p>The BGP purge timer configures the time before stale routes are purged.</p> <p>The no form of the command reverts to the default.</p>
Default	10
Parameters	<i>minutes</i> — Specifies the maximum time before stale routes are purged.
	Values 1 — 60

rapid-update

Syntax	rapid-update {[l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [evpn]} no rapid-update { [l2-vpn] [mvpn-ipv4] [mvpn-ipv6] [mdt-safi] [evpn]}
Context	config>router>bgp
Description	This command enables and disables BGP rapid update for specified address-families. When no parameter is given for the no rapid-update statement, rapid update is disabled for all address-families.
Default	no rapid-update
Parameters	<p>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.</p> <p>mvpn-ipv4 — Specifies BGP rapid update for the mvpn-ipv4 address family. The mvpn-pv4 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.</p> <p>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.</p> <p>mvpn-ipv6 — Specifies BGP rapid update for the mvpn-ipv6 address family.</p> <p>evpn — Specifies BGP rapid update for the evpn address family.</p>

rapid-withdrawal

Syntax	[no] rapid-withdrawal
Context	config>router>bgp
Description	<p>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.</p> <p>The no form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.</p>
Default	no rapid-withdrawal

prefix-limit

Syntax	prefix-limit <i>family limit</i> [log-only] [threshold <i>percentage</i>] [idle-timeout {<i>minutes</i> forever}] [post-import] no prefix-limit <i>family</i>
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>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.</p> <p>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.</p> <p>The no form of the command removes the prefix-limit.</p>
Default	No prefix limits for any address family.
Parameters	<p>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.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p>Values 1 — 100</p> <p><i>family</i> — The address family to which the limit applies.</p> <p>Values ipv4 vpn-ipv4 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</p> <p><i>limit</i> — The number of routes that can be learned from a peer expressed as a decimal integer.</p> <p>Values 1 — 4294967295</p> <p><i>minutes</i> — Specifies duration in minutes before automatically re-establishing a session.</p> <p>Values 1 — 1024</p>

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	<p>This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.</p> <p>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.</p> <p>The router software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.</p> <p>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.</p>
Parameters	<p>limited — This optional keyword removes private ASNs up to the first public ASN encountered. It then stops removing private ASNs.</p> <p>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.</p>

router-id

Syntax	router-id <i>ip-address</i> no router-id
Context	config>router>bgp
Description	<p>This command specifies the router ID to be used with this BGP instance.</p> <p>Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID. The router ID must be set to a valid host address.</p> <p>It is possible to configure an SR OS node 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.</p>
Default	No router-id is configured for BGP by default. The system interface IP address is used.

Parameters *ip-address* — The router ID expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. It is highly recommended that this address be the system IP 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**

transport-tunnel

Syntax **transport-tunnel** *ldp* | *rsvp-te* | *mpls*

Context config>router>bgp

Description This command selects the transport LSP option to provide model B or C connectivity.

 The **no** form of the command defaults to LDP as transport LSP method for model B or C connectivity.

Default transport-tunnel ldp

Parameters *ldp* — Allows LDP-based LSPs to be used as transport from the ASBR to local PE routers.

rsvp-te — Allows RSVP-TE based LSPs to be used as transport from the ASBR to local PE routers.

mpls — Specifies that both LDP and RSVP-TE can be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance.

ttl-security

Syntax	ttl-security <i>min-ttl-value</i> no ttl-security
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP/LDP 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.</p> <p>The no form of the command disables TTL security.</p>
Parameters	<i>min-ttl-value</i> — Specify the minimum TTL value for an incoming packet.
Values	1 — 255
Default	1

type

Syntax	[no] type { internal external }
Context	config>router>bgp>group config>router>bgp>group>neighbor
Description	<p>This command designates the BGP peer as type internal or external.</p> <p>The type of internal indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.</p> <p>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.</p> <p>The no form of the command used at the group level reverts to the default value.</p> <p>The no form of the command used at the neighbor level reverts to the value defined at the group level.</p>
Default	no type — Type of neighbor is derived on the local AS specified.
Parameters	internal — Configures the peer as internal. external — Configures the peer as external.

Other BGP-Related Commands

autonomous-system

Syntax	autonomous-system <i>autonomous-system</i> no autonomous-system
Context	config>router
Description	<p>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.</p> <p>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.</p>
Default	No autonomous system number is defined.
Parameters	<i>as-number</i> — The autonomous system number expressed as a decimal integer.
Values	1 — 4294967295

mh-primary-interface

Syntax	mh-primary-interface <i>interface-name</i> no mh-primary-interface
Context	config>router
Description	<p>This command creates a loopback interface for the use in multihoming resiliency. Once active this interface can be used to advertise reachability information to the rest of the network using the primary address which is backed up by the secondary</p> <p>This reachability for this address is advertised via IGP and LDP protocols to allow the resolution of BGP routes advertised with this address.</p> <p>The no form of the command disables this setting.</p>
Default	no mh-primary-interface

mh-secondary-interface

Syntax	mh-secondary-interface <i>interface-name</i> no mh-secondary-interface
Context	config>router
Description	<p>This command creates a loopback interface for the use in multihoming resiliency. This address is considered the secondary multihoming address and is only used to resolve routes advertised by the primary router in the event that router becomes unavailable. For this purpose, the Reachability for this address is advertised via IGPs and LDP protocols to allow the resolution of BGP routes advertised with this address by the primary multihoming router.</p> <p>The no form of the command disables this setting.</p>
Default	no mh-secondary-interface

address

Syntax	address { <i>ip-address/mask</i> <i>ip-address netmask</i> } no address
Context	config>router>mh-primary-interface config>router>mh-secondary-interface
Description	<p>This command assigns an IP address, IP subnet, and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface.</p> <p>An IP address must be assigned to each IP interface for the interface to be active. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.</p> <p>The local subnet that the address command defines must not be part of the services address space within the routing context by use of the config router service-prefix command. Once a portion of the address space is allocated as a service prefix, that portion is not available to IP interfaces for network core connectivity.</p> <p>The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. Show commands display CIDR notation and are stored in configuration files.</p> <p>By default, no IP address or subnet association exists on an IP interface until it is explicitly created.</p> <p>The no form of the command removes the IP address assignment from the IP interface. Interface specific configurations for IGP protocols like OSPF are also removed. The no form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface will operationally stop any protocol interfaces or MPLS LSPs that explicitly reference that IP address. When a new IP address is defined, the IP interface can be administratively enabled (no shutdown), which reinitializes the protocol interfaces and MPLS LSPs associated with that IP interface.</p> <p>If a new address is entered while another address is still active, the new address will be rejected.</p>
Default	no address

Parameters	<p><i>ip-address</i> — The IP address of the IP interface. The ip-addr portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.</p> <p>Values 1.0.0.0 — 223.255.255.255</p> <p>/ — The forward slash is a parameter delimiter that separates the ip-addr portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the ipaddr, the “/” and the mask-length parameter. If a forward slash does not ediatly follow the ipaddr, a dotted decimal mask must follow the prefix.</p> <p><i>mask-length</i> — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the ip-addr from the masklength parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1— 32. Note that a mask length of 32 is reserved for system IP addresses.</p> <p>Values 1— 3</p> <p><i>mask</i> — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the ip-addr from a traditional dotted decimal mask. The mask parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Note that a mask of 255.255.255.255 is reserved for system IP addresses.</p> <p>Values 128.0.0.0 — 255.255.255.255</p> <p><i>net-mask</i> — he subnet mask in dotted decimal notation.</p> <p>Values 0.0.0.0 — 223.255.255.255 (network bits all 1 and host bits all 0)</p>
-------------------	---

description

Syntax	<p>description <i>description-string</i></p> <p>no description</p>
Context	<p>config>router>mh-primary-interface</p> <p>config>router>mh-secondary-interface</p>
Description	<p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The no form of the command removes the description string from the context.</p>
Default	no description
Parameters	<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, etc.), the entire string must be enclosed within double quotes.</p>

shutdown

Syntax	shutdown no shutdown
Context	config>router>mh-primary-interface config>router>mh-secondary-interface
Description	<p>The shutdown 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, shutdown and no shutdown are always indicated in system generated configuration files.</p> <p>The no form of the command puts an entity into the administratively enabled state.</p>
Default	no shutdown

hold-time

Syntax	hold-time <i>holdover-time</i> no hold-time				
Context	config>router>mh-secondary-interface				
Description	<p>The optional hold-time parameter is only applicable for the secondary context and specifies how long label information learned about the secondary anycast address should be kept after that peer is declared down. This timer should be set to a value large enough for the remainder of the network to detect the failure and complete the reconvergence process.</p> <p>The no form of the command resets the hold-time back to the default value.</p>				
Default	no hold-time				
Parameters	<p><i>holdover-time</i> — (seconds) specifies the number of seconds the router should hold label information learned from the alternate router in it's secondary label table. This is to allow the reset of the network to reconverge after a router failure before the anycast based label assignments are flushed from the forwarding plane.</p> <table><tr><td>Values</td><td>0—65535</td></tr><tr><td>Default</td><td>90</td></tr></table>	Values	0—65535	Default	90
Values	0—65535				
Default	90				

router-id

Syntax **router-id** *ip-address*
 no router-id

Context config>router

Description This command configures the router ID for the router instance.

The router ID is used by OSPF 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 **no** form of the command reverts to the default value.

Default 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.

Parameters *router-id* — The 32 bit router ID expressed in dotted decimal notation or as a decimal value.

Show Commands

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
Values	router-name: Base, management service-id: 1 — 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 particular authentication key-chain.
Parameters	<i>keychain</i> — Specifies an existing keychain name.

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
30.1.1.0.2             1          eta_keychain1
```

Show Commands

```
=====
*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
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
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              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
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         : Disabled
Advertise Label       : None
Auth key chain        : testname
Local Capability      : RouteRefresh MP-BGP
Remote Capability     :
Import Policy         : None Specified / Inherited

```

```
-----
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** [*damp-type*] [**detail**]
damping [*ip-prefix* | *prefix-length*] [**detail**]

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 *ip-prefix* — 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 — 32

damp-type — 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.

Output Damping Output Fields — The following table describes BGP damping output fields.

Label	Description
BGP Router ID	The local BGP router ID.
The local BGP router ID.	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
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.
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 length of time in hour/minute/second (HH:MM:SS) format.
Last update	The time when BGP was updated last in day/hour/minute (DD:HH:MM) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of route 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.

Sample Output

```

A:ALA-12# show router 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  12.149.7.0/24       10.0.28.1     00h00m00s      60203 65001 19855 3356
                        1239  22406
si    24.155.6.0/23      10.0.28.1     00h43m41s      60203 65001 19855 3356
                        2914  7459
si    24.155.8.0/22      10.0.28.1     00h38m31s      60203 65001 19855 3356
                        2914  7459
si    24.155.12.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    24.155.22.0/23     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    24.155.24.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914  7459
si    24.155.28.0/22     10.0.28.1     00h34m31s      60203 65001 19855 3356
                        2914  7459
si    24.155.40.0/21     10.0.28.1     00h28m24s      60203 65001 19855 3356
                        7911  7459
si    24.155.48.0/20     10.0.28.1     00h28m24s      60203 65001 19855 3356
                        7911  7459
ud*i  61.8.140.0/24       10.0.28.1     00h00m00s      60203 65001 19855 3356
                        4637  17447
ud*i  61.8.141.0/24      10.0.28.1     00h00m00s      60203 65001 19855 3356
                        4637  17447
ud*i  61.9.0.0/18        10.0.28.1     00h00m00s      60203 65001 19855 3356
                        3561  9658  6163
. . .
ud*i  62.213.184.0/23  10.0.28.1     00h00m00s      60203 65001 19855 3356
                        6774  6774  9154
=====
A:ALA-12#

```

Show Commands

```
A:ALA-12# show router 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 : 12.149.7.0/24
-----
Network          : 12.149.7.0/24          Peer          : 10.0.28.1
NextHop          : 10.0.28.1              Reuse time    : 00h00m00s
Peer AS         : 60203                  Peer Router-Id : 32.32.27.203
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 : 15.142.48.0/20
-----
Network          : 15.142.48.0/20        Peer          : 10.0.28.1
NextHop          : 10.0.28.1              Reuse time    : 00h00m00s
Peer AS         : 60203                  Peer Router-Id : 32.32.27.203
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 : 15.200.128.0/19
-----
Network          : 15.200.128.0/19       Peer          : 10.0.28.1
NextHop          : 10.0.28.1              Reuse time    : 00h00m00s
Peer AS         : 60203                  Peer Router-Id : 32.32.27.203
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 : 15.203.192.0/18
-----
Network          : 15.203.192.0/18       Peer          : 10.0.28.1
NextHop          : 10.0.28.1              Reuse time    : 00h00m00s
Peer AS         : 60203                  Peer Router-Id : 32.32.27.203
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 bgp damping 15.203.192.0/18 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 15.203.192.0/18
=====
Network : 15.203.192.0/18
-----
Network      : 15.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h00m00s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
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 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 : 15.142.48.0/20
-----
Network      : 15.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
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 : 15.200.128.0/19
-----
Network      : 15.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time : 00h29m22s
Peer AS      : 60203               Peer Router-Id : 32.32.27.203
Local Pref   : none

```

Show Commands

```
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 : 15.203.240.0/20
-----
Network      : 15.203.240.0/20      Peer           : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time     : 00h29m22s
Peer AS      : 60203               Peer Router-Id  : 32.32.27.203
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 : 15.206.0.0/17
-----
Network      : 15.206.0.0/17      Peer           : 10.0.28.1
NextHop      : 10.0.28.1           Reuse time     : 00h29m22s
Peer AS      : 60203               Peer Router-Id  : 32.32.27.203
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#
```

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.

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.
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 — 4294967295 — The maximum number of routes BGP can learn from a peer.

Label	Description (Continued)
Passive	<p>Disabled — BGP attempts to establish a BGP connection with neighbor in the specified peer group.</p> <p>Enabled — BGP will not actively attempt to establish a BGP connection with neighbor in the specified peer group.</p>
Next Hop Self	<p>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.</p> <p>Enabled — BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.</p>
Aggregator ID 0	<p>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.</p> <p>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.</p>
Remove Private	<p>Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.</p> <p>Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.</p>
Damping	<p>Disabled — The peer group is configured not to dampen route flaps.</p> <p>Enabled — The peer group is configured to dampen route flaps.</p>
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.
Cluster Id	<p>The configured route reflector cluster ID.</p> <p>None — No cluster ID has been configured</p>
Client Reflect	<p>Disabled — The BGP route reflector will not reflect routes to this neighbor.</p> <p>Enabled — The BGP route reflector is configured to reflect routes to this neighbor.</p>
NLRI	<p>The type of NLRI information that the specified peer group can accept.</p> <p>Unicast — IPv4 unicast routing information can be carried.</p>
Preference	The configured route preference value for the peer group.

Label	Description (Continued)
List of Peers	A list of 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           State           : Up
Peer AS          : 40000             Local AS         : 65206
Local Address     : n/a              Loop Detect      : Ignore
Export Policy     : direct2bgp
Hold Time        : 90                Keep Alive       : 30
Cluster Id       : None              Client Reflect   : Enabled
NLRI             : Unicast           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           State           : Up
Peer AS          : 40000             Local AS         : 65206
Local Address     : n/a              Loop Detect      : Ignore
Connect Retry    : 20                Authentication   : None
Local Pref       : 100                MED Out         : 0
Multihop         : 0 (Default)
Min Route Advt.  : 30                Min AS Originate : 15
Prefix Limit     : No Limit          Passive          : Disabled
Next Hop Self    : Disabled          Aggregator ID 0  : Disabled
Remove Private   : Disabled          Damping          : Disabled
Export Policy     : direct2bgp
```

Show Commands

```
Hold Time      : 90                      Keep Alive      : 30
Cluster Id     : None                    Client Reflect    : Enabled
NLRI           : Unicast                  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                      State           : Up
Peer AS            : n/a                          Local AS          : 100
Local Address      : n/a                          Loop Detect       : Ignore
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
```

```
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#
```

neighbor

Syntax

```
neighbor [ip-address [detail]]
neighbor [as-number [detail]]
neighbor ip-address [family [type mvpn-type]] filter1 [brief]
neighbor ip-address [family] filter2
neighbor as-number [family] filter2
neighbor ip-address orf [filter3]
neighbor ip-address graceful-restart
```

Context show>router>bgp

Description This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS displays.

When either **received-routes** or **advertised-routes** is specified, then the routes received from or sent to the specified peer is listed (see second output example).

Note: This information is not available by SNMP.

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.

Parameters

ip-address — Display information for the specified IP address.

Values

- ipv4-address: a.b.c.d (host bits must be 0)
- interface: 32 characters maximum, mandatory for link local addresses.

as-number — Display information for the specified AS number.

Values 1 — 65535

family — Specify the type of routing information to be distributed by this peer group.

Values

- evpn** — Displays the BGP Ethernet VPN routes.
- 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 content of the multicast routing table.
- mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.

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 by this peer.
- history** — Displays statistics for dampened routes.
- suppressed** — Displays the number of paths from this peer that have been suppressed by

damping.

detail — Displays detailed information pertaining to *filter1*.

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.

detail — Displays detailed information pertaining to *filter2*

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.

brief — Displays information in a brief format. This parameter is only supported with received-routes and advertised-routes.

orf — Displays outbound route filtering for the BGP instance. ORF (Outbound Route Filtering) 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/ignored.

graceful-restart — Displays neighbors configured for graceful restart.

Output

Standard and Detailed Neighbor — The following table describes the standard and detailed command output fields for a BGP neighbor.

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.
State	Idle — The BGP peer is not accepting connections.

Label	Description (Continued)
Last State	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.

Label	Description (Continued)
	None — No events have occurred.
Last Error	Displays the last BGP error and subcode to occur on the BGP neighbor.
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.

Label	Description (Continued)
	Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
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 — 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.

Label	Description (Continued)
o/p Messages	Total number of packets sent to the BGP neighbor.
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      : 300          Peer Port      : 0
Peer Address : 10.0.0.5
Local AS     : 200          Local Port    : 0
Local Address : 10.0.0.104
Peer Type    : External
State        : Active       Last State    : Idle
Last Event   : stop
Last Error   : Cease
Local Family : IPv4
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      : 100          Peer Port      : 0
Peer Address : 10.0.0.91
Local AS     : 200          Local 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                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           : 1
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    : Disabled          Stale Routes Time: n/a
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 Port	: 49725
Peer Address	: 10.20.1.3		
Local AS	: 100	Local Port	: 179
Local Address	: 10.20.1.2		
Peer Type	: Internal		
State	: Established	Last State	: Established
Last Event	: rcvKeepAlive		
Last Error	: Cease		
Local Family	: IPv4		
Remote Family	: IPv4		
Hold Time	: 3	Keep Alive	: 1
Active Hold Time	: 3	Active Keep Alive	: 1
Cluster Id	: None		
Preference	: 170	Num of Flaps	: 0
Recd. Paths	: 1		
IPv4 Recd. Prefixes	: 11	IPv4 Active Prefixes	: 10
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		
Input Queue	: 0	Output Queue	: 0
i/p Messages	: 471	o/p Messages	: 473
i/p Octets	: 3241	o/p Octets	: 3241
i/p Updates	: 4	o/p Updates	: 4
TTL Security	: Disabled	Min TTL Value	: n/a
Advertise Inactive	: Disabled	Peer Tracking	: Disabled
Advertise Label	: None		
Auth key chain	: eta_keychain1		
Local Capability	: RouteRefresh MP-BGP		
Remote Capability	: RouteRefresh MP-BGP		
Import Policy	: None Specified / Inherited		
Export Policy	: static2bgp		

Show Commands

```
-----
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             Active Prefixes : 0
Recd. Paths    : 0             Suppressed Paths : 0
```

```

Input Queue      : 0
i/p Messages    : 0
i/p Octets      : 0
i/p Updates     : 0
Export Policy    : direct2bgp
=====
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 Address  : 3.3.3.3
Local AS     : 100
Local Address : 0.0.0.0
Peer Type    : Internal
State        : Active
Last Event   : stop
Last Error   : Cease
Local Family  : IPv4
Remote Family : Unused
Hold Time    : 10
Active Hold Time : 0
Cluster Id   : 2.2.3.4
Preference   : 101
Recd. Paths  : 0
IPv4 Recd. Prefixes : 0
IPv4 Suppressed Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0
Mc IPv4 Recd. Pfxs : 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 : 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

```

Show Commands

```
=====
*A:vRR>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                 Backup Vpn IPv6 Pfxs : 0
Input Queue          : 0                 Output 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

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

IPv6 Rej. Pfxs    : 0
VPN-IPv6 Rej. Pfxs : 0
Mc IPv6 Rej. Pfxs  : 0
MVPN-IPv6 Rej. Pfxs : 0
Flow-IPv6 Rej. Pfxs : 0
MDT-SAFI Rej. Pfxs : 0
MS-PW Rej. Pfxs   : 0
Evpn Rej. Pfxs    : 0

```

```

=====
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 Ouput — The following table describes the command output for both the standard and detailed information for a neighbor.

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:ALA-12# show router bgp neighbor 10.0.0.16 received-routes
=====
BGP Router ID : 10.0.0.16          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 Neighbor
=====
Flag  Network          Nexthop          LocalPref  MED      As-Path
-----
?    10.0.0.16/32        10.0.0.16       100        none     No As-Path
?    10.0.6.0/24         10.0.0.16       100        none     No As-Path
?    10.0.8.0/24         10.0.0.16       100        none     No As-Path
```



```

?      10.0.12.0/24      10.0.0.16      100      none      No As-Path
?      10.0.13.0/24      10.0.0.16      100      none      No As-Path
?      10.0.204.0/24     10.0.0.16      100      none      No As-Path
=====
A:ALA-12#

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] [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** — 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.
 - mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.
 - ip-address** — Displays the next hop information for the specified IP address.
 - Values**
 - ipv4-address: a.b.c.d (host bits must be 0)
 - 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.

Label	Description
BGP ID	The local BGP router ID.
AS	The configured autonomous system number.

Label	Description (Continued)
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.
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
```

paths

Syntax paths**Context** show>router>bgp**Description** This command displays a summary of BGP path attributes.**Output** **Show Path Output** — The following table describes the command output fields for a BGP path.

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 nexthop.
MED	The Multi-Exit Discriminator value.
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

=====

Show Commands

```

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
```

route-target

Syntax	route-target
Context	show>router>bgp
Description	This command displays a summary of route-target.

Sample Output

```

*A:Dut-D# show router bgp routes route-target
=====
BGP Router ID:10.20.1.4          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 RT Constrain Routes
=====
Flag  Route Target                      LocalPref  MED
      Nexthop
      As-Path
-----
u*?   0:0:0/0                          None       0
      10.10.9.6
      106 106
-----
```

Routes : 1

=====

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] | [siteid site-id] | [veid veid] [offset vp/s-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]
```

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

family — 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.
- mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.

mvpn-type — 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,

prefix — Specifies the type of routing information to display.

Values	Syntax:	<rd>[<rd>:]<ip-prefix[/ip-prefix-length]>	
	rd	{<ip-addr:comm-val> <2byte-asnumber:ext-comm-val> <4byte-asnumber:comm-val>}	
	comm-val	[0..65535]	
	2byte-asnumber	[0..65535]	
	ext-comm-val	[0..4294967295]	
	4byte-asnumber	asn1.asn2 (two 2-byte pieces)	
		asn1 - [1..65535]	
		asn2 - [0..65535]	
	ip-address	a.b.c.d	
	ipv4-prefix	a.b.c.d	
	ipv4-prefix-le	[0..32]	

filter — 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 *comm-id* — Displays all routes with the specified BGP community.

Values	[<i>as-number1:comm-val1</i> <i>ext-comm</i> <i>well-known-comm</i>]	
	ext-comm	type: {ip-address:comm-val1 as-number1:comm-val2 as-number2:comm-val1}
	as-number1	0 — 65535
	comm-val1	0 — 65535
	type	target, origin
	ip-address	a.b.c.d
	comm-val2	0 — 4294967295
	as-number2	0 — 4294967295
	well-known-comm no-export, no-export-subconfed, no-advertise	

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 — 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 — 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 (SAP(s), pseudowire (s) 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 — 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] [aai-type2 aii-type2] [brief] — Displays routes for ms-pw family.

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 evpn family.

tag vni-id — Displays all routes with the specified ethernet-tag. For VXLAN tunnerls, the ethernet-tag encodes the VNI (VXLAN Network Identifier).

Values 1 — 16777215

— Output

BGP Route — The following table describes the command output fields for BGP routes.

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 PE's 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.

Label	Description (Continued)
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.
Aggregator	The aggregator attribute value. none — Aggregator attributes are not present.
Atomic Aggr.	Atomic — The atomic aggregator flag is set. Not Atomic — The atomic aggregator flag is not 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.
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.

Label	Description (Continued)
	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 215.0.0.0/24 detail
=====
BGP IPv4 Routes
=====
-----
Original Attributes
Network       : 215.0.0.0/24
Nexthop       : 202.50.0.2
Path Id       : None
From          : 202.50.0.2
Res. Nexthop  : 202.50.0.2
Local Pref.   : n/a
Aggregator AS : None
Atomic Aggr.  : Not Atomic
Community     : No Community Members
Cluster       : No Cluster Members
Originator Id : None
Fwd Class     : None
Flags         : Used Valid Best IGP
Route Source  : External
AS-Path       : 5000
Modified Attributes
Network       : 215.0.0.0/24
Nexthop       : 202.50.0.2
Path Id       : None
From          : 202.50.0.2
Interface Name : GE-3/2/1
Aggregator     : None
MED            : None
Peer Router Id : 150.0.0.245
Priority       : None
```

Show Commands

```
Res. Nexthop      : 202.50.0.2
Local Pref.       : 150
Aggregator AS    : None
Atomic Aggr.     : Not Atomic
Community        : 4713:10 4713:510
Cluster          : No Cluster Members
Originator Id    : None
Fwd Class        : None
Flags            : Used Valid Best IGP Sticky
TieBreakReason   : MED
Route Source     : External
AS-Path          : 5000
-----
PMSI Tunnel Attribute :
Tunnel-type      : LDP P2MP LSP
MPLS Label       : 0
Root-Node        : 10.20.1.2
Flags            : Leaf not required
LSP-ID           : 8193
-----

*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
Nexthop         : 10.20.1.1
From            : 10.20.1.1
Res. Nexthop    : 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
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
MPLS Label      : 0
P2MP-ID         : 1001
Extended-Tunne*: 10.20.1.1
Flags           : Leaf not required
Tunnel-ID       : 61440

*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
 Nexthop : 10.20.1.1
 From : 10.20.1.1
 Res. Nexthop : 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 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 l2-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
 Nexthop : 10.20.1.1
 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:1.20.30.40:6543
 l2-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

Show Commands

```
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
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : target:100:100
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
Flags        : Used Valid Best Incomplete
Route Source : Internal
AS-Path      : 106
VPRN Imported : 1
-----
Interface Name : int_to_D
Aggregator     : None
MED            : None
Peer Router Id : 10.20.1.4
Priority        : None
-----
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 SPEC-
IFIED.I.E. ANY RD (RD is optional).
=====
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              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

```

Show Commands

```
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#

```

Show Commands

```
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.   : 10.20.1.2:1      VPN Label    : 131070
From         : 10.20.1.2
Res. Nexthop : 10.10.1.2
Local Pref.   : 100              Interface Name: to-sr7
```



```

Aggregator AS : none
Atomic Aggr. : Not Atomic
Community : target:10.20.1.2:1
Cluster : No Cluster Members
Originator Id : None
Flags : Used Valid Best IGP
AS-Path : No As-Path
VPRN Imported : 1 2 10 12
Aggregator : none
MED : none
Peer Router Id: 10.20.1.2

-----
RIB Out Entries
-----
Routes : 1
=====
A:SR-12#

*A:praragon-siml# /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 SourceAS
Nexthop SourceIP
As-Path 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-siml#

*A:praragon-siml# 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
=====

```

Show Commands

```

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
Nexthop         : 10.20.1.4
From            : 10.20.1.4
Res. Nexthop    : 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 -

```

Show Commands

```
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 Interface Name : net
Aggregator AS : None Aggregator : None
Atomic Aggr. : Not Atomic MED : 5000
AIGP Metric : 100
Community : None
Cluster : No Cluster Members
Originator Id : None Peer Router Id : 2.2.2.2
Fwd Class : None Priority : None
Flags : Used Valid Best Incomplete
Route Source : External
AS-Path : 200 400 500

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

```

```

-----
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
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

```

```
RIB Out Entries
```

```

-----
Network      : 11.1.1.1/32
Nexthop      : 1.1.1.1
Path Id      : None
To           : 3.3.3.3
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : 150
Interface Name : NotAvailable
Aggregator    : None
MED           : 5000

```

Show Commands

```
Community      : None
Cluster        : No Cluster Members
Originator Id  : None                      Peer Router Id : 3.3.3.3
Origin         : Incomplete
AS-Path        : 200 400 500
-----
Routes : 2
=====

*A:DUT#
=====
*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

Interface Name : NotAvailable
Aggregator     : None
MED            : 0
```

```

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              : 0:0:0:0:0:0:0:0:0:0      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                      Interface Name : NotAvailable
Aggregator AS    : None                      Aggregator     : None
Atomic Aggr.     : Not Atomic                MED           : 0
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                      Peer Router Id : 10.20.1.2
Flags            : Used Valid Best IGP
Route Source     : Internal
AS-Path          : No As-Path
EVPN type        : IP-PREFIX
ESI              : N/A                      Tag           : 1

```

Show Commands

```
Gateway Address: 00:00:01:00:01:02
Prefix          : 3.0.1.6/32          Route Dist.    : 10.20.1.2:1
MPLS Label     : X
Route Tag       : Z
Neighbor-AS     : N/A
Orig Validation : N/A
Source Class    : 0                  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                  Interface Name : NotAvailable
Aggregator AS  : None                 Aggregator     : None
Atomic Aggr.   : Not Atomic           MED            : 0
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                 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                  Tag            : 1
Gateway Address: 00:00:01:00:01:02
Prefix          : 3.0.1.6/32          Route Dist.    : 10.20.1.2:1
MPLS Label     : X
Route Tag       : W
Neighbor-AS     : 111
Orig Validation : N/A
Source Class    : 0                  Dest Class     : 0

-----
Routes : 1
=====
```

policy-test

Syntax **policy-test** *policy-name* **family** *family* **prefix** *prefix/pfxlen* [**longer**] [**neighbor** *neighbor*] [**display-rejects**] [**detail**]

Context show>router *router-id*>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 **policy-name** — Must be the name of an existing configured and committed policy.

family — ipv4 or ipv6

Default **ipv4**

prefix — The IPv4 or IPv6 prefix/mask to be evaluated. The keyword longer may be specified to evaluate longer prefix matches. (optional)

vr-id — The router ID.

Default base router if not specified

neighbor — The BGP neighbor (optional)

display-rejects — Display routes that were rejected by the policy. If not specified, only a count of rejected routes will be shown. (optional)

detail — When the policy modifies route attributes, it displays the modifications made to the routes. This command requires an exact prefix to be specified. (optional)

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                            None       None
      220.0.0.2                            None       -
      14
-----
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
=====
```

Show Commands

```
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 NextHop As-Path	LocalPref Path-Id	MED VPNLabel

	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	-

Show Commands

```
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
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    : target:65530:20
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
Flags        : Used Valid Best Incomplete
Route Source : External
AS-Path      : 14
Peer Router Id : 14.14.14.10
Priority      : None
Interface Name : to-sim-6
Aggregator    : None
MED           : 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
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

-----
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

```

Show Commands

```
=====
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   VPNLabel
      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

```

Show Commands

```
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    VPLabel
      As-Path
-----
Accepted by Policy
u*>?  4.0.0.6/32                None      None
```



```

220.0.0.2
14
None -

```

```

-----
Total Routes : 6 Routes rejected : 5

```

```

A:sim-1# show router bgp policy-test bgpprefix44rej 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
      220.0.0.2                  None       None
      14                        None       131069
Accepted by Policy
u*>i  65530:20:8.0.0.1/32
      220.0.0.2                  None       None
      14                        None       131070
Accepted by Policy
u*>i  65530:20:10.0.3.0/24
      220.0.0.2                  None       None
      14 101                    None       131070
Accepted by Policy
u*>i  65530:20:10.13.0.0/24
      220.0.0.2                  None       None
      14 101                    None       131070
Accepted by Policy
u*>i  65530:20:10.23.0.0/24
      220.0.0.2                  None       None
      14 101                    None       131070
Accepted by Policy
u*>i  65530:20:13.13.13.13/32
      220.0.0.2                  None       None
      14 101                    None       131070
Accepted by Policy
u*>i  65530:20:20.20.20.5/32
      220.0.0.2                  None       None
      14                        None       131070
Accepted by Policy
u*>i  65530:20:20.20.20.6/32
      220.0.0.2                  None       None
      14                        None       131070
Rejected by Policy Entry = 10
u*>i  65530:20:44.44.44.0/24
      220.0.0.2                  None       None
      14 101                    None       131070

```

Show Commands

```
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
=====
```

summary

Syntax **summary** [**all**]
summary [**family** *family*] [**neighbor** *ip-address*]

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.
 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

Output **BGP Summary Output** — The following table describes the command output fields for a BGP summary.

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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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 Rcv/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: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

=====
BGP Summary
=====
Neighbor
      AS      PktRcvd InQ  Up/Down  State|Rcv/Act/Sent (IPv4)
              PktSent OutQ                Rcv/Act/Sent (VpnIPv4)
              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**]
fib *slot-number* [**ipv4|ipv6**] **summary**
fib *slot-number* **nh-table-usage**

Context show>router>fib

Description	This command displays FIB information for a specific IOM.
--------------------	---

Parameters *slot-number* — Specifies the slot number.

Values 1 — 10

family — Specify 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.
---------------	---

ip-prefix — The IP prefix for prefix list entry in dotted decimal notation.

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 — FFFF]H
		d: [0 — 255]D

prefix-length — Specifies prefix length.

```

ipv4-prefix-length:      0 — 32
ipv6-prefix-length:      0 — 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 — Shows next-hop table usage.

qos — Specifies the QoS.

accounting-class — Specifies the accounting class.

Output	Sample Output
<p>1. Input: A list of integers: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]</p> <p>2. Process: The list is iterated over, and each element is squared.</p> <p>3. Output: A list of squared integers: [1, 4, 9, 16, 25, 36, 49, 64, 81, 100]</p>	<p>1. Input: A list of integers: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]</p> <p>2. Process: The list is iterated over, and each element is squared.</p> <p>3. Output: A list of squared integers: [1, 4, 9, 16, 25, 36, 49, 64, 81, 100]</p>

```
*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
=====

```

mvpn

Syntax mvpn

Context show>router

Description This command displays Multicast VPN related 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#

```

Show Commands

route-table

Syntax **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*]

Context show>router>route-table

Description This command displays route-table information.

Output **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                      0
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

```

Show Commands

```
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
=====
```

Clear Commands

damping

Syntax	damping <i>[[ip-prefix/ip-prefix-length] [neighbor ip-address]] [group name]</i>								
Context	clear>router>bgp								
Description	This command clears or resets the route damping information for received routes.								
Parameters	<p><i>ip-prefix/ip-prefix-length</i> — Clears damping information for entries that match the IP prefix and prefix length.</p> <p>Values</p> <table> <tr> <td>ipv4-prefix:</td><td>a.b.c.d (host bits must be 0)</td></tr> <tr> <td>ipv4-prefix-length:</td><td>0 — 328</td></tr> </table> <p><i>neighbor ip-address</i> — Clears damping information for entries received from the BGP neighbor.</p> <p>Values</p> <table> <tr> <td>ipv4-address:</td><td>a.b.c.d</td></tr> </table> <p><i>group name</i> — Clears damping information for entries received from any BGP neighbors in the peer group.</p> <p>Values</p> <table> <tr> <td></td><td>32 characters maximum</td></tr> </table>	ipv4-prefix:	a.b.c.d (host bits must be 0)	ipv4-prefix-length:	0 — 328	ipv4-address:	a.b.c.d		32 characters maximum
ipv4-prefix:	a.b.c.d (host bits must be 0)								
ipv4-prefix-length:	0 — 328								
ipv4-address:	a.b.c.d								
	32 characters maximum								

flap-statistics

Syntax	flap-statistics <i>[[ip-prefix/mask] [neighbor ip-address]] [group group-name] [regex reg-exp] [policy policy-name]</i>						
Context	clear>router>bgp						
Description	This command clears route flap statistics.						
Parameters	<p><i>ip-prefix/mask</i> — Clears route flap statistics for entries that match the specified IP prefix and mask length.</p> <p>Values</p> <table> <tr> <td>ip-prefix:</td><td>a.b.c.d (host bits must be 0)</td></tr> <tr> <td>mask:</td><td>0 — 32</td></tr> </table> <p><i>neighbor ip-address</i> — Clears route flap statistics for entries received from the specified BGP neighbor.</p> <p>Values</p> <table> <tr> <td>ipv4-address:</td><td>a.b.c.d</td></tr> </table> <p><i>group group-name</i> — Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.</p> <p><i>regex reg-exp</i> — Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression.</p> <p><i>policy policy-name</i> — Clears route flap statistics for entries that match the specified route policy.</p>	ip-prefix:	a.b.c.d (host bits must be 0)	mask:	0 — 32	ipv4-address:	a.b.c.d
ip-prefix:	a.b.c.d (host bits must be 0)						
mask:	0 — 32						
ipv4-address:	a.b.c.d						

neighbor

Syntax	neighbor { <i>ip-address</i> as <i>as-number</i> external all } [soft soft-inbound] 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	<p><i>ip-address</i> — Resets the BGP neighbor with the specified IP address.</p> <p>Values ipv4-address: a.b.c.d</p> <p>as <i>as-number</i> — Resets all BGP neighbors with the specified peer AS.</p> <p>Values 1 — 65535</p> <p>external — Resets all EBGp neighbors.</p> <p>all — Resets all BGP neighbors.</p> <p>soft — The specified BGP neighbor(s) re-evaluates all routes in the Local-RIB against the configured export policies.</p> <p>soft-inbound — The specified BGP neighbor(s) re-evaluates all routes in the RIB-In against the configured import policies.</p> <p>statistics — The BGP neighbor statistics.</p> <p>end-of-rib — Clears the routing information base (RIB). This command applies when the SR OS node 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.</p>

protocol

Syntax	protocol
Context	clear>router>bgp
Description	Resets the entire BGP protocol.

Debug Commands

events

Syntax	events [<i>neighbor ip-address</i> group name] no events
Context	debug>router>bgp
Description	This command logs all events changing the state of a BGP peer.
Parameters	neighbor ip-address — Debugs only events affecting the specified BGP neighbor. <div> Values ipv4-address: a.b.c.d (host bits must be 0) </div> group name — Debugs only events affecting the specified peer group and associated neighbors.

graceful-restart

Syntax	graceful-restart [<i>neighbor ip-address</i> group name] 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 ip-address — Debugs only events affecting the specified BGP neighbor. <div> Values ipv4-address: a.b.c.d (host bits must be 0) </div> group name — Debugs only events affecting the specified peer group and associated neighbors.

keepalive

Syntax	keepalive [<i>neighbor ip-addr</i> group name] no keepalive
Context	debug>router>bgp
Description	This command decodes and logs all sent and received keepalive messages in the debug log.
Parameters	neighbor ip-address — Debugs only events affecting the specified BGP neighbor. <div> Values ipv4-address: a.b.c.d (host bits must be 0) </div> group name — Debugs only events affecting the specified peer group and associated neighbors.

Debug Commands

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) 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) 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 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** *ip-address* — Debugs only events affecting the specified BGP neighbor.

Values **ipv4-address:** a.b.c.d (host bits must be 0)

group *name* — Debugs only events affecting the specified peer group and associated neighbors.

route-refresh

Syntax **route-refresh** [**neighbor** *ip-address* | **group** *name*]
no route-refresh

Context debug>router>bgp

Description This command enables and disables debugging for BGP route-refresh.

Parameters **neighbor** *ip-address* — Debugs only events affecting the specified BGP neighbor.

Values **ipv4-address:** a.b.c.d (host bits must be 0)

group *name* — Debugs only events affecting the specified peer group and associated neighbors.

rtm

Syntax **rtm** [**neighbor** *ip-address* | **group** *name*]
no rtm

Context debug>router>bgp

Description This command logs RTM changes 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)

group *name* — Debugs only events affecting the specified peer group and associated neighbors.

socket

Syntax **socket** [**neighbor** *ip-address* | **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)

group *name* — Debugs only events affecting the specified peer group and associated neighbors.

Debug Commands

timers

Syntax	timers [neighbor <i>ip-address</i> group <i>name</i>] no timers
Context	debug>router>bgp
Description	This command logs all BGP timer events to 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) group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

update

Syntax	update [neighbor <i>ip-address</i> group <i>name</i>] no update
Context	debug>router>bgp
Description	This command decodes and logs all sent and received update 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) group <i>name</i> — Debugs only events affecting the specified peer group and associated neighbors.

Route Policies

In This Chapter

This chapter provides information about configuring route policies.

Topics in this chapter include:

- [Configuring Route Policies on page 878](#)
 - [Policy Statements on page 879](#)
 - [Default Action Behavior on page 881](#)
 - [BGP and OSPF Route Policy Support on page 890](#)
 - [BGP Route Policies on page 891](#)
 - [Re-advertised Route Policies on page 892](#)
 - [Triggered Policies on page 892](#)
 - [When to Use Route Policies on page 899](#)
- [Route Policy Configuration Process Overview on page 900](#)
- [Configuration Notes on page 901](#)

Configuring Route Policies

Alcatel-Lucent'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.

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.

Routing Policy Subroutines

It is possible to reference a routing policy from within a routing policy to construct powerful subroutine based policies.

A single level of policy subroutines is supported. Policy subroutines may evaluate true or false through matching and policy entry actions. A policy entry action of ‘accept’ will evaluate as true while a policy entry action of ‘reject’ will evaluate as 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 section [Default Action Behavior](#) for more 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.

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.

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.

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.
-

Denied IP Prefixes

The following IP address prefixes are not allowed by the routing protocols and the Route Table Manager and are not be populated within the forwarding table:

- 0.0.0.0/8 or longer
- 127.0.0.0/8 or longer
- 224.0.0.0/4 or longer
- 240.0.0.0/4 or longer

Any other prefixes that need to be filtered can be filtered explicitly using route policies.

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.

Alcatel-Lucent'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.

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 is:

1. Regular expressions should always be enclosed in quotes.
2. An elementary term; for example, an AS number "200"
3. A range term composed of two elementary terms separated by the '-' character like "200-300".
4. The '.' dot wild-card character which matches any elementary term.
5. A regular expression enclosed in parenthesis "()".
6. 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:

1. An elementary term which for a community string is any single digit like “4”.
2. A range term composed of two elementary terms separated by the ‘-’ character like “2-3”.
3. A colon ‘:’ to delimit the AS number from the community value
4. The ‘.’ dot wild-card character which matches any elementary term or ‘.’.
5. A regular expression enclosed in parenthesis “()”.
6. 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.
7. Extended communities such as “target:” and “origin:” may consist of two regular expressions separated by the ampersand (‘&’) character. The first expression is applied to the as-value of the community string and the second to the local administrative value.

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 20](#).

Table 20: 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.

Table 20: Regular Expression Operators (Continued)

Operator	Description
{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 21](#).

Table 21: Community Strings Examples

Example Expression	Example Matches
"ext:...&f(.*?) [af] \$"	Matches the community "ext:0002:fffa0000001a".
'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 22](#).

Table 22: AS Path and Community Regular Expression Examples

AS Path to Match Criteria	Regular Expression	Example Matches
Null AS path	null ^a	Null AS path
AS path is 11	11	11
AS path is 11 22 33	11 22 33	11 22 33

Table 22: AS Path and Community Regular Expression Examples (Continued)

AS Path to Match Criteria	Regular Expression	Example Matches
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
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

Table 22: AS Path and Community Regular Expression Examples (Continued)

AS Path to Match Criteria	Regular Expression	Example Matches
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 ...
AS number is any number. Community value is any number that ends with 1, 2, or 3.	^(.*):(.*[1-3])\$	11:1 100:2002 333:55553 ...

Table 22: AS Path and Community Regular Expression Examples (Continued)

AS Path to Match Criteria	Regular Expression	Example Matches
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

a. The `null` keyword matches an empty AS path.

BGP and OSPF Route Policy Support

OSPF and BGP requires route policy support. [Figure 35](#) and [Figure 37](#) display where route policies are evaluated in the protocol. [Figure 35](#) depicts BGP which applies a route policy as an internal part of the BGP route selection process. [Figure 37](#) 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).

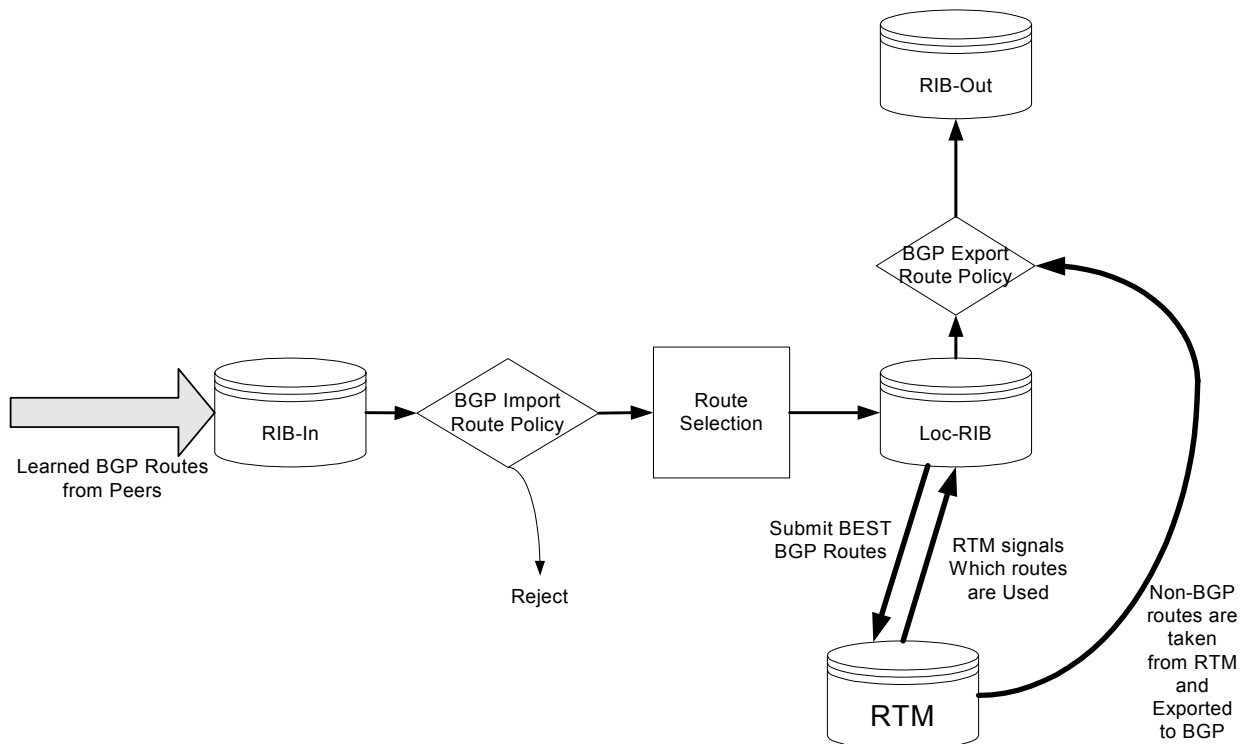


Figure 35: BGP Route Policy Diagram

BGP Route Policies

Alcatel-Lucent'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 36: BGP Route Policy Diagram

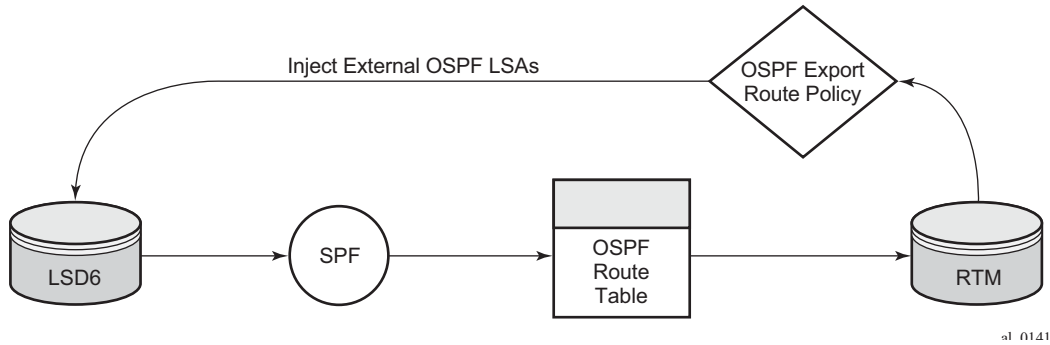


Figure 37: OSPF Route Policy Diagram

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.

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 sent 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.

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 | number 1} | {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 23](#).

Table 23: 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.

BGP Policy Subroutines

Currently, BGP policies only support a single level/tier of configuration that makes configuring complicated policies difficult to meet Internet transit/peering policy requirements.

This feature allows an operator to configure a policy entry that may use a sub-policy as match criteria to improve the flexibility of match criteria. Only a single level of policy nesting is supported (the sub-policy must not have any sub-policies of its own).

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.

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 38](#), 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.

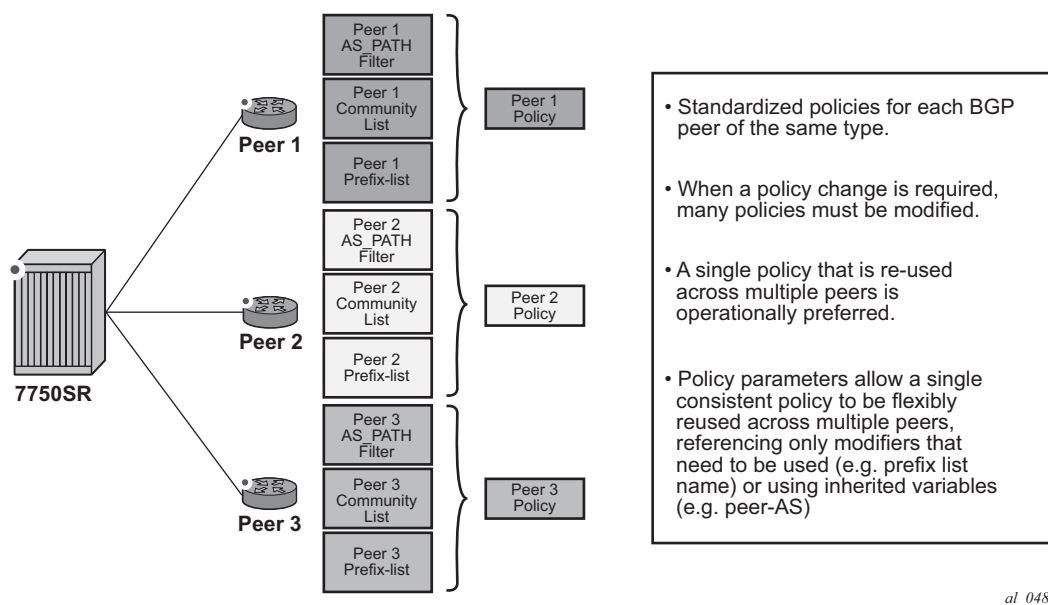


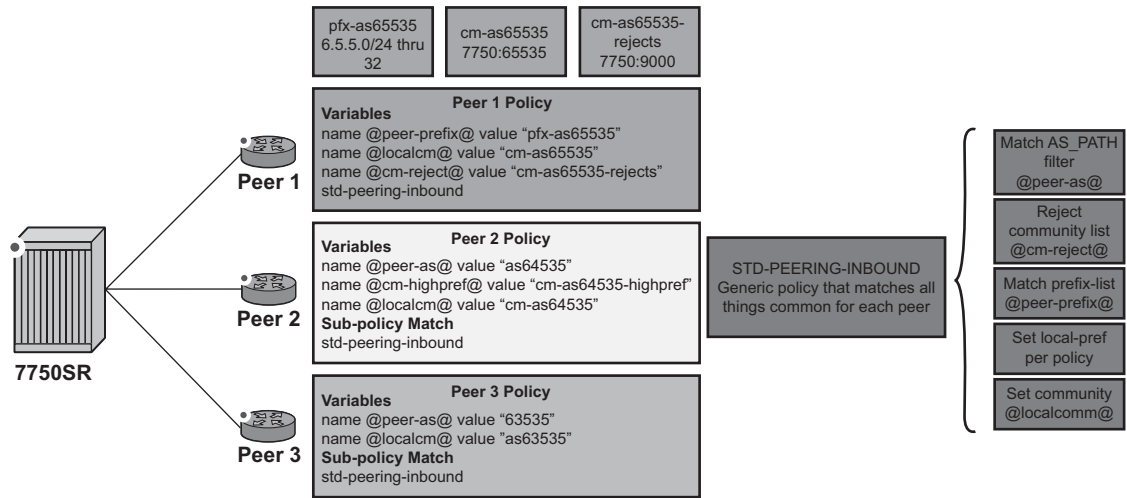
Figure 38: Route Policy Past Mode of Operation

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

Routing Policy Parameterization

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.



al_0484

Figure 39: Route Policy Parameterization using sub-policies

The approach taken to implement route policy parameterization is to re-use the sub-policy functionality introduced in SR OS 11.0.

The logical flow of this is to configure a per-peer policy in which the variable names and values are defined. Using [Figure 39](#) 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 10
    from
      policy-variables
```

```

        name "@localcm@" value "cm-as65535"
        name "@peer-as@" value "as65535"
        name "@cm-reject@" value "cm-as65535-rejects"
        name "@cm-highpref@" value "cm-as65535-highpref"
        name "@peer-prefix@" value "pfx-as65535"
    exit
    policy "std-peering-inbound"
exit
action accept
exit
exit
exit
policy-statement "peer2"
description "peer2 inbound at IXP ABC, using std-peering-inbound"
entry 10
    from
        policy-variables
            name "@localcm@" value "cm-as64535"
            name "@peer-as@" value "as64535"
            name "@cm-reject@" value "cm-as64535-rejects"
            name "@cm-highpref@" value "cm-as64535-highpref"
            name "@peer-prefix@" value "pfx-as64535"
        exit
        policy "std-peering-inbound"
    exit
    action accept
    exit
exit
exit
policy-statement "peer3"
description "peer3 inbound at IXP ABC, using std-peering-inbound"
entry 10
    from
        policy-variables
            name "@localcm@" value "cm-as63535"
            name "@peer-as@" value "as63535"
            name "@cm-reject@" value "cm-as63535-rejects"
            name "@cm-highpref@" value "cm-as63535-highpref"
            name "@peer-prefix@" value "pfx-as63535"
        exit
        policy "std-peering-inbound"
    exit
    action accept
    exit
exit
exit
policy-statement "std-peering-inbound"
description "Standard inbound peering policy for all standard IXP peers"
entry 10
    from
        community "@cm-reject@"
    exit
    action reject
exit
entry 20
    from
        prefix-list "@peer-prefix@"
        as-path "@peer-as@"
    exit
    action accept
        community add "@localcm@"
        local-preference 400
    exit

```

Routing Policy Parameterization

```
exit
entry 30
  from
    community "@cm-highpref@"
  exit
  action accept
    community add "@localcm@"
    local-preference 4000
  exit
exit
```

This configuration would take slightly different actions depending on the peer.

Peer 1

- Prefixes that have a community matching ‘cm-as65535-rejects’ are specifically rejected.
- Prefix list ‘pfx-as65535’ is evaluated and prefixes accepted based on that prefix-list.
- Local-preference on accepted prefixes is set to 400.
- Community ‘7750:65535’ is added to accepted prefixes.
- As community-list ‘cm-65535-highpref’ doesn’t exist, this entry is not evaluated.

Peer 2

- As community-list ‘cm-64535-rejects’ doesn’t exist, this entry is not evaluated.
- Prefix list ‘pfx-as64535’ and AS-path ‘as64535’ is evaluated and prefixes accepted based on that prefix-list and AS-path combo.
- Local-preference on accepted prefixes is set to 400.
- Community ‘7750:64535’ is added to accepted prefixes.
- Prefixes matching ‘cm-as64535-highpref’ are set to a local-preference of 4000.

Peer 3

- As community-list ‘cm-as63535-rejects’ doesn’t exist, this entry is not evaluated.
- Prefix-list ‘pfx-as63535’ and AS-path ‘as63535’ is evaluated and prefixes accepted based on that prefix-list and AS-path combo.
- Local-preference on accepted prefixes is set to 400.
- Community ‘7750:63533’ is added on accepted prefixes.
- As community-list ‘cm-63535-highpref’ doesn’t exist, this entry is not evaluated.

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).

Route Policy Configuration Process Overview

Figure 40 displays the process to provision basic route policy parameters.

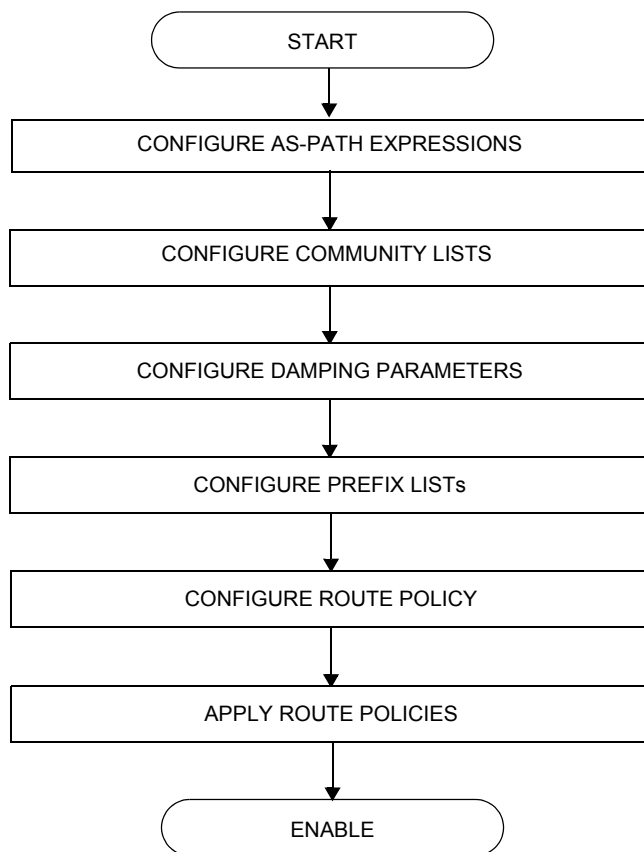


Figure 40: Route Policy Configuration and Implementation Flow

Configuration Notes

This section describes route policy configuration caveats.

General

- When configuring policy statements, the policy statement name must be unique.

Configuring Route Policies with CLI

This section provides information to configure route policies using the command line interface.

Topics in this section include:

- [Route Policy Configuration Overview on page 904](#)
 - [When to Create Routing Policies on page 904](#)
 - [Policy Evaluation on page 906](#)
 - [Damping on page 909](#)
- [Configuring Route Policy Components on page 912](#)
 - [Creating a Route Policy on page 914](#)
 - [Beginning the Policy Statement on page 913](#)
 - [Configuring an Entry on page 916](#)
 - [Configuring a Community List on page 917](#)
 - [Configuring Damping on page 918](#)
 - [Configuring a Prefix List on page 919](#)
 - [Configuring PIM Join/Register Policies on page 920](#)
- [Route Policy Configuration Management Tasks on page 923](#)

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.

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.

Default Route Policy Actions

Each routing protocol has default behaviors for the import and export of routing information. [Table 24](#) shows the default behavior for each routing protocol.

Table 24: 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.	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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.	<ul style="list-style-type: none"> • 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.	<ul style="list-style-type: none"> • 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.

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 41 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 42 depicts the next-policy and next-entry route processes.

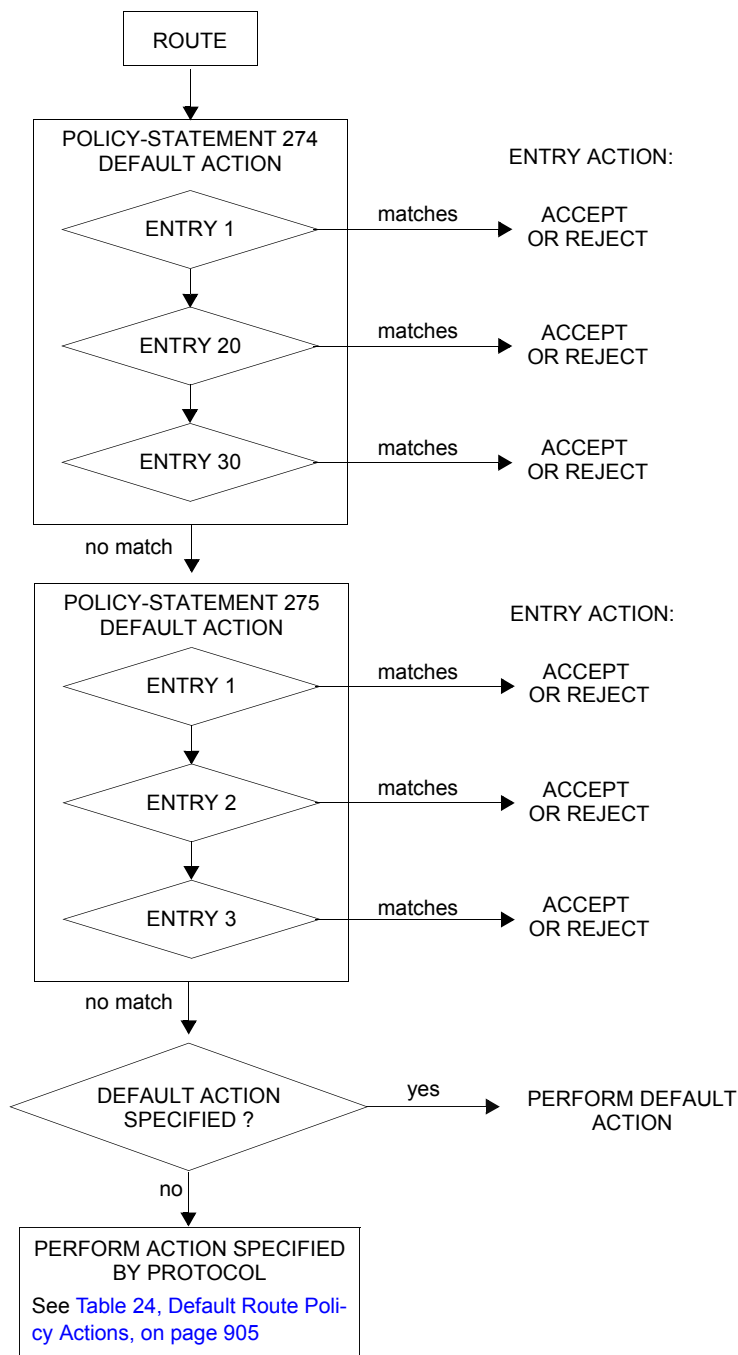


Figure 41: Route Policy Process Example

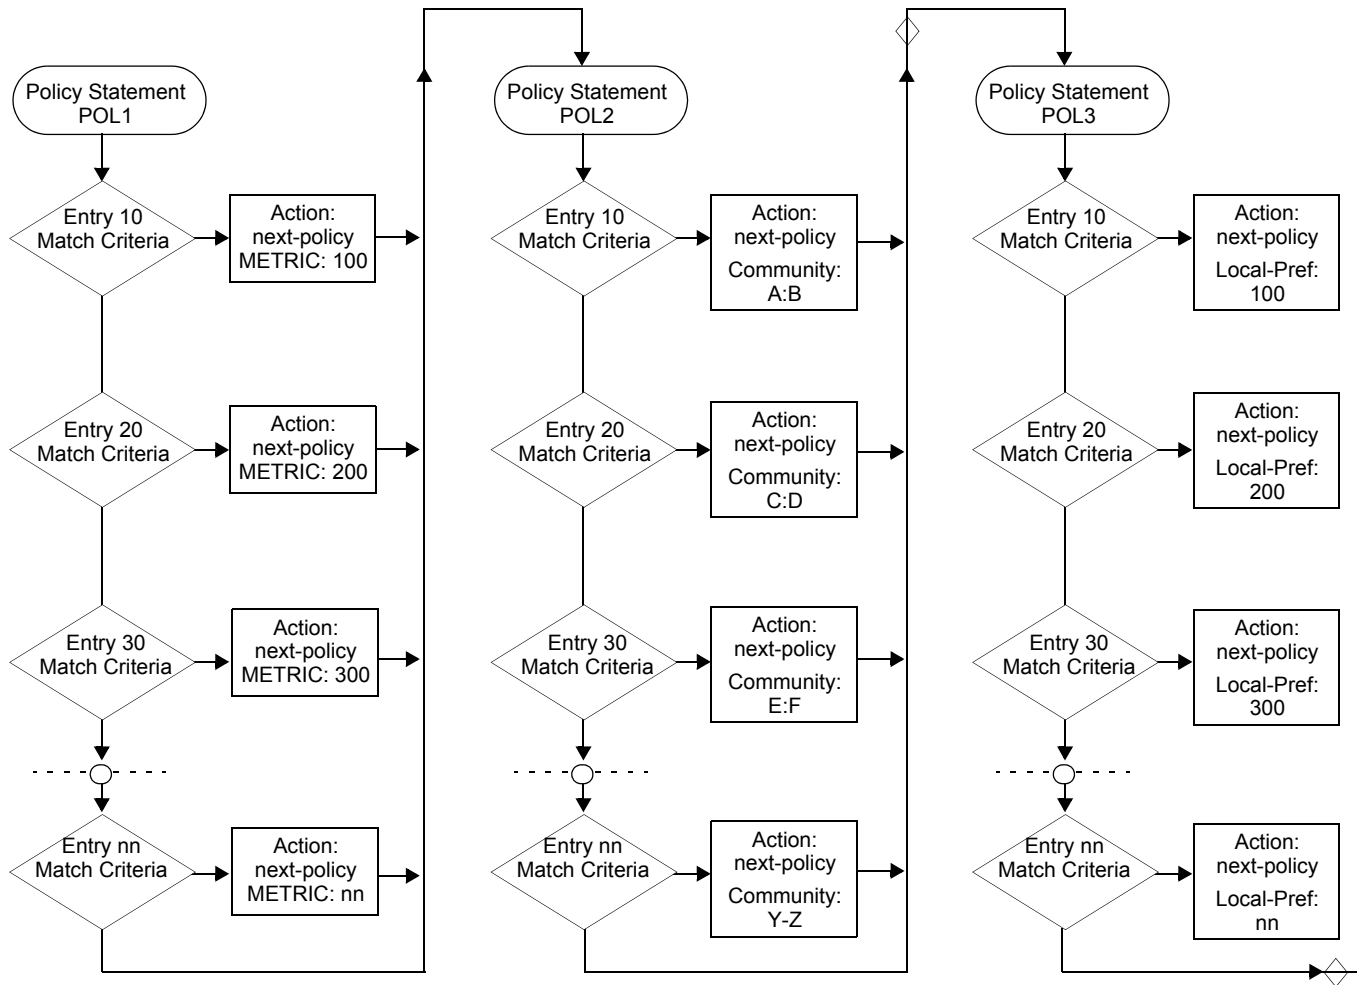


Figure 42: Next Policy Logic Example

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 43](#) 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 43: Damping Example

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#
```

Configuring Route Policy Components

Use the CLI syntax displayed below to configure:

- [Creating a Route Policy on page 914](#)
- [Beginning the Policy Statement on page 913](#)
- [Configuring an Entry on page 916](#)
- [Configuring a Community List on page 917](#)
- [Configuring Damping on page 918](#)
- [Configuring a Prefix List on page 919](#)
- [Configuring PIM Join/Register Policies on page 920](#)

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 on page 916](#)).

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.

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#
```

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 on page 916](#)).

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#
```

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 on page 927](#) 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
          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
-----
A:ALA-B>config>router>policy-options#
    policy-statement "exporttmsgt"
      entry 1
        from
          protocol vpn-leak
        exit
        action accept
        exit
      exit
      entry 2
        from
          protocol tms
        exit
        action accept
        exit
      exit
    exit
  commit
exit
exit
```


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#
```

Configuring Damping

NOTES:

- For each damping profile, all parameters must be configured.
- The `suppress` value must be greater than the `reuse` value (see [Figure 43 on page 909](#)).
- 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#
```

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 "dampetest123"
          half-life 15
          max-suppress 60
          reuse 750
      exit
-----
A:ALA-B>config>router>policy-options#
```

Configuring PIM Join/Register Policies

Join policies are used in Protocol Independent Multicast (PIM) configurations to prevent the transportation of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. PIM Join filters reduce the potential for denial of service (DoS) attacks and PIM state explosion—large numbers of Joins forwarded to each router on the RPT, resulting in memory consumption. See [Importing PIM Join/Register Policies on page 74](#).

*,G or S,G is the information used to forward unicast or multicast packets.

- **group-address** matches the group in join/prune messages
group-address 229.55.150.208/32 exact
- **source-address** matches the source in join/prune messages
source-address 192.168.0.0/16 longer
- **interface** matches any join message received on the specified interface
interface port 1/1/1
- **neighbor** matches any join message received from the specified neighbor
neighbor 1.1.1.1

The following configuration example will not allow join messages for group 229.50.50.208/32 and source 192.168.0.0/16 but allows other join messages.

Configuring policy-statement

```
A:ALA-B>config>router# policy-options
A:ALA-B>config>router>policy-options# begin
A:ALA-B>config>router>policy-options# policy-statement foo
A:ALA-B>config>router>policy-options>policy-statement$ entry 10
A:ALA-B>config>router>policy-options>policy-statement>entry$ from
A:ALA-B>config>router>policy-options>policy-statement>entry>from$ group-address
229.50.50.208/32
A:ALA-B>config>router>policy-options>policy-statement>entry>from$ source-address
192.168.0.0
A:ALA-B>config>router>policy-options>policy-statement>entry>from$ exit
A:ALA-B>config>router>policy-options>policy-statement>entry# action reject
A:ALA-B>config>router>policy-options>policy-statement>entry#
```

The following configuration example allows registers for *, 224.0.0.0/8.

```
A:ALA-B>config>router>policy-options# policy-statement reg-pol
A:ALA-B>config>router>policy-options>policy-statement$ entry 10
A:ALA-B>config>router>policy-options>policy-statement>entry$ from
A:ALA-B>config>router>policy-options>policy-statement>entry>from$ group-address 224.0.0.0/
8
A:ALA-B>config>router>policy-options>policy-statement>entry# action accept
A:ALA-B>config>router>policy-options>policy-statement>entry>action# exit
A:ALA-B>config>router>policy-options>policy-statement>entry# exit
A:ALA-B>config>router>policy-options>policy-statement# exit
```

```
A:ALA-B>config>router>policy-options# info
-----
...
    policy-statement "foo"
        entry 10
            from
                group-address "229.50.50.208/32"
                source-address 192.168.0.0
            exit
            action reject
        exit
    exit
    policy-statement "reg-pol"
        entry 10
            from
                group-address "224.0.0.0/8"
            exit
            action accept
        exit
    exit
exit
...
-----
A:ALA-B>config>router>policy-options#
```

Configuring Bootstrap Message Import and Export Policies

Bootstrap import and export policies are used to control the flow of bootstrap messages to and from the RP.

The following configuration example specifies that no BSR messages received or sent out of interface port 1/1/1.

```
A:ALA-B>config>router>policy-options# policy-statement pim-import
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ from
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ interface port 1/1/1
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ exit
:A:ALA-B>config>router>policy-options>policy-statement>entry# action reject
:A:ALA-B>config>router>policy-options>policy-statement>entry# exit
:A:ALA-B>config>router>policy-options>policy-statement# exit

:A:ALA-B>config>router>policy-options# policy-statement pim-export
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ to
:A:ALA-B>config>router>policy-options>policy-statement>entry>to$ interface port 1/1/1
:A:ALA-B>config>router>policy-options>policy-statement>entry# action reject
:A:ALA-B>config>router>policy-options>policy-statement>entry# exit
:A:ALA-B>config>router>policy-options>policy-statement# exit

:A:ALA-B>configure router pim rp bootstrap-import pim-import
:A:ALA-B>configure router pim rp bootstrap-export pim-export
```

Route Policy Configuration Management Tasks

This section discusses the following route policy configuration management tasks:

- [Editing Policy Statements and Parameters on page 923](#)
 - [Deleting an Entry on page 925](#)
 - [Deleting a Policy Statement on page 925](#)
-

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
```

Route Policies

```
exit
default-action accept
  as-path add "test"
  community add "365"
  damping "flapper"
  next-hop 10.10.10.104
exit
-----
```


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

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

Route Policy Command Reference

Command Hierarchies

- [Route Policy Configuration Commands on page 927](#)
- [Show Commands on page 931](#)

Route Policy Configuration Commands

```

config
  — [no] router
    — [no] triggered-policy
    — [no] policy-options
      — abort
      — as-path (policy options) name {regular-expression | null}
      — no as-path (policy options) name
      — as-path-group (policy options) name
      — no as-path-group (policy options)
        — [no] entry entry-id expression reg-exp
      — begin [exclusive]
      — commit
      — community name members comm-id [comm-id] (up to 15 max)
      — community name expression expression
      — no community name [members 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] policy-statement name
        — default-action {accept | next-entry | next-policy | reject}
        — no default-action
          — add-paths-send-limit send-limit
          — no add-paths-send-limit
          — aigp-metric {number1 | igp | add number2}
          — as-path {add | replace} name
          — no as-path
          — as-path-prepend as-number [repeat]
          — no as-path-prepend
          — bgp-leak
          — no bgp-leak
          — community {{add name [remove name]} | {remove name [add name]} | {replace name}}
```

```

— no community
— 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}
— no origin
— origin-validation-state state
— no origin-validation-state
— preference preference
— no preference
— [no] sticky-ecmp
— tag hex-string
— no tag
— type {type}
— no type
— description description-string
— no description
— [no] entry entry-id
— action {accept| next-entry | next-policy | reject}
— no action
— add-paths-send-limit send-limit
— no add-paths-send-limit
— aigp-metric {number1 | igp | add number2}
— as-path {add | replace} name
— no as-path
— as-path-prepend as-number [ repeat ]
— no as-path-prepend
— bgp-leak
— no bgp-leak
— community { {add name [remove name]} | {remove name
[add name]} | {replace name}}
— no community
— 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 {set {igp | number1} | { add | subtract} number2}
— no metric
— next-hop ip-address

```

```

— no next-hop
— [no] next-hop-self
— [no] next-hop-self
— origin {igp | egp | incomplete}
— no origin
— origin-validation-state state
— no origin-validation-state
— policy-variables name
— preference preference
— no preference
— [no] sticky-ecmp
— tag tag
— no tag
— type {type}
— no type
— description description-string
— no description
— [no] from
    — area area-id
    — no area
    — as-path name
    — no as-path
    — community name
    — no community
    — [no] external
    — family [ipv4][mcast-ipv4] [vpn-ipv4] [l2-vpn] [mvpn-
      ipv4] [mdt-safi] [flow-ipv4] [route-target] [mcast-vpn-
      ipv4]
    — no family
    — 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
    — mvpn-type {1 | 2 | 3 | 4 | 5 | 6 | 7}
    — no mvpn-type
    — neighbor {ip-address | prefix-list name}
    — no neighbor
    — origin {igp | egp | incomplete | any}
    — no origin
    — origin-validation-state state
    — no origin-validation-state
    — policy-variables name name-string value value-string
    — prefix-list name [name...(up to 5 max)]
    — no prefix-list
    — protocol protocol [all | instance instance]
    — no protocol
    — source-address ip-address
    — source-address prefix-list prefix-list-name
    — no source-address
    — state state

```

- **no state**
 - **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*...(up to 5 max)]
 - **protocol** *protocol*
 - **no protocol**

```

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]

```

Show Commands

```

show
— router router-name
— policy [name | damping | prefix-list name | as-path name | community name | admin]

show
— router router-name
— bgp
— route
— policy-test policy-name family family prefix prefix/pfxlen [longer]
— neighbor neighbor [display-rejects] [detail]

```

Route Policy Command Reference

Generic Commands

abort

Syntax	abort
Context	config>router>policy-options This command is required to discard changes made to a route policy.
Default	none

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.
Default	none
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.
Default	none

description

Syntax	description <i>string</i> no description
Context	config>router>policy-options>policy-statement config>router>policy-options>policy-statement>entry
Description	<p>This command creates a text description which is stored in the configuration file to help identify the content of the entity.</p> <p>The no form of the command removes the string from the configuration.</p>
Default	none
Parameters	<i>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, etc.), the entire string must be enclosed within double quotes.

Route Policy Options

as-path (policy options)

Syntax	as-path <i>name</i> { <i>reg-exp</i> null} no as-path <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 regular expression statement is defined.
Parameters	<p><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, etc.), the entire string must be enclosed within double quotes.</p> <p><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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”.</p> <p>null — The AS path expressed as an empty regular expression string.</p>

as-path-group (policy options)

Syntax	as-path-group <i>name</i> no as-path <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 regular expression statement is defined.
Parameters	<p><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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”.</p>

entry

Syntax	entry <i>entry-id</i> expression <i>reg-exp</i> no entry
Context	config>router>policy-options>as-path-group
Description	<p>This command creates the context to edit route policy entries within an autonomous system path group.</p> <p>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.</p> <p>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.</p> <p>The no form of the command removes the specified entry from the autonomous system path group.</p>
Default	none
Parameters	<p><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.</p> <p>Values 1 — 4294967295</p> <p><i>reg-exp</i> — 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, etc.), the entire string must be enclosed within double quotes.</p>

community

Syntax	community <i>name</i> members <i>comm-id</i> [<i>comm-id</i>]...(up to 15 max) community <i>name</i> expression <i>expression</i> no community <i>name</i> [members <i>comm-id</i>]
Context	config>router>policy-options
Description	<p>This command creates a route policy community list to use in route policy entries.</p> <p>The no form of the command deletes the community list or the provided community ID.</p>
Default	no community — No community names or members are specified.
Parameters	<p><i>name</i> — The community 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, etc.), the entire string must be enclosed within double quotes.</p>

comm-id — The community ID. Note that up to 15 community ID strings can be specified up to a total maximum of 72 characters. A community ID can be specified in different forms.

Values 72 chars max

```
<2byte-asnumber:comm-val>|<reg-ex>|<ext-comm>|<well-known-comm>
ext-comm      <type>:  {<ip-address:comm-val>|
                        <reg-ex1&reg-ex2>|
                        <ip-address&reg-ex2>|
                        <2byte-asnumber:ext-comm-val>|
                        <4byte-asnumber:comm-val>|<as-number:val-in-mbps>}

ext:4300:<ovstate>
extL<value1>:<value2>
2byte-asnumber [0..65535]
comm-val       [0..65535]
reg-ex         [72 chars max]
type           target|origin
ip-address     a.b.c.d
ext-comm-val   [0..4294967295]
4byte-asnumber [0..4294967295]
reg-ex1        [63 chars max]
reg-ex2        [63 chars max]
well-known-comm null|no-export|no-export-subconfed| no-advertise
as-number      [0..65535]
val-in-mbps    [0..16777215]
ovstate        0, 1 or 2 (0 for valid), (1 for Not-Found), or (2 for Invalid)
value1         [0000..FFFF]
value2         [0..FFFFFFFFFFFFFF]
```

- *as-num:comm -value* — The *as-num* is the Autonomous System Number (ASN)

Values as-num: 1 — 65535
comm-value: 0 — 65535

- *type {target | origin}:as-num:comm-value* — The keywords *target* or *origin* denote the community as an extended community of type route target or route origin respectively. The *as-num* and *comm-value* allow the same values as described above for regular community values, including regular expressions.
- *reg-ex1 reg-ex2* — A regular expression string. Allowed values are any string up to 63 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.
- *well-known-comm* — keywords **null**, **no-export**, **no-export-subconfed**, **no-advertise**.

expression *expression* — Specifies a logical expression containing terms and operators. It can contain sub-expressions enclosed in round brackets.

Values 900 chars max <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])"

Route Policy Options

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	300 seconds
Parameters	<i>seconds</i> — Specifies the duration the session with exclusive lock may be inactive.
	Values Values 1 - 3600

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.
Default	none

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 effect 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 <i>soft</i> or <i>soft-inbound</i> 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	Non-dynamic route policy is disabled.

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 profiles are defined.
Parameters	<i>name</i> — 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, etc.), 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	<p>This command configures the half-life parameter for the route damping profile.</p> <p>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).</p> <p>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.</p> <p>The no form of the command removes the half life parameter from the damping profile.</p>
Default	No half life value is specified. The half life value must be explicitly configured.
Parameters	<i>minutes</i> — The half life in minutes expressed as a decimal integer. Values 1 — 45

max-suppress

Syntax	max-suppress <i>minutes</i> no max-suppress
Context	config>router>policy-options>damping
Description	<p>This command configures the maximum suppression parameter for the route damping profile.</p> <p>This value indicates the maximum time, expressed in minutes, that a route can remain suppressed.</p> <p>The no form of the command removes the maximum suppression parameter from the damping profile.</p>
Default	No maximum suppression time is configured.
Parameters	<i>minutes</i> — The maximum suppression time, in minutes, expressed as a decimal integer.
Values	1 — 720

reuse

Syntax	reuse <i>integer</i> no reuse
Context	config>router>policy-options>damping
Description	<p>This command configures the reuse parameter for the route damping profile.</p> <p>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.</p> <p>The no form of the command removes the reuse parameter from the damping profile.</p>
Default	No reuse parameter is configured.
Parameters	<i>integer</i> — The reuse value expressed as a decimal integer.
Values	1 — 20000

suppress

Syntax **suppress** *integer*
 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 parameter is configured.**

Parameters *integer* — The suppress value expressed as a decimal integer.

Values 1 — 20000

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.
Default	none
Parameters	<i>name</i> — 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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”.

prefix

Syntax	[no] prefix <i>ip-prefix/prefix-length</i> { [exact longer through <i>length</i>] [prefix-length-range <i>length1-length2</i>] } no prefix [<i>ipv-prefix/prefix-length</i>] [exact longer through <i>length</i> prefix-length-range <i>length1-length2</i>]												
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	<i>ip-prefix</i> — The IP prefix for prefix list entry in dotted decimal notation. <table><tr><td>Values</td><td>ipv4-prefix:</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td></td><td>ipv4-prefix-length:</td><td>0 — 32</td></tr></table> exact — Specifies the prefix list entry only matches the route with the specified <i>ip-prefix</i> and prefix <i>mask</i> (<i>length</i>) values. longer — Specifies the prefix list entry matches any route that matches the specified <i>ip-prefix</i> and prefix <i>mask</i> <i>length</i> values equal to or greater than the specified mask. through <i>length</i> — Specifies the prefix list entry matches any route that matches the specified ip-prefix and has a prefix length between the specified <i>length</i> values inclusive. <table><tr><td>Values</td><td>0 — 32</td></tr></table> prefix-length-range <i>length1 - length2</i> — 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. <table><tr><td>Values</td><td>0 — 32, <i>length2</i> > <i>length1</i></td></tr></table>			Values	ipv4-prefix:	a.b.c.d (host bits must be 0)		ipv4-prefix-length:	0 — 32	Values	0 — 32	Values	0 — 32, <i>length2</i> > <i>length1</i>
Values	ipv4-prefix:	a.b.c.d (host bits must be 0)											
	ipv4-prefix-length:	0 — 32											
Values	0 — 32												
Values	0 — 32, <i>length2</i> > <i>length1</i>												

Route Policy Entry Match Commands

entry

Syntax	entry <i>entry-id</i> no entry
Context	config>router>policy-options>policy-statement
Description	<p>This command creates the context to edit route policy entries within the route policy statement.</p> <p>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.</p> <p>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.</p> <p>The no form of the command removes the specified entry from the route policy statement.</p>
Default	none
Parameters	<p><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.</p> <p>Values 1 — 4294967295</p>

area

Syntax	area <i>area-id</i> no area
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command configures an OSPF area as a route policy match criterion.</p> <p>This match criterion is only used in export policies.</p> <p>All OSPF routes (internal and external) are matched using this criterion if the best path for the route is by the specified area.</p> <p>The no form of the command removes the OSPF area match criterion.</p>
Default	none
Parameters	<p><i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.</p> <p>Values 0.0.0.0 — 255.255.255.255 (dotted decimal), 0 — 4294967295 (decimal)</p>

Route Policy Entry Match Commands

as-path

Syntax	as-path <i>name</i> no as-path
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command configures an AS path regular expression statement as a match criterion for the route policy entry.</p> <p>If no AS path criterion is specified, any AS path is considered to match.</p> <p>AS path regular expression statements are configured at the global route policy level (config>router>policy-options>as-path <i>name</i>).</p> <p>The no form of the command removes the AS path regular expression statement as a match criterion.</p>
Default	no as-path — Matches any AS path.
Parameters	<i>name</i> — Specifies an existing name. 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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”.

community

Syntax	community <i>name</i> no community
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command configures a community list as a match criterion for the route policy entry.</p> <p>If no community list is specified, any community is considered a match.</p> <p>The no form of the command removes the community list match criterion.</p>
Default	no community — Matches any community.
Parameters	<p><i>name</i> — The community 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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”.</p> <p>The <i>name</i> specified must already be defined.</p>

from

Syntax	[no] from
Context	config>router>policy-options>policy-statement>entry
Description	<p>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.</p> <p>If no condition is specified, all route sources are considered to match.</p> <p>The no form of the command deletes the source match criteria for the route policy statement entry.</p>

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] [mcast-ipv4] [vpn-ipv4] [l2-vpn] [mvpn-ipv4] [mdt-safi] [flow-ipv4] [flow-ipv6] [route-target] [mcast-vpn-ipv4] [evpn] no family
Context	config>router>policy-options>policy-statement>entry>from
Description	This command specifies address families as matching conditions.
Parameters	<p>ipv4 — Specifies IPv4 routing information.</p> <p>mcast-ipv4 — Specifies multicast IPv4 routing information.</p> <p>vpn-ipv4 — Specifies IPv4 VPN routing information.</p> <p>l2-vpn — Exchanges Layer 2 VPN information.</p> <p>mvpn-ipv4 — Exchanges Multicast VPN related information</p> <p>mdt-safi — Exchange Multicast VPN (MDT-SAFI) related information</p> <p>flow-ipv4 — Exchanges IPv4 flowspec routes belonging to AFI 1 and SAFI 133</p> <p>flow-ipv6 — Exchange Ipv6 flowspec routes belonging to AFI 2 and SAFI 133</p> <p>route-target — Specifies to use route targets to be advertised to the peers if ORF is enabled for this peer group</p> <p>mcast-vpn-ipv4 — – Exchanges Multicast Routes in VPN using SAFI 129.</p>

Route Policy Entry Match Commands

group-address

Syntax	group-address <i>prefix-list-name</i> no group-address
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command specifies the multicast group-address prefix list containing multicast group-addresses that are imbedded 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.</p> <p>The no form of the command removes the criterion from the configuration.</p>
Default	no group-address
Parameters	<p><i>prefix-list-name</i> — 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, etc.), the entire string must be enclosed within double quotes.</p> <p>The <i>prefix-list-name</i> is defined in the config>router>policy-options>prefix-list context.</p>

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	<p><i>prefix-list-name</i> — 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, etc.), the entire string must be enclosed within double quotes.</p> <p>The <i>prefix-list-name</i> is defined in the config>router>policy-options>prefix-list context.</p>

interface

Syntax	interface <i>interface-name</i> no interface
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command specifies the router interface, specified either by name or address, as a filter criterion.</p> <p>The no form of the command removes the criterion from the configuration.</p>
Default	no interface
Parameters	<p><i>ip-int-name</i> — Specify the name of the interface as a match criterion for this entry. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p>

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>polic-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 RFC6514.

neighbor

Syntax	neighbor { <i>ip-address</i> prefix-list <i>name</i> } no neighbor			
Context	config>router>policy-options>policy-statement>entry>to config>router>policy-options>policy-statement>entry>from			
Description	<p>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.</p> <p>The no form of the of the command removes the neighbor IP match criterion from the configuration.</p>			
Default	no neighbor — Matches any neighbor.			
Parameters	<i>ip-addr</i> — The neighbor IP address in dotted decimal notation.			
	<table><tr><td>Values</td><td>ipv4-address:</td><td>a.b.c.d</td></tr></table>	Values	ipv4-address:	a.b.c.d
Values	ipv4-address:	a.b.c.d		
	prefix-list <i>name</i> — 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, etc.), the entire string must be enclosed within double quotes.			

Route Policy Entry Match Commands

The *name* specified must already be defined.

origin

Syntax	origin {igp egp incomplete any aaa dhcp lubd} no origin
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command will configure a match criteria for the origin attribute. Originally, the origin attribute was applicable only to BGP as a mandatory well-known BGP attribute.</p> <p>The functionality of the origin attribute has expanded to subscriber-management routes (/32 IPv4 host and IPv6 PD wan-host routes). Each subscriber-management route will internally (local to the node) by default carry the origin attribute with one of the three new values (aaa, dhcp and lubd). The value of the attribute will depend on the origin of the subscriber-management route. The aaa, dhcp or lubd values will never be carried in BGP updates as part of the BGP origin attribute or be otherwise visible within the BGP process.</p> <p>This introduction of the three new values for the origin attribute in the subscriber-management routes will allow customized advertisement of the subscriber-management routes via routing policy.</p>
Default	no origin — Matches any BGP origin attribute
Parameters	<p>igp — Configures matching path information originating within the local AS.</p> <p>egp — Configures matching path information originating in another AS.</p> <p>incomplete — Configures matching path information learned by another method.</p> <p>any — Specifies to ignore this criteria.</p> <p>aaa — IPv4</p> <p>subscriber-management /32 host routes that are originated via 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.</p> <p>IPv6</p> <p>subscriber-management routes that are originated through framed-ipv6-prefix (SLAAC), delegated-ipv6-prefix (IA_PD) or alc-ipv6-address (IA_NA) Radius attributes . This is valid for IPoE and PPPoE type host.</p> <p>dhcp — IPv4</p> <p>subscriber-management /32 host routes that are originated via DHCP server (local or remote) and also Radius framed-ip-address=255.255.255.254 (RFC 2865).</p> <p>IPv6</p> <p>subscriber-management routes that are assigned via local DHCPv6 server pools whose name is obtained through Alc-Delegated-IPv6-Pool (PD pool) and Framed-IPv6-Pool (NA pool) Radius attributes. This is valid for IPoE and PPPoE type hosts.</p> <p>In addition, for IPoEv6 only, the pool name can be also obtained via ipv6-delegated-prefix-pool (PD pool) and ipv6-wan-address-pool (NA pool) from LUDB.</p>

ludb — IPv4

subscriber-management /32 host routes that are originated via LUDB. This should also cover Radius fallback category (Radius falls back to system-defaults or to LUDB).

IPv6

subscriber-management routes obtained from LUDB via ipv6-address (IA_NA) or ipv6-prefix (IA_PD). This is supported only for IPoE.

origin-validation-state

Syntax	origin-validation-state <i>state</i> 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: <ul style="list-style-type: none"> • Valid (0) • Not-Found (1) • Invalid (2)
Default	no origin-validation-stat
Parameters	valid — Match routes with an origin validation state of valid. notFound — Match routes with an origin validation state of Not Found. invalid — Match routes with 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 — No route policy statements are defined.

Route Policy Entry Match Commands

Parameters *name* — The route policy statement 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, etc.), the entire string must be enclosed within double quotes.

policy-statement

Syntax **[no] policy-statement** *name*

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** — No route policy statements are defined.

Parameters *name* — The route policy statement 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, etc.), the entire string must be enclosed within double quotes

policy-variables

Syntax **policy-variables** *name-string* **value** *value-string*
no policy-variables

Context config>router>policy-options>policy-statement>from

Description Routing policy variable 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). Using policy variables 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 variable names and the variable value.

The **no** form of the command removes the **policy-variables** statement.

Parameters *name-string* — The name of the policy 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 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.

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, etc.), the entire string must be enclosed within double quotes.

prefix-list

Syntax	prefix-list <i>name</i> [<i>name</i> ...up to 5 max] no prefix-list
Context	config>router>policy-options>policy-statement>entry>from config>router>policy-options>policy-statement>entry>to
Description	<p>This command configures a prefix list as a match criterion for a route policy statement entry.</p> <p>If no prefix list is specified, any network prefix is considered a match.</p> <p>An empty prefix list will evaluate as if 'no match' was found.</p> <p>The prefix lists specify the network prefix (this includes the prefix and length) a specific policy entry applies.</p> <p>A maximum of five prefix names can be specified.</p> <p>The no form of the command removes the prefix list match criterion.</p>
Default	no prefix-list — Matches any network prefix.
Parameters	<i>name</i> — 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, etc.), the entire string must be enclosed within double quotes.

protocol

Syntax	protocol { <i>protocol</i> } [all instance <i>instance</i>] no protocol
Context	config>router>policy-options>policy-statement>entry>from config>router>policy-options>policy-statement>entry>from
Description	<p>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.</p> <p>If no protocol criterion is specified, any protocol is considered a match.</p> <p>The no form of the command removes the protocol match criterion.</p>
Default	no protocol — Matches any protocol.
Parameters	<p>protocol — The protocol name to match on.</p> <p>Values direct, static, bgp, isis, ospf, rip, aggregate, bgp-vpn, igmp, pim, ospf3, ldp, sub-mgmt, mld, managed, vpn-leak, tms, nat, periodic, ipsec</p> <p>instance — The OSPF or IS-IS instance.</p> <p>Values 1 — 31</p> <p>all — OSPF- or ISIS-only keyword.</p>

Route Policy Entry Match Commands

protocol

Syntax	protocol { <i>protocol</i> } [all instance <i>instance</i>] no protocol
Context	config>router>policy-options>policy-statement>entry>to config>router>policy-options>policy-statement>entry>from
Description	<p>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.</p> <p>If no protocol criterion is specified, any protocol is considered a match.</p> <p>The no form of the command removes the protocol match criterion.</p>
Default	no protocol — Matches any protocol.
Parameters	<p>protocol — The protocol name to match on.</p> <p>Values bgp, isis, ospf, rip, bgp-vpn, ospf3, vpn-leak, ldp</p> <p>instance — The OSPF or IS-IS instance.</p> <p>Values 1 — 31</p> <p>all — OSPF- or ISIS-only keyword.</p>

source-address

Syntax	source-address <i>ip-address</i> source-address prefix-list <i>prefix-list-name</i> no source-address										
Context	config>router>policy-options>policy-statement>entry>from										
Description	<p>This command specifies the source address that is embedded in the join or prune packet as a filter criterion.</p> <p>The no form of the command removes the criterion from the configuration.</p>										
Default	none										
Description	This command specifies a multicast data source address as a match criterion for this entry.										
Parameters	<p><i>ip-address</i> — The IP prefix for the IP match criterion in dotted decimal notation.</p> <p>Values</p> <table><tr><td>ipv4-address</td><td>a.b.c.d</td></tr><tr><td>ipv6-address</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr><tr><td></td><td>x:x:x:x:x:d.d.d.d</td></tr><tr><td></td><td>x - [0..FFFF]H</td></tr><tr><td></td><td>d - [0..255]D</td></tr></table> <p><i>prefix-list-name</i> — Specifies the prefix list name up to 32 characters in length.</p>	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..FFFF]H		d - [0..255]D
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..FFFF]H										
	d - [0..255]D										

state

Syntax	state <i>state</i> no state
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>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:</p> <ul style="list-style-type: none"> • subscriber-interface routes • subscriber-management routes (/32 IPv4 and IPv6 PD wan-host) • managed-routes (applicable only to IPv4). <p>Based on the state attribute of the route we can manipulate the route advertisement into the network.</p> <p>We can enable or disable (in case there is no SRRP running) tracking of SRRP state by routes.</p> <p>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).</p> <p>For subscriber-management and managed-routes, tracking is enabled per group interface under which SRRP is enabled.</p>
Default	none
Description	This command specifies a multicast data source address as a match criterion for this entry.
Parameters	<p>srrp-master — Track routes with the state attribute carrying srrp-master state.</p> <p>srrp-non-master — Track routes with the state attribute carrying srrp-non-master state.</p> <p>ipsec-master-with-peer — Track routes with the state attribute carrying ipsec-master-with-peer state.</p> <p>ipsec-non-master — Track routes with the state attribute carrying ipsec-non-master state.</p> <p>ipsec-master-without-peer — Track routes with the state attribute carrying ipsec-master-without-peer state.</p>

tag

Syntax	tag <i>tag</i> no tag
Context	config>router>policy-options>policy-statement>entry>from
Description	<p>This command matches the tag value on routes of type static, periodic as well as the tag field on routes learned through an external LSA.</p> <p>The no form of the command removes the tag field match criterion.</p>
Default	no tag — Matches any external LSA tag field.

Route Policy Entry Match Commands

Parameters *tag* — Matches a specific external LSA tag field.

Values **no-tag**, 1 — 4294967295

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.

Parameters **1** — Matches OSPF routes with type 1 LSAs.
2 — Matches OSPF routes with type 2 LSAs.

Route Policy Action Commands

action

Syntax	action { accept next-entry next-policy reject } no action
Context	config>router>policy-options>policy-statement>entry
Description	<p>This command creates the context to configure actions to take for routes matching a route policy statement entry.</p> <p>This command is required and must be entered for the entry to be active.</p> <p>Any route policy entry without the action command will be considered incomplete and will be inactive.</p> <p>The no form of the command deletes the action context from the entry.</p>
Default	no action — No action is defined.
Parameters	<p>accept — Specifies routes matching the entry match criteria will be accepted and propagated.</p> <p>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).</p> <p>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).</p> <p>reject — Specifies routes matching the entry match criteria would be rejected.</p>

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	<p>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.</p> <p>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.</p>

Route Policy Action Commands

Default **no default**

Parameters *send-limit* — Specify the maximum number of paths to advertise for matched routes to an Add-Paths peer.
 Values 1—16

aigp-metric

Syntax **aigp-metric** {*number1* | **igp** | *add number2*}
 no aigp-metric

Context config>router>policy-options>policy-statement>entry>action
 config>router>policy-options>policy-statement>default-action

Description 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 *number1* — Administratively defined metric.

Values 0 — 4294967295

Default none

number2 — Administratively defined metric.

Values 1 — 4294967295

Default none

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 — The AS path attribute is not changed.
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> — 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, etc.), the entire string must be enclosed within double quotes. Policy parameters must be enclosed by “@” symbols; for example, “@peerAS@”. The <i>name</i> specified must already be defined.

as-path-prepend

Syntax	as-path-prepend <i>as-num</i> [<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 — no AS number prepending configured.
Parameters	<i>as-num</i> — The AS number to prepend expressed as a decimal integer. Values 1 — 4294967295 <i>repeat</i> — The number of times to prepend the specified AS number expressed as a decimal integer. Values 1 — 50

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	<p>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 that 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.</p> <p>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.</p>
Default	no default

community

Syntax	community {{ add <i>name</i> [remove <i>name</i>]} { remove <i>name</i> [add <i>name</i>]} { replace <i>name</i> }} no community
Context	config>router>policy-options>policy-statement>default-action config>router>policy-options>policy-statement>entry>action
Description	<p>This command adds or removes a BGP community list to or from routes matching the route policy statement entry.</p> <p>If no community list is specified, the community path attribute is not changed.</p> <p>The community list changes the community path attribute according to the add and remove keywords.</p> <p>The no form of the command disables the action to edit the community path attribute for the route policy entry.</p>
Default	no community — The community path attribute is not changed.
Parameters	<p>add — The specified community list is added to any existing list of communities.</p> <p>remove — The specified community list is removed from the existing list of communities.</p> <p>replace — The specified community list replaces any existing community attribute.</p> <p><i>name</i> — The community 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, etc.), the entire string must be enclosed within double quotes.</p>

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 — Use the default damping profile.
Parameters	<i>name</i> — 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, etc.), the entire string must be enclosed within double quotes. The <i>name</i> specified must already be defined. none — Disables route damping for the route policy.

fc

Syntax	fc <i>fc</i> [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	<i>fc</i> — Specify the name of one of the predefined forwarding classes in the system. <table> <tr> <td>Values</td><td>be, l2, af, l1, h2, ef, h1, nc</td></tr> <tr> <td>Default</td><td>none (no QoS information is associated with matched routes)</td></tr> </table> <p>priority {low high} — This parameter associates an enqueueing priority with routes matched by the policy entry. Specifying a priority is optional.</p> <table> <tr> <td>Values</td><td>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</td></tr> </table>	Values	be, l2, af, l1, h2, ef, h1, nc	Default	none (no QoS information is associated with matched routes)	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
Values	be, l2, af, l1, h2, ef, h1, nc						
Default	none (no QoS information is associated with matched routes)						
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						

Route Policy Action Commands

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 reject} no default-action
Context	config>router>policy-options>policy-statement
Description	<p>This command enables the context to configure actions for routes that do not match any route policy statement entries when the accept parameter is specified.</p> <p>The default action clause 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.</p> <p>If a default action is defined and no match(es) occurred with the entries in the policy then the default action clause is used.</p> <p>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.</p> <p>The no form of the command deletes the default-action context for the policy statement.</p>
Default	no default-action — No default action is specified.
Parameters	<p>accept — Specifies routes matching the entry match criteria will be accepted and propagated.</p> <p>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).</p> <p>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).</p> <p>reject — Specifies routes matching the entry match criteria would be rejected.</p>

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	<p>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.</p> <p>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</p>

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, 6PE, VPN-IPv4, and VPN-IPv6.

Default **no default-action** — No default action is specified.

local-preference

Syntax **local-preference** *preference*
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** — BGP default preference is assigned.

Parameters *preference* — The local preference expressed as a decimal integer.

Values 0 — 4294967295

metric

Syntax **metric** {add|subtract} *metric*
metric set [igp|*metric-value*]
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** — Uses the configured metric (if defined) or do not advertise a metric.

Route Policy Action Commands

Parameters	add — Specified <i>integer</i> 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 — 4294967295	

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	<p>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.</p> <p>Once 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.</p> <p>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.</p> <p>Note that when 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.</p>
Default	disabled
Parameters	<p>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.</p> <p><i>ip-int-name</i> — specifies the alternate interface to which IGMP messages are redirected.</p>

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 — The next hop attribute is not changed.		
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 — 32

next-hop-self

Syntax	[no] next-hop-self		
Context	config>router>policy-options>policy-statement <i>name</i> >default-action config>router>policy-options>policy-statement>entry>action		
Description	This command advertises a next hop IP address belonging to this router even if a third-party next hop is available to routes matching the policy statement entry. The no form of the command disables advertising the next-hop-self option for the route policy entry.		
Default	no next-hop-self — The next hop IP address is not changed.		

next-hop-self

Syntax	[no] next-hop-self [multihoming primary-anycast secondary-anycast]		
Context	config>router>policy-option>policy-statement>entry>action		
Description	This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer. This is primarily used to avoid third-party route advertisements when connected to a multi-access network. In addition, this command can be used to enable and configure the multi-homing reliency mechanism replacing the usual BGP nexthop with a configured anycast address. The no form of the command returns the setting of the BGP next-hop attribute to the default value determined by the BGP protocol.		
Default	no next-hop-self		

Route Policy Action Commands

Parameters *primary-anycast* — Specifies the anycast address that the local node will use to replace the BGP nexthop address in route updates associated peers.

secondary-address — Specifies the anycast address that the local node is to track.

origin

Syntax **origin {igp | egp | incomplete}**
 no origin

Context config>router>policy-options>policy-statement *name*>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.

origin-validation-state

Syntax **origin-validation-state state**
 no origin-validation-state

Context config>router>policy-options>policy-statement *name*>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 **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 <i>name</i> >default-action config>router>policy-options>policy-statement>entry>action>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.
Default	no preference — No route preference is assigned by the policy entry. The protocol default preference is used.
Parameters	<i>preference</i> — The route preference expressed as a decimal integer. Values 1 — 255 (0 represents unset - MIB only)

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 preference

tag

Syntax	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 an OSPF tag to routes matching the entry. The tag value is used to apply a tag to a route for either an OSPF or RIP route. A hexadecimal value of 4 octets can be entered. For OSPF, all four octets can be used. For RIP, only the two most significant octets are used if more than two octets are configured.

Route Policy Action Commands

The **no** form of the command removes the tag.

Default **no tag**

Parameters *tag* — Assigns an OSPF, RIP or ISIS tag to routes matching the entry.

Values Accepts decimal or hex formats:
OSPF and ISIS: [0x0..0xFFFFFFFF]H
RIP: [0x0..0xFFFF]H

type

Syntax **type** {*type*}
 no type

Context config>router>policy-options>policy-statement *name*>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 2

Parameters *type* — 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.

Show Commands

policy

Syntax `policy [name | prefix-list name | admin]`

Context `show>router`

Description This command displays configured policy statement information.

Parameters **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.

prefix-list *name* — Displays the prefix lists configured in the route policy.

admin — If the keyword **admin** is included, the entire policy option configuration displays, including any un-committed configuration changes. This command is similar to the **info** command.

Output **Route Policy Output** — The following table describes 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 on page 967](#)
- [show router policy admin on page 968](#)
- [show router policy “BGP To RIP” on page 969](#)
- [show router policy damping on page 969](#)

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 :
=====
```

Show Commands

```
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 0x8008135
    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#
```

policy-test

Syntax	policy <i>policy-name</i> family <i>family</i> prefix <i>prefix/pfxlen</i> [longer] neighbor <i>neighbor</i> [display-rejects] [detail]
Context	show>router>bpg>routes
Description	<p>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. When a prefix is rejected, the entry ID that rejected it is displayed.</p> <p>When the policy evaluation tool is applied to vpn-ipv4 or vpn-ipv6 family types, a warning will be displayed to the operator about the possibility of hidden routes. If the vrf-import policy that is being evaluated is attempting to import RTs that were never before imported on the node, the routes will not be displayed (accepted or rejected) because they are not available in the RTM until the vrf-import policy is activated. This is due to the architecture of SR-OS storing MP-BGP routes in the RTM.</p> <p>Configuration of ‘mp-bgp-keep’ under the BGP hierarchy will no longer suppress these routes by default, at the cost of higher memory requirements to store all MP-BGP routes received from MP-BGP neighbors.</p>
Parameters	<p><i>policy name</i> — Specifies information of an existing configured and committed policy.</p> <p><i>family</i> — Specifies ipv4 or ipv6 information. Default is ipv4.</p> <p><i>prefix</i> — Specifies IPv4 or IPv6 prefix/mask to be evaluated. The longer keyword may be specified to evaluate longer prefix matches. (Optional)</p> <p><i>neighbor</i> — Specifies BGP neighbor (optional).</p> <p>display-rejects — Displays routes that were rejected by the policy. If this is not specified, only a count of rejected routes is shown.</p> <p>detail — Displays changes in the policy attributes. When the policy modifies route attributes, display the modifications made to the routes. This command requires an exact prefix to be specified. (Optional)</p>

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
  4.0.0.6/32                               None       None
  220.0.0.2                               None       -
  14
```

```

-----
Total Routes : 17 Routes rejected : 16
=====

A:sim-1# show router bgp policy-test bgpprefix6 prefix 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
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
2.2.2.2/32
Rejected by Default action
4.0.0.1/32
Rejected by Default action
4.0.0.2/32
Rejected by Default action
4.0.0.3/32
Rejected by Default action
4.0.0.4/32
Rejected by Default action
4.0.0.5/32
Accepted by Policy
4.0.0.6/32
Rejected by Default action
6.0.0.1/32
Rejected by Default action
7.0.0.1/32
Rejected by Default action
10.0.4.0/24
Rejected by Default action
10.12.0.0/24
Rejected by Default action

```

Show Commands

```
10.14.0.0/24
Rejected by Default action
10.24.0.0/24
Rejected by Default action
12.12.12.12/32
Rejected by Default action
220.0.0.2/32
Rejected by Default action
220.0.0.3/32
Rejected by Default action
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
2.2.2.2/32                        None       None
220.0.0.2                        None       -
14
Rejected by Default action
4.0.0.1/32                       None       None
220.0.0.2                        None       -
14
Rejected by Default action
4.0.0.2/32                       None       None
220.0.0.2                        None       -
14
Rejected by Default action
4.0.0.3/32                       None       None
220.0.0.2                        None       -
14
Rejected by Default action
4.0.0.4/32                       None       None
220.0.0.2                        None       -
14
Rejected by Default action
4.0.0.5/32                       None       None
220.0.0.2                        None       -
14
Accepted by Policy
4.0.0.6/32                       None       None
220.0.0.2                        None       -
14
Rejected by Default action
```



```

        6.0.0.1/32                                None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        7.0.0.1/32                                None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        10.0.4.0/24                               None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        10.12.0.0/24                              None      20
        220.0.0.2                                None      -
        14
Rejected by Default action
        10.14.0.0/24                              None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        10.24.0.0/24                              None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        12.12.12.12/32                            None      20
        220.0.0.2                                None      -
        14
Rejected by Default action
        220.0.0.2/32                              None      None
        220.0.0.2                                None      -
        14
Rejected by Default action
        220.0.0.3/32                              None      10
        220.0.0.2                                None      -
        14
Rejected by Default action
        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
Nexthop      : 220.0.0.2
Path Id      : None
From         : 220.0.0.2

```

Show Commands

```
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            :
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       :
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
Fwd Class       : None
Flags           :
Route Source     : External
AS-Path         : 14

Interface Name   : to-sim-6
Aggregator      : None
MED             : None
Peer Router Id  : 14.14.14.10
Priority         : None

```

```

-----
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

```

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

```

```

4.0.0.6/32
-----

```

```

Routes : 1

```

Show Commands

```
=====
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
  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
  4.0.0.1/32
Rejected by Default action
  4.0.0.2/32
Rejected by Default action
  4.0.0.3/32
Rejected by Default action
  4.0.0.4/32
Rejected by Default action
  4.0.0.5/32
Accepted by Policy
  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    VPNLabel
      As-Path
-----
Rejected by Default action
  4.0.0.1/32                               None       None
    220.0.0.2                             None       -
      14
Rejected by Default action
  4.0.0.2/32                               None       None
    220.0.0.2                             None       -
      14
Rejected by Default action
  4.0.0.3/32                               None       None
    220.0.0.2                             None       -
      14
Rejected by Default action
  4.0.0.4/32                               None       None
    220.0.0.2                             None       -
      14
Rejected by Default action
  4.0.0.5/32                               None       None
    220.0.0.2                             None       -
      14
Accepted by Policy
  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
  4.0.0.1/32
Rejected by Default action
  4.0.0.2/32
Rejected by Default action
  4.0.0.3/32
Rejected by Default action
  4.0.0.4/32
Rejected by Default action
  4.0.0.5/32
Accepted by Policy
  4.0.0.6/32

```

Show Commands

```
-----
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
  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 bgpprefix44rej 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
  1:30:192.14.15.0/24                    None       None
  220.0.0.2                              None       131069
  14
Accepted by Policy
  65530:20:8.0.0.1/32                    None       None
  220.0.0.2                              None       131070
  14
Accepted by Policy
  65530:20:10.0.3.0/24                    None       None
  220.0.0.2                              None       131070
  14 101
Accepted by Policy
  65530:20:10.13.0.0/24                    None       None
```

Route Policies

220.0.0.2	None	131070
14 101		
Accepted by Policy		
65530:20:10.23.0.0/24	None	None
220.0.0.2	None	131070
14 101		
Accepted by Policy		
65530:20:13.13.13.13/32	None	None
220.0.0.2	None	131070
14 101		
Accepted by Policy		
65530:20:20.20.20.5/32	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:20.20.20.6/32	None	None
220.0.0.2	None	131070
14		
Rejected by Policy Entry = 10		
65530:20:44.44.44.0/24	None	None
220.0.0.2	None	131070
14 101		
Accepted by Policy		
65530:20:192.14.15.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.16.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.17.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.18.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.19.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.20.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.21.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.22.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		
65530:20:192.14.23.0/24	None	None
220.0.0.2	None	131070
14		
Accepted by Policy		

Show Commands

```

65530:20:192.14.25.0/24          None      None
220.0.0.2                      None      131070
14
Accepted by Policy
65530:20:196.34.0.0/24          None      None
220.0.0.2                      None      131070
14
Accepted by Policy
220.0.0.2:50:192.50.50.0/24     None      None
220.0.0.2                      None      131067
14
Accepted by Policy
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
=====

```

```
*A:Dut-C>config>router# show router bgp routes leakable
```

```

=====
BGP Router ID:10.20.1.3      AS:1      Local AS:1
=====

```

```
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 IPv4 Routes
=====

```

Flag	Network Nexthop (Router) As-Path	LocalPref Path-Id	MED Label
u*>i	9.9.9.9/32 1.3.5.5 3 6500	None None	None -
u*>i	10.10.10.10/32 1.3.5.5 3 6500	None None	None -
u*>i	10.20.1.5/32 1.3.5.5 3 6500	None None	None -
u*>i	11.11.11.11/32 1.3.5.5 3 6500	None None	None -
u*>i	12.12.12.12/32 1.3.5.5 3 6500	None None	None -
u*>i	13.13.13.13/32 10.20.1.4 10 6011	100 17	None -
u*>i	14.14.14.14/32 10.20.1.4 10 6012	100 14	None -
u*>i	15.15.15.15/32 10.20.1.4 10 6021	100 15	None -


```

u*>i 16.16.16.16/32          100      None
      10.20.1.4             16        -
      10 6022

```

```

-----
Routes : 9
=====

```

```

*A:Dut-C>config>router# show router bgp routes leaked

```

```

=====
BGP Router ID:10.20.1.3      AS:1      Local AS:1
=====
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 IPv4 Routes
=====

```

Flag	Network Nexthop (Router) As-Path	LocalPref Path-Id	MED Label
u*>li	1.1.1.1/32	100	None
	39.1.1.2 (VPRN 1)	None	-
	3011		
u*>li	2.2.2.2/32	100	None
	39.1.2.2 (VPRN 1)	None	-
	3012		
u*>li	3.3.3.3/32	100	None
	39.2.1.2 (VPRN 1)	None	-
	3021		
u*>li	4.4.4.4/32	100	None
	39.2.2.2 (VPRN 1)	None	-
	3022		

```

-----
Routes : 4
=====

```

```

A:pe1# show router bgp routes vpn-ipv4 2.0.0.0/8 detail

```

```

=====
BGP Router ID:192.0.2.1      AS:65001      Local AS:65001
=====
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
=====

```

```

-----
Original Attributes

```

Network	: 2.0.0.0/8		
Nexthop	: 192.0.2.2		
Route Dist.	: 65001:2	VPN Label	: 262142
Path Id	: None		
From	: 192.0.2.2		

Show Commands

```
Res. Nexthop      : n/a
Local Pref.       : 100
Aggregator AS    : None
Atomic Aggr.     : Not Atomic
AIGP Metric      : None
Connector        : None
Community        : target:65001:1
Cluster          : No Cluster Members
Originator Id    : None
Fwd Class        : None
Flags            : Used 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    : 01h48m27s
VPRN Imported    : 1
```

Interface Name : to_4008
Aggregator : None
MED : None

Peer Router Id : 192.0.2.2
Priority : None

Dest Class : 0

Modified Attributes

```
Network          : 2.0.0.0/8
Nexthop          : 192.0.2.2
Route Dist.      : 65001:2
Path Id          : None
From             : 192.0.2.2
Res. Nexthop     : n/a
Local Pref.      : 100
Aggregator AS    : None
Atomic Aggr.     : Not Atomic
AIGP Metric      : None
Connector        : None
Community        : target:65001:1
Cluster          : No Cluster Members
Originator Id    : None
Fwd Class        : None
Flags            : Used 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   : 3
Last Modified    : 01h49m47s
VPRN Imported    : 1
```

VPN Label : 262142

Interface Name : to_4008
Aggregator : None
MED : None

Peer Router Id : 192.0.2.2
Priority : None

Dest Class : 0

Standards and Protocol Support

Note that the information presented is subject to change without notice.
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Ethernet Standards

IEEE 1588 Precision Clock Synchronization Protocol
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 Ethernet Linear Protection Switching
ITU-T G.8032 Ethernet Ring Protection Switching
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks

OSPF

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 (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 4576 Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)
RFC 4970 Extensions to OSPF for Advertising Optional Router Capabilities
RFC 5185 OSPF Multi-Area Adjacency
RFC 5243 OSPF Database Exchange Summary List Optimization
RFC 5250 The OSPF Opaque LSA Option
RFC 5709 OSPFv2 HMAC-SHA Cryptographic Authentication
RFC 6987 OSPF Stub Router Advertisement

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening

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 BGP4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP
RFC 4486 Subcodes for BGP Cease Notification Message
RFC 4577 OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)
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 – GR helper
RFC 4760 Multi-protocol Extensions for BGP
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 Confederations for BGP (obsoletes 3065)
RFC 5291 Outbound Route Filtering Capability for BGP-4

RFC 5575 Dissemination of Flow Specification Rules
RFC 5668 4-Octet AS Specific BGP Extended Community
draft-ietf-idr-add-paths Advertisement of Multiple Paths in BGP
draft-ietf-idr-best-external Advertisement of the Best External Route in BGP

IS-IS

ISO/IEC 10589:2002, Second Edition, Nov. 2002 Intermediate System to Intermediate System Intra-Domain Routing Information Exchange Protocol
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 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 6329 IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
draft-ietf-isis-mi-02 IS-IS Multi-Instance

IP, LDP, and Segment Routing Fast Reroute (FRR)

RFC 5286 Basic Specification for IP Fast Reroute: Loop-Free Alternates
draft-ietf-isis-segment-routing-extensions-03 IS-IS Extensions for Segment Routing
draft-ietf-rtgwg-lfa-manageability-07 Operational management of Loop Free Alternates
draft-ietf-rtgwg-remote-lfa-09 Remote LFA FRR
draft-kratn-mofrr-02 Multicast only Fast Re-Route

IPSec

RFC 2401 Security Architecture for the Internet Protocol
RFC 2406 IP Encapsulating Security Payload (ESP)
RFC 2409 The Internet Key Exchange (IKE)
RFC 2560 X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP
RFC 3706 IKE Dead Peer Detection
RFC 3947 Negotiation of NAT-Traversal in the IKE
RFC 3948 UDP Encapsulation of IPsec ESP Packets
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 5996 Internet Key Exchange Protocol Version 2 (IKEv2)
RFC 5998 An Extension for EAP-Only Authentication in IKEv2

draft-ietf-ipsec-isakmp-xauth-06 Extended Authentication within ISAKMP/Oakley (XAUTH)
draft-ietf-ipsec-isakmp-modecfg-05 The ISAKMP Configuration Method

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2375 IPv6 Multicast Address Assignments
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing
RFC 2710 Multicast Listener Discovery (MLD) for IPv6
RFC 2740 OSPF for IPv6
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses
RFC 3315 Dynamic Host Configuration Protocol for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC 3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) 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 IPv6 Addressing Architecture
RFC 4443 Internet Control Message Protocol (ICMPv6)
for the Internet Protocol Version 6 (IPv6) Specification
RFC 4552 Authentication/Confidentiality for OSPFv3

RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
 RFC 5072 IP Version 6 over PPP
 RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
 RFC 5187 OSPFv3 Graceful Restart (Helper Mode)
 RFC 5308 Routing IPv6 with IS-IS
 RFC 5340 OSPF for IPv6
 RFC 5838 Support of Address Families in OSPFv3

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
 RFC 2236 Internet Group Management Protocol, (Snooping)
 RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)
 RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
 RFC 3446 Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
 RFC 3618 Multicast Source Discovery Protocol (MSDP)
 RFC 3956 Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address
 RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
 RFC 4604 Using IGMPv3 and 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 4624 Multicast Source Discovery Protocol (MSDP) MIB
 RFC 5059 Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)
 RFC 5059 Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)

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 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 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)
 RFC 7385 IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points
 draft-dolganow-l3vpn-mvpn-expl-track-00 Explicit tracking in MPLS/BGP IP VPN

MPLS — GENERAL

RFC 2430 A Provider Architecture DiffServ & TE
 RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)
 RFC 2597 Assured Forwarding PHB Group (rev3260)
 RFC 2598 An Expedited Forwarding PHB
 RFC 3031 MPLS Architecture
 RFC 3032 MPLS Label Stack Encoding
 RFC 3140 Per-Hop Behavior Identification Codes
 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

MPLS — LDP

RFC 3037 LDP Applicability
 RFC 3478 Graceful Restart Mechanism for LDP – GR helper
 RFC 5036 LDP Specification
 RFC 5283 LDP extension for Inter-Area LSP
 RFC 5443 LDP IGP Synchronization
 RFC 5561 LDP Capabilities
 RFC 6388 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP
 RFC 6826 Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
 draft-ietf-mpls-ldp-ip-pw-capability-09 Disabling IPoMPLS and P2P PW LDP Application's State Advertisement
 draft-ietf-mpls-ldp-ipv6-15 Updates to LDP for IPv6
 draft-pdutta-mpls-ldp-adj-capability-00 LDP Adjacency Capabilities
 draft-pdutta-mpls-ldp-v2-00 LDP Version 2
 draft-pdutta-mpls-multi-ldp-instance-00 Multiple LDP Instances
 draft-pdutta-mpls-tldp-hello-reduce-04 Targeted LDP Hello Reduction

MPLS/RSVP — TE

RFC 2702 Requirements for Traffic Engineering over MPLS
 RFC2747 RSVP Cryptographic Authentication
 RFC 2961 RSVP Refresh Overhead Reduction Extensions
 RFC3097 RSVP Cryptographic Authentication - Updated Message Type Value
 RFC 3209 Extensions to RSVP for Tunnels

RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling

Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions – (support of of IF_ID RSVP_HOP object with unnumbered interface and RSVP-TE Graceful Restart Helper Procedures)

RFC 3477 Signalling Unnumbered Links in Resource Reservation Protocol-Traffic Engineering (RSVP-TE)

RFC 3564 Requirements for Diff-Serv-aware TE

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 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 – RSVP-TE Extensions

RFC 5712 MPLS Traffic Engineering Soft Preemption

RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

MPLS — OAM

RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

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

MPLS — TP (7750/7450 only)

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-TP Identifiers

RFC 6378 MPLS-TP Linear Protection

RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile

RFC 6426 MPLS On-Demand Connectivity and Route Tracing

RFC 6478 Pseudowire Status for Static Pseudowires

RFC 7213 MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing

MPLS — GMPLS

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 GMPLS recovery

draft-ietf-ccamp-rsvp-te-srlg-collect-04 RSVP-TE Extensions for Collecting SRLG Information

RIP

RFC 1058 RIP Version 1

RFC 2080 RIPng for IPv6

RFC 2082 RIP-2 MD5 Authentication

RFC 2453 RIP Version 2

TCP/IP

RFC 768 UDP

RFC 791 IP

RFC 792 ICMP

RFC 793 TCP

RFC 826 ARP

RFC 854 Telnet

RFC 951 Bootstrap Protocol (BOOTP)

RFC 1350 The Tftp Protocol (revision 2)

RFC 1519 CIDR

RFC 1542 Clarifications and Extensions for the Bootstrap Protocol

RFC 1812 Requirements for IPv4 Routers

RFC 2347 TFTP option Extension

RFC 2328 TFTP Blocksize Option

RFC 2349 TFTP Timeout Interval and Transfer Size option

RFC 2401 Security Architecture for Internet Protocol

RFC 2428 FTP Extensions for IPv6 and NATs

RFC 3596 DNS Extensions to Support IP version 6

RFC 5880 Bidirectional Forwarding Detection

RFC 5881 BFD IPv4 and IPv6 (Single Hop)

RFC 5883 BFD for Multihop Paths

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

draft-ietf-vrrp-unified-spec-02 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

PPP

RFC 1332 PPP IPCP
 RFC 1377 PPP OSINLCP
 RFC 1638/2878 PPP BCP
 RFC 1661 PPP (rev RFC2151)
 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 2516 A Method for Transmitting PPP Over Ethernet
 RFC 2615 PPP over SONET/SDH
 RFC 2686 The Multi-Class Extension to Multi-Link PPP

Frame Relay

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement
 FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation
 ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.
 FRF2.2 PVC Network-to- Network Interface (NNI) Implementation Agreement.
 FRF.12 Frame Relay Fragmentation Implementation Agreement
 FRF.16.1 Multilink Frame Relay UNI/ NNI Implementation Agreement
 ITU-T Q.933, Annex A Additional procedures for Permanent Virtual Connection (PVC) status management

ATM

RFC 1626 Default IP MTU for use over ATM AAL5
 RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management
 RFC 2515 Definition of Managed Objects for ATM Management
 RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5

AF-TM-0121.000 Traffic Management Specification Version 4.1
 ITU-T Recommendation I.610 B-ISDN Operation and Maintenance Principles and Functions version 11/95
 ITU-T Recommendation I.432.1 BISDN user-network interface – Physical layer specification: General characteristics
 GR-1248-CORE Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3
 GR-1113-CORE Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1
 AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0
 AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR
 AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA) Specification Version 1.1

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)
 RFC 3046 DHCP Relay Agent Information Option (Option 82)
 RFC 1534 Interoperation between DHCP and BOOTP

Policy Management and Credit Control

3GPP TS 29.212 Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11 and Release 12) - Gx support as it applies to wireline environment (BNG)
 RFC 3588 Diameter Base Protocol
 RFC 4006 Diameter Credit Control Application

NAT

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 6888 Common Requirements For Carrier-Grade NATs (CGNs)

VPLS

RFC 4761 Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling
 RFC 4762 Virtual Private LAN Services Using LDP
 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)

Pseudowire

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 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)
 RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks
 RFC 4816 PWE3 ATM Transparent Cell Transport Service
 RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks
 RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks
 RFC 4446 IANA Allocations for PWE3
 RFC 4447 Pseudowire Setup and Maintenance Using LDP

Standards and Protocols

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) OAM Message Mapping
RFC 6391 Flow Aware Transport of Pseudowires over an MPLS PSN
RFC 6575 ARP Mediation for IP Interworking of Layer 2 VPN
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 OAM Interworking
RFC 7267 Dynamic Placement of Multi-Segment Pseudowires
draft-ietf-l2vpn-vpws-iw-oam-04 OAM Procedures for VPWS Interworking
MFA Forum 9.0.0 The Use of Virtual trunks for ATM/MPLS Control Plane 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

ANCP/L2CP

RFC 5851 ANCP framework
draft-ietf-ancp-protocol-02 ANCP Protocol

Voice /Video Performance:

ITU-T G.107 The E Model- A computational model for use in planning.
ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an

Equipment Impairment Factor using Passive Monitoring
ITU-T Rec. P.564 Conformance testing for voice over IP transmission quality assessment models
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.
RFC 3550, Appendix A.8 RTP: A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter.

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)
MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004
RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

SONET/SDH

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

AAA

RFC 2865 Remote Authentication Dial In User Service
RFC 2866 RADIUS Accounting
draft-grant-tacacs-02 The TACACS+ Protocol

SSH

RFC 4250 The Secure Shell (SSH) Protocol Assigned Numbers
RFC 4251 The Secure Shell (SSH) Protocol Architecture

RFC 4254 The Secure Shell (SSH) Connection Protocol

OpenFlow

ONF OpenFlow Switch Specification Version 1.3.1 (Hybrid-switch/FlowTable)

Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000
ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008
ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.
GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005
ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.
ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.
ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.
ITU-T G.8265.1 Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010.
IEEE 1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

Network Management

ITU-T X.721 Information technology-
OSI-Structure of Management
Information

ITU-T X.734 Information technology-
OSI-Systems Management: Event
Report Management Function

M.3100/3120 Equipment and Connection
Models

TMF 509/613 Network Connectivity
Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining
Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information
Base for the Transmission Control
Protocol

RFC 2465 Management Information
Base for IPv6: Textual Conventions
and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-
NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-
SMMIB

RFC 2575 SNMP-VIEW-BASED-ACM-
MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2578 Structure of Management
Information Version 2 (SMIv2)

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOGMIB

RFC 3019 IP Version 6 Management
Information Base for The Multicast
Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for
Describing Simple Network

Management Protocol (SNMP)
Management Frameworks

RFC 3412 Message Processing and
Dispatching for the Simple Network
Management Protocol (SNMP)

RFC 3413 Simple Network Management
Protocol (SNMP) Applications

RFC 3414 User-based Security Model
(USM) for version 3 of the Simple
Network Management Protocol
(SNMPv3)

RFC 3418 SNMP MIB

RFC 3826 The Advanced Encryption
Standard (AES) Cipher Algorithm in
the SNMP User-based Security
Model

RFC 4113 Management Information
Base for the User Datagram Protocol
(UDP)

RFC 4292 IP Forwarding Table MIB

RFC 4293 MIB for the Internet Protocol

RFC 5101 Specification of the IP Flow
Information Export (IPFIX)
Protocol for the Exchange of IP
Traffic Flow Information

RFC 6241 Network Configuration
Protocol (NETCONF)

RFC 6242 Using the NETCONF Protocol
over Secure Shell (SSH)

draft-ietf-bfd-mib-00 Bidirectional
Forwarding Detection Management
Information Base

draft-ietf-isis-wg-mib-06 Management
Information Base for Intermediate
System to Intermediate System (IS-
IS)

draft-ietf-ospf-mib-update-04 OSPF
Version 2 Management Information
Base

draft-ietf-mboned-msdp-mib-01
Multicast Source Discovery protocol
MIB

draft-ietf-mppls-lsr-mib-06 Multiprotocol
Label Switching (MPLS) Label
Switching Router (LSR)
Management Information Base

draft-ietf-mppls-te-mib-04 Multiprotocol
Label Switching (MPLS) Traffic
Engineering Management
Information Base

draft-ietf-mppls-ldp-mib-07 Definitions of
Managed Objects for the
Multiprotocol Label Switching,
Label Distribution Protocol (LDP)

IEEE 802.3ad MIB

Customer documentation and product support



Customer documentation

<http://documentation.alcatel-lucent.com>



Technical support

<http://support.alcatel-lucent.com>



Documentation feedback

documentation.feedback@alcatel-lucent.com

