

Junos® OS

OSPF User Guide

Published
2025-12-16

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos® OS OSPF User Guide

Copyright © 2025 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

[About This Guide | xviii](#)

1

OSPF Overview

[Introduction to OSPF | 2](#)

[OSPF Overview | 2](#)

[OSPF Packets Overview | 7](#)

[Understanding OSPF External Metrics | 10](#)

[Supported OSPF and OSPFv3 Standards | 11](#)

2

Understand OSPF Configurations

[Understanding OSPF Configurations | 14](#)

3

Configure OSPF Interfaces

[Configuring OSPF Interfaces | 16](#)

[About OSPF Interfaces | 16](#)

[Example: Configuring an Interface on a Broadcast or Point-to-Point Network | 18](#)

[Requirements | 18](#)

[Overview | 18](#)

[Configuration | 19](#)

[Verification | 22](#)

[Example: Configuring OSPF Demand Circuits | 23](#)

[Requirements | 23](#)

[Overview | 23](#)

[Configuration | 24](#)

[Verification | 26](#)

[Example: Configuring a Passive OSPF Interface | 27](#)

[Requirements | 27](#)

[Overview | 27](#)

[Configuration | 28](#)

[Verification | 29](#)

Example: Configuring OSPFv2 Peer interfaces | 30

Requirements | 30

Overview | 31

Configuration | 31

Verification | 32

Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network | 33

Requirements | 33

Overview | 34

Configuration | 35

Verification | 37

Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network | 38

Requirements | 38

Overview | 38

Configuration | 39

Verification | 41

Understanding Multiple Address Families for OSPFv3 | 41

Example: Configuring Multiple Address Families for OSPFv3 | 42

Requirements | 42

Overview | 43

Configuration | 44

Verification | 47

4

Configure OSPF Areas

Configuring OSPF Areas | 49

Understanding OSPF Areas | 50

OSPF Designated Router Overview | 53

Example: Configuring an OSPF Router Identifier | 54

Requirements | 54

Overview | 55

Configuration | 55

Verification | 57

Example: Controlling OSPF Designated Router Election | 57

Requirements | 57

Overview | 57

Configuration | 58

Verification | 59

Understanding OSPF Areas and Backbone Areas | 60

Example: Configuring a Single-Area OSPF Network | 62

Requirements | 62

Overview | 62

Configuration | 63

Verification | 64

Example: Configuring a Multiarea OSPF Network | 65

Requirements | 65

Overview | 66

Configuration | 67

Verification | 71

Understanding Multiarea Adjacency for OSPF | 71

Example: Configuring Multiarea Adjacency for OSPF | 72

Requirements | 72

Overview | 72

Configuration | 74

Verification | 77

Understanding Multiarea Adjacencies for OSPFv3 | 79

Example: Configuring a Multiarea Adjacency for OSPFv3 | 79

Requirements | 80

Overview | 80

Configuration | 81

Verification | 87

Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas | 89

Example: Configuring OSPF Stub and Totally Stubby Areas | 91

Requirements | 91

Overview | 92

Configuration | 93

Verification | 96

Example: Configuring OSPF Not-So-Stubby Areas | 97

Requirements | 97

Overview | 97

Configuration | 99

Verification | 103

Understanding OSPFv3 Stub and Totally Stubby Areas | 105

Example: Configuring OSPFv3 Stub and Totally Stubby Areas | 105

Requirements | 105

Overview | 106

Configuration | 107

Verification | 117

Understanding OSPFv3 Not-So-Stubby Areas | 121

Example: Configuring OSPFv3 Not-So-Stubby Areas | 121

Requirements | 121

Overview | 121

Configuration | 123

Verification | 134

Understanding Not-So-Stubby Areas Filtering | 140

Example: Configuring OSPFv3 Not-So-Stubby Areas with Filtering | 141

Requirements | 141

Overview | 141

Configuration | 142

Verification | 148

Understanding OSPF Virtual Links for Noncontiguous Areas | 152

Example: Configuring OSPF Virtual Links to Connect Noncontiguous Areas | 153

Requirements | 153

Overview | 153

Configuration | 154

Verification | 158

Example: Configuring OSPFv3 Virtual Links | 159

Requirements | 159

Overview | 159

Configuration | 160

Verification | 174

Configure OSPF Route Control

Configuring OSPF Route Control | 195

Understanding OSPF Route Summarization | 195

Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements Sent into the Backbone Area | 196

Requirements | 196

Overview | 197

Configuration | 199

Verification | 204

Example: Limiting the Number of Prefixes Exported to OSPF | 205

Requirements | 205

Overview | 206

Configuration | 206

Verification | 208

Understanding OSPF Traffic Control | 208

Example: Controlling the Cost of Individual OSPF Network Segments | 211

Requirements | 211

Overview | 211

Configuration | 213

Verification | 216

Understanding Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors | 217

Example: Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors | 218

Example Prerequisites | 219

Before You Begin | 219

Functional Overview | 220

Topology Overview | 220

Topology Illustration | 221

R0 Configuration Steps | 221

Verification | 224

Appendix 1: Set Commands on All Devices | 229

Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth | 231

Requirements | 233

Overview | 234

Verification | 234

Example: Controlling OSPF Route Preferences | 235

Requirements | 236

Overview | 237

Verification | 238

Understanding OSPF Overload Function | 238

Example: Configuring OSPF to Make Routing Devices Appear Overloaded | 240

Requirements | 241

Overview | 241

Configuration | 242

Verification | 243

Understanding the SPF Algorithm Options for OSPF | 245

Example: Configuring SPF Algorithm Options for OSPF | 246

Requirements | 246

Overview | 246

Configuration | 247

Verification | 249

Configuring OSPF Refresh and Flooding Reduction in Stable Topologies | 250

Understanding Synchronization Between LDP and IGP | 252

Example: Configuring Synchronization Between LDP and OSPF | 252

Requirements | 252

Overview | 253

Configuration | 254

Verification | 257

OSPFv2 Compatibility with RFC 1583 Overview | 257

Example: Disabling OSPFv2 Compatibility with RFC 1583 | 258

Requirements | 258

Overview | 258

Configuration | 259

Verification | 260

Configure OSPF Authentication

Configuring OSPF Authentication | 262

Understanding IPsec Authentication for OSPF Packets on EX Series Switches | 262

Understanding OSPFv2 Authentication | 265

Understanding OSPFv3 Authentication | 268

Example: Configuring Simple Authentication for OSPFv2 Exchanges | 270

Requirements | 270

Overview | 270

Configuration | 271

Verification | 272

Example: Configuring MD5 Authentication for OSPFv2 Exchanges | 273

Requirements | 273

Overview | 274

Configuration | 274

Verification | 276

Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface | 277

Requirements | 277

Overview | 278

Configuration | 279

Verification | 281

Using IPsec to Secure OSPFv3 Networks (CLI Procedure) | 282

Configuring Security Associations | 282

Securing OPSFv3 Networks | 283

Example: Configuring IPsec Authentication for an OSPF Interface | 284

Requirements | 284

Overview | 284

Configuration | 287

7

Verification | 291

Configure OSPF Routing Instances

Configuring OSPF Routing Instances | 294

Understanding OSPF Routing Instances | 294

Installing Routes from OSPF Routing Instances into the OSPF Routing Table Group | 296

Example: Configuring Multiple Routing Instances of OSPF | 297

Requirements | 297

Overview | 297

Configuration | 300

Verification | 306

Multiple Independent IGP Instances of OSPFv2 | 306

Example: Configure Multiple Independent Instances of OSPFv2 with Segment Routing | 309

Example Prerequisites | 309

Before You Begin | 310

Functional Overview | 310

Topology Overview | 311

Topology Illustration | 311

R2 Configuration Steps | 311

Verification | 315

Appendix 1: Set Commands on All Devices | 321

8

Configure OSPF Timers

Configuring OSPF Timers | 325

OSPF Timers Overview | 325

Example: Configuring OSPF Timers | 326

Requirements | 327

Overview | 327

Configuration | 328

Verification | 334

9

Configure OSPF Fault Detection using BFD

Configuring OSPF Fault Detection using BFD | 336

Understanding BFD for OSPF | 336

Example: Configuring BFD for OSPF | 338

Requirements | 338

Overview | 339

Configuration | 340

Verification | 343

Understanding BFD Authentication for OSPF | 344

Configuring BFD Authentication for OSPF | 346

Configuring BFD Authentication Parameters | 346

Viewing Authentication Information for BFD Sessions | 348

10

Configure Graceful Restart for OSPF

Configuring Graceful Restart for OSPF | 352

Graceful Restart for OSPF Overview | 352

Example: Configuring Graceful Restart for OSPF | 354

Requirements | 354

Overview | 354

Configuration | 355

Verification | 359

Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart | 361

Requirements | 361

Overview | 361

Configuration | 362

Verification | 366

Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart | 367

Requirements | 367

Overview | 367

Configuration | 368

Verification | 371

Example: Disabling Strict LSA Checking for OSPF Graceful Restart | 372

Requirements | 372

Overview | 373

Configuration | 373

Verification | 376

Configure Loop-Free Alternate Routes for OSPF

Configuring Loop-Free Alternate Routes for OSPF | 379

Per Prefix Loop Free Alternates for OSPF | 379

Configuring Per-Prefix LFA for OSPF | 380

Loop-Free Alternate Routes for OSPF Overview | 381

Example: Configuring Loop-Free Alternate Routes for OSPF | 382

Requirements | 383

Overview | 383

Configuration | 384

Verification | 406

Configuring Link Protection for OSPF | 421

Configuring Node-Link Protection for OSPF | 422

Configuring Node to Link Protection Fallback for OSPF | 423

Excluding an OSPF Interface as a Backup for a Protected Interface | 424

Configuring Backup SPF Options for Protected OSPF Interfaces | 425

Configuring RSVP Label-Switched Paths as Backup Paths for OSPF | 427

Remote LFA over LDP Tunnels in OSPF Networks Overview | 428

Configuring Remote LFA Backup over LDP Tunnels in an OSPF Network | 429

Example: Configuring Remote LFA Over LDP Tunnels in OSPF Networks | 431

Requirements | 432

Overview | 432

Configuration | 433

Verification | 443

Configure OSPF Support for Traffic Engineering

Configuring OSPF Support for Traffic Engineering | 453

OSPF Support for Traffic Engineering | 453

Example: Enabling OSPF Traffic Engineering Support | 455

Requirements | 455

Overview | 456

Configuration | 456

Verification | 462

Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface | 464

Requirements | 464

Overview | 464

Configuration | 464

Verification | 466

OSPF Passive Traffic Engineering Mode | 467

Example: Configuring OSPF Passive Traffic Engineering Mode | 467

Requirements | 467

Overview | 468

Configuration | 468

Verification | 470

Advertising Label-Switched Paths into OSPFv2 | 471

Example: Advertising Label-Switched Paths into OSPFv2 | 471

Requirements | 471

Overview | 472

Configuration | 474

Verification | 490

Static Adjacency Segment Identifier for OSPF | 491

Overview of Segment Routing Statistics | 494

How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering | 496

Understanding OSPF Flexible Algorithm for Segment Routing | 496

Example: OSPF Flexible Algorithm | 505

Overview | 506

Requirements | 507

Configuration | 507

- Verification | 528

- | 539

- | 539

- | 539

Configuring Application-Specific Link Attribute on an OSPF Interface | 540

How to Enable Link Delay Measurement and Advertising in OSPF | 546

- Understanding Link Delay Measurement and Advertising in OSPF | 546

- Configuring OSPF Link Delay and Delay Normalization on an OSPF Interface | 548

How to Configure Microloop Avoidance in OSPFv2 Segment Routing Networks | 554

- Understanding OSPF Microloop Avoidance | 554

- Configuring Segment Routing Microloop Avoidance in OSPFv2 Networks | 557

- Overview | 557

- Requirements | 557

- Topology | 557

- Configuration | 558

- Verification | 577

13

Configure OSPF Database Protection

Configuring OSPF Database Protection | 587

- OSPF Database Protection Overview | 587

- Configuring OSPF Database Protection | 588

14

Configure OSPF Routing Policy

Configuring OSPF Routing Policy | 591

- Understanding Routing Policies | 591

- Understanding OSPF Routing Policy | 595

- Understanding Backup Selection Policy for OSPF Protocol | 597

- Configuring Backup Selection Policy for the OSPF Protocol | 599

- Topology-Independent Loop-Free Alternate with Segment Routing for OSPF | 607

- Topology-Independent Loop-Free Alternate with Segment Routing for OSPF Overview | 607

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for OSPF | 609

Example: Configuring Backup Selection Policy for the OSPF or OSPF3 Protocol | 610

Requirements | 610

Overview | 611

Configuration | 612

Verification | 637

Example: Injecting OSPF Routes into the BGP Routing Table | 644

Requirements | 644

Overview | 644

Configuration | 645

Verification | 648

Troubleshooting | 649

Example: Redistributing Static Routes into OSPF | 649

Requirements | 650

Overview | 650

Configuration | 651

Verification | 653

Example: Configuring an OSPF Import Policy | 654

Requirements | 654

Overview | 654

Configuration | 655

Verification | 659

Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF | 660

Requirements | 661

Overview | 661

Configuration | 662

Verification | 666

Import and Export Policies for Network Summaries Overview | 667

Example: Configuring an OSPF Export Policy for Network Summaries | 667

Requirements | 667

Overview | 668

Configuration | 670

Verification | 679

Example: Configuring an OSPF Import Policy for Network Summaries | 680

Requirements | 680

Overview | 680

Configuration | 682

Verification | 692

Example: Redistributing OSPF Routes into IS-IS | 693

Requirements | 693

Overview | 693

Configuration | 694

Verification | 702

15

Configure OSPFv2 Sham Links

Configuring OSPFv2 Sham Links | 707

OSPFv2 Sham Links Overview | 707

Example: Configuring OSPFv2 Sham Links | 709

Requirements | 709

Overview | 709

Configuration | 710

Verification | 718

16

Configure OSPF on Logical Systems

Configuring OSPF on Logical Systems | 723

OSPF Support for Logical Systems | 723

Example: Configuring OSPF on Logical Systems Within the Same Router | 724

Requirements | 724

Overview | 724

Configuration | 726

Verification | 731

17

Troubleshooting Network Issues

Troubleshooting Network Issues | 737

Working with Problems on Your Network | 737

- Isolating a Broken Network Connection | 738
- Identifying the Symptoms of a Broken Network Connection | 740
- Isolating the Causes of a Network Problem | 742
- Taking Appropriate Action for Resolving the Network Problem | 743
- Evaluating the Solution to Check Whether the Network Problem Is Resolved | 745
- Checklist for Tracking Error Conditions | 747
- Configure Routing Protocol Process Tracing | 749
- Configure Routing Protocol Tracing for a Specific Routing Protocol | 752
- Monitor Trace File Messages Written in Near-Real Time | 755
- Stop Trace File Monitoring | 756

18

Verifying and Monitoring OSPF

Verifying and Monitoring OSPF Configuration | 759

Verifying an OSPF Configuration | 759

- Verifying OSPF-Enabled Interfaces | 759
- Verifying OSPF Neighbors | 761
- Verifying the Number of OSPF Routes | 762
- Verifying Reachability of All Hosts in an OSPF Network | 764

Tracing OSPF Protocol Traffic | 765

Example: Tracing OSPF Protocol Traffic | 767

- Requirements | 767
- Overview | 767
- Configuration | 769
- Verification | 774

19

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 777

About This Guide

Use this guide to configure, monitor, and troubleshoot the OSPF routing protocol on your Juniper Networks devices.

[Junos OS Routing Protocols Library for Routing Devices](#)

1

CHAPTER

OSPF Overview

IN THIS CHAPTER

- [Introduction to OSPF | 2](#)
-

Introduction to OSPF

IN THIS SECTION

- [OSPF Overview | 2](#)
- [OSPF Packets Overview | 7](#)
- [Understanding OSPF External Metrics | 10](#)
- [Supported OSPF and OSPFv3 Standards | 11](#)

OSPF Overview

IN THIS SECTION

- [OSPF Default Route Preference Values | 3](#)
- [OSPF Routing Algorithm | 4](#)
- [OSPF Three-Way Handshake | 5](#)
- [OSPF Version 3 | 5](#)

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.



NOTE: On SRX Series Firewalls, when only one link-protection is configured under the OSPF interface, the device does not install an alternative route in the forwarding table. When the per-packet load-balancing is enabled as a workaround, the device does not observe both the OSPF metric and sending the traffic through both the interfaces.

An OSPF AS can consist of a single area, or it can be subdivided into multiple areas. In a single-area OSPF network topology, each router maintains a database that describes the topology of the AS. Link-state information for each router is flooded throughout the AS. In a multiarea OSPF topology, each router maintains a database that describes the topology of its area, and link-state information for each router is flooded throughout that area. All routers maintain summarized topologies of other areas within an AS. Within each area, OSPF routers have identical topological databases. When the AS or area topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPFv2 protocol exchanges can be authenticated. OSPFv3 relies on IPsec to provide this functionality. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. A single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF 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.



NOTE: By default, Junos OS is compatible with RFC 1583, *OSPF Version 2*. In Junos OS Release 8.5 and later, you can disable compatibility with RFC 1583 by including the `no-rfc-1583` statement. For more information, see "[Example: Disabling OSPFv2 Compatibility with RFC 1583](#)" on page 258.

This topic describes the following information:

OSPF Default Route Preference Values

The Junos OS routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference value is from 0 through 4,294,967,295 (232 - 1), with a lower value indicating a more preferred route. [Table 1 on page 4](#) lists the default preference values for OSPF.

Table 1: Default Route Preference Values for OSPF

How Route Is Learned	Default Preference	Statement to Modify Default Preference
OSPF internal route	10	OSPF preference
OSPF AS external routes	150	OSPF external-preference

OSPF Routing Algorithm

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated router form adjacencies with other routing devices.) Adjacencies determine the distribution of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.

The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area


(*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

OSPF Three-Way Handshake

OSPF creates a topology map by flooding LSAs across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in [Figure 1 on page 5](#).

Figure 1: OSPF Three-Way Handshake



In [Figure 1 on page 5](#), Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

OSPF Version 3

OSPFv3 is a modified version of OSPF that supports IP version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.

- The protocol runs per link rather than per subnet.
- Router and network link-state advertisements (LSAs) do not carry prefix information.
- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.
- Flooding scopes are as follows:
 - Link-local
 - Area
 - AS
- Link-local addresses are used for all neighbor exchanges except virtual links.
- Authentication is removed. The IPv6 authentication header relies on the IP layer.
- The packet format has changed as follows:
 - Version number 2 is now version number 3.
 - The **db** option field has been expanded to 24 bits.
 - Authentication information has been removed.
 - Hello messages do not have address information.
 - Two new option bits are included: **R** and **V6**.
- Type 3 summary LSAs have been renamed *inter-area-prefix-LSAs*.
- Type 4 summary LSAs have been renamed *inter-area-router-LSAs*.

SEE ALSO

[Understanding OSPF Areas and Backbone Areas | 60](#)

[Example: Disabling OSPFv2 Compatibility with RFC 1583 | 258](#)

OSPF Packets Overview

IN THIS SECTION

- [OSPF Packet Header | 7](#)
- [Hello Packets | 8](#)
- [Database Description Packets | 8](#)
- [Link-State Request Packets | 8](#)
- [Link-State Update Packets | 9](#)
- [Link-State Acknowledgment Packets | 9](#)
- [Link-State Advertisement Packet Types | 9](#)

There are several types of link-state advertisement (LSA) packets.

This topic describes the following information:

OSPF Packet Header

All OSPFv2 packets have a common 24-byte header, and OSPFv3 packets have a common 16-byte header, that contains all information necessary to determine whether OSPF should accept the packet. The header consists of the following fields:

- Version number—The current OSPF version number. This can be either **2** or **3**.
- Type—Type of OSPF packet.
- Packet length—Length of the packet, in bytes, including the header.
- Router ID—IP address of the router from which the packet originated.
- Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0.
- Checksum—Fletcher checksum.
- Authentication—(OSPFv2 only) Authentication scheme and authentication information.
- Instance ID—(OSPFv3 only) Identifier used when there are multiple OSPFv3 realms configured on a link.

Hello Packets

Routers periodically send hello packets on all interfaces, including virtual links, to establish and maintain neighbor relationships. Hello packets are multicast on physical networks that have a multicast or broadcast capability, which enables dynamic discovery of neighboring routers. (On nonbroadcast networks, dynamic neighbor discovery is not possible, so you must configure all neighbors statically as described in ["Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network" on page 33.](#))

Hello packets consist of the OSPF header plus the following fields:

- Network mask—(OSPFv2 only) Network mask associated with the interface.
- Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval.
- Options—Optional capabilities of the router.
- Router priority—The router's priority to become the designated router.
- Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be down. All routers on a shared network must use the same router dead interval.
- Designated router—IP address of the designated router.
- Backup designated router—IP address of the backup designated router.
- Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

Database Description Packets

When initializing an adjacency, OSPF exchanges database description packets, which describe the contents of the topological database. These packets consist of the OSPF header, packet sequence number, and the link-state advertisement's header.

Link-State Request Packets

When a router detects that portions of its topological database are out of date, it sends a link-state request packet to a neighbor requesting a precise instance of the database. These packets consist of the OSPF header plus fields that uniquely identify the database information that the router is seeking.

Link-State Update Packets

Link-state update packets carry one or more link-state advertisements one hop farther from their origin. The router multicasts (floods) these packets on physical networks that support multicast or broadcast mode. The router acknowledges all link-state update packets and, if retransmission is necessary, sends the retransmitted advertisements unicast.

Link-state update packets consist of the OSPF header plus the following fields:

- Number of advertisements—Number of link-state advertisements included in this packet.
- Link-state advertisements—The link-state advertisements themselves.

Link-State Acknowledgment Packets

The router sends link-state acknowledgment packets in response to link-state update packets to verify that the update packets have been received successfully. A single acknowledgment packet can include responses to multiple update packets.

Link-state acknowledgment packets consist of the OSPF header plus the link-state advertisement header.

Link-State Advertisement Packet Types

Link-state request, link-state update, and link-state acknowledgment packets are used to reliably flood link-state advertisement packets. OSPF sends the following types of link-state advertisements:

- Router link advertisements—Are sent by all routers to describe the state and cost of the router's links to the area. These link-state advertisements are flooded throughout a single area only.
- Network link advertisements—Are sent by designated routers to describe all the routers attached to the network. These link-state advertisements are flooded throughout a single area only.
- Summary link advertisements—Are sent by area border routers to describe the routes that they know about in other areas. There are two types of summary link advertisements: those used when the destination is an IP network, and those used when the destination is an AS boundary router. Summary link advertisements describe interarea routes, that is, routes to destinations outside the area but within the AS. These link-state advertisements are flooded throughout the advertisement's associated areas.
- AS external link advertisement—Are sent by AS boundary routers to describe external routes that they know about. These link-state advertisements are flooded throughout the AS (except for stub areas).

Each link-state advertisement type describes a portion of the OSPF routing domain. All link-state advertisements are flooded throughout the AS.

Each link-state advertisement packet begins with a common 20-byte header.

SEE ALSO

[Understanding OSPF Areas | 50](#)

[Understanding OSPF Configurations | 14](#)

[OSPF Designated Router Overview | 53](#)

[Understanding OSPFv2 Authentication | 265](#)

[OSPF Timers Overview | 325](#)

Understanding OSPF External Metrics

When OSPF exports route information from external autonomous systems (ASs), it includes a cost, or *external metric*, in the route. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route.

- Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router.
- Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router.

By default, OSPF uses the Type 2 external metric.

Both Type 1 and Type 2 external metrics can be present in the AS at the same time. In that event, Type 1 external metrics always takes the precedence.

Type 1 external paths are always preferred over Type 2 external paths. When all paths are Type 2 external paths, the paths with the smallest advertised Type 2 metric are always preferred.

SEE ALSO

[Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth | 231](#)

Supported OSPF and OSPFv3 Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for OSPF and OSPF version 3 (OSPFv3).

- RFC 1583, *OSPF Version 2*
- RFC 1765, *OSPF Database Overflow*
- RFC 1793, *Extending OSPF to Support Demand Circuits*
- RFC 1850, *OSPF Version 2 Management Information Base*
- RFC 2154, *OSPF with Digital Signatures*
- RFC 2328, *OSPF Version 2*
- RFC 2370, *The OSPF Opaque LSA Option*

Support is provided by the `update-threshold` configuration statement at the `[edit protocols rsvp interface interface-name]` hierarchy level.

- RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
- RFC 3623, *Graceful OSPF Restart*
- RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
- RFC 4136, *OSPF Refresh and Flooding Reduction in Stable Topologies*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

Only interface switching is supported.

- RFC 4552, *Authentication/Confidentiality for OSPFv3*
- RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*
- RFC 4812, *OSPF Restart Signaling*
- RFC 4813, *OSPF Link-Local Signaling*
- RFC 4915, *Multi-Topology (MT) Routing in OSPF*

- RFC 5185, *OSPF Multi-Area Adjacency*
- RFC 5187, *OSPFv3 Graceful Restart*
- RFC 5250, *The OSPF Opaque LSA Option*



NOTE: RFC 4750, mentioned in this RFC as a "should" requirement is not supported. However, RFC 1850, the predecessor to RFC 4750 is supported.

- RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
- RFC 5340, *OSPF for IPv6* (RFC 2740 is obsoleted by RFC 5340)
- RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
- RFC 5838, *Support of Address Families in OSPFv3*
- Internet draft draft-ietf-ospf-af-alt-10.txt, *Support of address families in OSPFv3*
- Internet draft draft-katz-ward-bfd-02.txt, *Bidirectional Forwarding Detection*
Transmission of echo packets is not supported.
- RFC 6549, *OSPFv2 Multi-Instance Extensions*
- RFC 8665, *OSPF Extensions for Segment Routing*
- Internet draft draft-ietf-lsr-flex-algo-07.txt, *IGP Flexible Algorithm*

The following RFCs do not define standards, but provide information about OSPF and related technologies. The IETF classifies them as "Informational."

- RFC 3137, *OSPF Stub Router Advertisement*
- RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
- RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
- RFC 8920, *OSPF Application-Specific Link Attributes*
- RFC 8920, *OSPFv2 Prefix/Link Attribute Advertisement*

SEE ALSO

[Supported IPv6 Standards](#)

[Accessing Standards Documents on the Internet](#)

2

CHAPTER

Understand OSPF Configurations

IN THIS CHAPTER

- [Understanding OSPF Configurations | 14](#)
-

Understanding OSPF Configurations

To activate OSPF on a network, you must enable the protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF link-state advertisements (LSAs) are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

To complete the minimum device configuration for a node in an OSPF network involves:

1. Configuring the device interfaces.
See the [Junos OS Network Interfaces Library for Routing Devices](#) or the [Junos OS Interfaces Configuration Guide for Security Devices](#).
2. Configuring the router identifiers for the devices in your OSPF network
3. Creating the backbone area (area 0) for your OSPF network and adding the appropriate interfaces to the area



NOTE: Once you complete this step, OSPF begins sending LSAs. No additional configuration is required to enable OSPF traffic on the network.

You can further define your OSPF network depending on your network requirements. Some optional configurations involve:

- Adding additional areas to your network and configure area border routers (ABRs)
- Reducing the amount of memory that the nodes use to maintain the topology database by configuring stub and not-so-stubby areas
- Ensuring that only trusted routing devices participate in the autonomous systems' routing by enabling authentication
- Controlling the flow of traffic across the network by configuring path metrics and route selection

When describing how to configure OSPF, the following terms are used as follows:

- OSPF refers to both OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3)
- OSPFv2 refers to OSPF version 2
- OSPFv3 refers to OSPF version 3

3

CHAPTER

Configure OSPF Interfaces

IN THIS CHAPTER

- [Configuring OSPF Interfaces | 16](#)
-

Configuring OSPF Interfaces

IN THIS SECTION

- [About OSPF Interfaces | 16](#)
- [Example: Configuring an Interface on a Broadcast or Point-to-Point Network | 18](#)
- [Example: Configuring OSPF Demand Circuits | 23](#)
- [Example: Configuring a Passive OSPF Interface | 27](#)
- [Example: Configuring OSPFv2 Peer interfaces | 30](#)
- [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network | 33](#)
- [Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network | 38](#)
- [Understanding Multiple Address Families for OSPFv3 | 41](#)
- [Example: Configuring Multiple Address Families for OSPFv3 | 42](#)

About OSPF Interfaces

To activate OSPF on a network, you must enable the OSPF protocol on one or more interfaces on each device within the network on which traffic is to travel. How you configure the interface depends on whether the interface is connected to a broadcast or point-to-point network, a point-to-multipoint network, a nonbroadcast multiaccess (NBMA) network, or across a demand circuit.

- A broadcast interface behaves as if the routing device is connected to a LAN.
- A point-to-point interface provides a connection between a single source and a single destination (there is only one OSPF adjacency).
- A point-to-multipoint interface provides a connection between a single source and multiple destinations.
- An NBMA interface behaves in a similar fashion to a point-to-multipoint interface, but you might configure an NBMA interface to interoperate with other equipment.
- A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You can also configure an OSPF interface to be passive, to operate in passive traffic engineering mode, or to be a peer interface.

- A passive interface advertises its address, but does not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- An interface operating in OSPF passive traffic engineering mode floods link address information within the autonomous system (AS) and makes it available for traffic engineering calculations.
- A peer interface can be configured for OSPFv2 routing devices. A peer interface is required for Generalized MPLS (GMPLS) to transport traffic engineering information through a link separate from the control channel. You establish this separate link by configuring a peer interface. The peer interface name must match the Link Management Protocol (LMP) peer name. A peer interface is optional for a hierarchy of RSVP label-switched paths (LSPs). After you configure the forwarding adjacency, you can configure OSPFv2 to advertise the traffic engineering properties of a forwarding adjacency to a specific peer.

Point-to-point interfaces differ from multipoint in that only one OSPF adjacency is possible. (A LAN, for instance, can have multiple addresses and can run OSPF on each subnet simultaneously.) As such, when you configure a numbered point-to-point interface to OSPF by name, multiple OSPF interfaces are created. One, which is unnumbered, is the interface on which the protocol is run. An additional OSPF interface is created for each address configured on the interface, if any, which is automatically marked as passive.

For OSPFv3, one OSPF-specific interface must be created per interface name configured under OSPFv3. OSPFv3 does not allow interfaces to be configured by IP address.

Enabling OSPF on an interface (by including the `interface` statement), disabling it (by including the `disable` statement), and not actually having OSPF run on an interface (by including the `passive` statement) are mutually exclusive states.



NOTE: When you configure OSPFv2 on an interface, you must also include the `family inet` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. When you configure OSPFv3 on an interface, you must also include the `family inet6` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support address families other than unicast IPv6.

SEE ALSO

| [Example: Configuring OSPF Passive Traffic Engineering Mode](#) | 467

Example: Configuring an Interface on a Broadcast or Point-to-Point Network

IN THIS SECTION

- Requirements | 18
- Overview | 18
- Configuration | 19
- Verification | 22

This example shows how to configure an OSPF interface on a broadcast or point-to-point network.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- Topology | 19

If the interface on which you are configuring OSPF supports broadcast mode (such as a LAN), or if the interface supports point-to-point mode (such as a PPP interface or a point-to-point logical interface on Frame Relay), you specify the interface by including the IP address or the interface name for OSPFv2, or only the interface name for OSPFv3.

In Junos OS Release 9.3 and later, an OSPF point-to-point interface can be an Ethernet interface without a subnet. If you configure an interface on a broadcast network, designated router and backup designated router election is performed.



NOTE: Using both the interface name and the IP address of the same interface produces an invalid configuration.

In this example, you configure interface **ge-0/2/0** as an OSPFv2 interface in OSPF area 0.0.0.1.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 19](#)
- [Procedure | 20](#)
- [Results | 21](#)

CLI Quick Configuration

To quickly configure an OSPF interface on a broadcast or point-to-point network and to allow the inbound OSPF into the interfaces that are active, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
set protocols ospf area 0.0.0.1 interface ge-0/2/0
set security zones security-zone Trust host-inbound-traffic protocols all
set security zones security-zone Trust host-inbound-traffic system-services all
set groups global security policies default-policy permit-all
set security zones security-zone Trust interfaces ge-0/2/0
```



NOTE: You can also enable specific protocols or services for a security zone instead of all protocols or services. For instance, OSPF protocol or SSH service in this example:

```
set security zones security-zone Trust host-inbound-traffic protocols ospf
set security zones security-zone Trust host-inbound-traffic system-services ssh
```

Procedure

Step-by-Step Procedure

To configure an OSPF interface on a broadcast or point-to-point network:

1. Configure the interface.



NOTE: For an OSPFv3 interface, specify an IPv6 address.

```
[edit]
user@host# set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
```

2. Create an OSPF area.



NOTE: For an OSPFv3 interface, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
```

5. Configure the security zone to allow the inbound OSPF traffic into the interfaces that are active For more information about security zone, see *Security Zones*.

```
[edit]
user@host# set security zones security-zone Trust host-inbound-traffic protocols all
user@host# set security zones security-zone Trust host-inbound-traffic system-services all
user@host# set groups global security policies default-policy permit-all
user@host# set security zones security-zone Trust interfaces ge-0/2/0
user@host# commit
```



NOTE: You can also enable specific protocols or services for a security zone instead of all protocols or services. For instance, OSPF protocol or SSH service in this example:

```
set security zones security-zone Trust host-inbound-traffic protocols ospf
set security zones security-zone Trust host-inbound-traffic system-services ssh
```

Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}
```

```
}  
}
```

```
user@host# show protocols ospf  
area 0.0.0.1 {  
    interface ge-0/2/0.0;  
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

Verification

IN THIS SECTION

- [Verifying the OSPF Interface | 22](#)

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose

Verify the interface configuration. Depending on your deployment, the Type field might display LAN or P2P.

Action

From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

SEE ALSO

| No Link Title

Example: Configuring OSPF Demand Circuits

IN THIS SECTION

- Requirements | 23
- Overview | 23
- Configuration | 24
- Verification | 26

This example shows how to configure an OSPF demand circuit interface.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).



NOTE: If you are using OSPF demand circuits over an ISDN link, you must configure an ISDN interface and enable dial-on-demand routing.

- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- Topology | 24

OSPF sends periodic hello packets to establish and maintain neighbor adjacencies and uses link-state advertisements (LSAs) to make routing calculations and decisions. OSPF support for demand circuits is defined in RFC 1793, *Extending OSPF to Support Demand Circuits*, and suppresses the periodic hello packets and LSAs. A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You configure demand circuits on an OSPF interface. When the interface becomes a demand circuit, all hello packets and LSAs are suppressed as soon as OSPF synchronization is achieved. LSAs have a DoNotAge bit that stops the LSA from aging and prevents periodic updates from being sent. Hello packets and LSAs are sent and received on a demand-circuit interface only when there is a change in the network topology. This reduces the amount of traffic through the OSPF interface.

Consider the following when configuring OSPF demand circuits:

- Periodic hellos are only suppressed on point-to-point and point-to-multipoint interfaces. If you configure demand circuits on an OSPF broadcast network or on an OSPF nonbroadcast multiaccess (NBMA) network, periodic hello packets are still sent.
- Demand circuit support on an OSPF point-to-multipoint interface resembles that for point-to-point interfaces. If you configure a point-to-multipoint interface as a demand circuit, the device negotiates hello suppression separately on each interface that is part of the point-to-multipoint network.

This example assumes that you have a point-to-point connection between two devices using SONET/SDH interfaces. A demand-circuit interface automatically negotiates the demand-circuit connection with its OSPF neighbor. If the neighbor does not support demand circuits, then no demand circuit connection is established.

In this example, you configure OSPF interface **so-0/1/0** in OSPF area 0.0.0.1 as a demand circuit.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 25](#)
- [Procedure | 25](#)
- [Results | 26](#)

CLI Quick Configuration

To quickly configure an OSPF demand circuit interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-0/1/0 demand-circuit
```

Procedure

Step-by-Step Procedure

To configure an OSPF demand circuit interface on one neighboring interface:

1. Create an OSPF area.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit ]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the neighboring interface as a demand circuit.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-0/1/0 demand-circuit
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```



NOTE: Repeat this entire configuration on the other neighboring interface.

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
ospf {
  area 0.0.0.1 {
    interface so-0/1/0.0 {
      demand-circuit;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

IN THIS SECTION

- [Verifying the Status of Neighboring Interfaces | 26](#)

Confirm that the configuration is working properly.

Verifying the Status of Neighboring Interfaces

Purpose

Verify information about the neighboring interface. When the neighbor is configured for demand circuits, a DC flag displays.

Action

From operational mode, enter the **show ospf neighbor detail** command for OSPFv2, and enter the **show ospf3 neighbor detail** command for OSPFv3.

Example: Configuring a Passive OSPF Interface

IN THIS SECTION

- Requirements | 27
- Overview | 27
- Configuration | 28
- Verification | 29

This example shows how to configure a passive OSPF interface. A passive OSPF interface advertises its address but does not run the OSPF protocol.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

By default, OSPF must be configured on an interface for direct interface addresses to be advertised as interior routes. To advertise the direct interface addresses without actually running OSPF on that interface (adjacencies are not formed and hello packets are not generated), you configure that interface as a passive interface.

Enabling OSPF on an interface (by including the `interface` statement), disabling it (by including the `disable` statement), and not actually having OSPF run on an interface (by including the `passive` statement) are mutually exclusive states.



NOTE: If you do not want to see notifications for state changes in a passive OSPF interface, you can disable the OSPF traps for the interface by including the `no-interface-state-traps` statement. The `no-interface-state-traps` statement is supported only for OSPFv2.

In this example, you configure interface **ge-0/2/0** as a passive OSPF interface in area 0.0.0.1 by including the `passive` statement.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 28](#)
- [Procedure | 28](#)
- [Results | 29](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]  
set protocols ospf area 0.0.0.1 interface ge-0/2/0 passive
```

Procedure

Step-by-Step Procedure

To configure a passive OSPF interface:

1. Create an OSPF area.



NOTE: For an OSPFv3 interface, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the passive interface.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0 passive
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.1 {
    interface ge-0/2/0.0 {
      passive;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Status of OSPF Interfaces](#) | 30

Confirm that the configuration is working properly.

Verifying the Status of OSPF Interfaces

Purpose

Verify the status of the OSPF interface. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Example: Configuring OSPFv2 Peer interfaces

IN THIS SECTION

- [Requirements | 30](#)
- [Overview | 31](#)
- [Configuration | 31](#)
- [Verification | 32](#)

This example shows how to configure an OSPFv2 peer interface.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.

- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).
- Configure Generalized MPLS per your network requirements. .

Overview

You can configure an OSPFv2 peer interface for many reasons, including when you configure Generalized MPLS (GMPLS). This example configures a peer interface for GMPLS. GMPLS requires traffic engineering information to be transported through a link separate from the control channel. You establish this separate link by configuring a peer interface. The OSPFv2 peer interface name must match the Link Management Protocol (LMP) peer name. You configure GMPLS and the LMP settings separately from OSPF.

This example assumes that GMPLS and the LMP peer named **oxc1** are already configured, and you need to configure the OSPFv2 peer interface in area 0.0.0.0.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 31](#)
- [Procedure | 32](#)
- [Results | 32](#)

CLI Quick Configuration

To quickly configure an OSPFv2 peer interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]  
set protocols ospf area 0.0.0.0 peer-interface oxc1
```

Procedure

Step-by-Step Procedure

To configure a peer OSPFv2 interface used by the LMP:

1. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the peer interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set peer-interface oxc1
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.0 {
    peer-interface oxc1;
  }
```

Verification

IN THIS SECTION

- [Verifying the Configured OSPFv2 Peer | 33](#)

Confirm that the configuration is working properly.

Verifying the Configured OSPFv2 Peer

Purpose

Verify the status of the OSPFv2 peer. When an OSPFv2 peer is configured for GMPLS, the Peer Name field displays the name of the LMP peer that you created for GMPLS, which is also the configured OSPFv2 peer.

Action

From operational mode, enter the **show link-management** command.

Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

IN THIS SECTION

- Requirements | 33
- Overview | 34
- Configuration | 35
- Verification | 37

This example shows how to configure an OSPFv2 interface on a nonbroadcast multiaccess (NBMA) network.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#).

- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

IN THIS SECTION

- [Topology | 35](#)

When you configure OSPFv2 on an NBMA network, you can use nonbroadcast mode rather than point-to-multipoint mode. Using this mode offers no advantages over point-to-multipoint mode, but it has more disadvantages than point-to-multipoint mode. Nevertheless, you might occasionally find it necessary to configure nonbroadcast mode to interoperate with other equipment. Because there is no autodiscovery mechanism, you must configure each neighbor.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routing devices must have a direct connection to both the designated and backup designated routers, or unpredictable results occur.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration. For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the interface name.

In this example, you configure the Asynchronous Transfer Mode (ATM) interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you and specify the following settings:

- **interface-type nbma**—Sets the interface to run in NBMA mode. You must explicitly configure the interface to run in NBMA mode.
- **neighbor *address* <eligible>**—Specifies the IP address of the neighboring device. OSPF routing devices normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the device cannot discover its neighbors dynamically, so you must configure all the neighbors statically. To configure multiple neighbors, include multiple `neighbor` statements. If you want the neighbor to be a designated router, include the **eligible** keyword.
- **poll-interval**—Specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before it establishes adjacency with a neighbor. Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The range is from 1 through 255 seconds. By default, the device sends hello packets out the interface every 120 seconds before it establishes adjacency with a neighbor.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the `poll-interval` statement to the time specified in the `hello-interval` statement.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 35](#)
- [Procedure | 35](#)
- [Results | 36](#)

CLI Quick Configuration

To quickly configure an OSPFv2 interface on an NBMA network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 interface-type nbma
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 neighbor 192.0.2.2 eligible
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 poll-interval 130
```

Procedure

Step-by-Step Procedure

To configure an OSPFv2 interface on an NBMA network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.

In this example, include the **eligible** keyword to allow the neighbor to be a designated router.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 interface-type nbma neighbor 192.0.2.2 eligible
```

4. Configure the poll interval.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 poll-interval 130
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/32;
    }
  }
}
```

```
}  
}
```

```
user@host# show protocols ospf  
area 0.0.0.1 {  
  interface at-0/1/0.0 {  
    interface-type nbma;  
    neighbor 192.0.2.2 eligible;  
    poll-interval 130;  
  }  
}
```

Verification

IN THIS SECTION

- [Verifying the OSPF Interface | 37](#)

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose

Verify the interface configuration. Confirm that the Type field displays NBMA.

Action

From operational mode, enter the `show ospf interface detail` command.

SEE ALSO

| [OSPF Timers Overview | 325](#)

Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network

IN THIS SECTION

- Requirements | 38
- Overview | 38
- Configuration | 39
- Verification | 41

This example shows how to configure an OSPFv2 interface on a point-to-multipoint network.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

IN THIS SECTION

- Topology | 39

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint Asynchronous Transfer Mode (ATM) or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, you must configure each neighbor.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration.

In this example, you configure ATM interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify 192.0.2.1 as the neighbor's IP address.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 39](#)
- [Procedure | 39](#)
- [Results | 40](#)

CLI Quick Configuration

To quickly configure an OSPFv2 interface on a point-to-multipoint network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
set protocols ospf area 0.0.0.1 interface at-0/1/0 neighbor 192.0.2.1
```

Procedure

Step-by-Step Procedure

To configure an OSPFv2 interface on a point-to-multipoint network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area and specify the neighbor.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 neighbor 192.0.2.1
```

To configure multiple neighbors, include a `neighbor` statement for each neighbor.

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.2/32;
    }
  }
}
```

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface at-0/1/0.0 {
    neighbor 192.0.2.1;
  }
}
```

```
}
}
```

Verification

IN THIS SECTION

- [Verifying the OSPF Interface | 41](#)

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose

Verify the interface configuration. Confirm that the Type field displays P2MP.

Action

From operational mode, enter the `show ospf interface detail` command.

Understanding Multiple Address Families for OSPFv3

By default, OSPFv3 supports only unicast IPv6 routes. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support multiple address families, including IPv4 unicast, IPv4 multicast, and IPv6 multicast. This multiple address family support allows OSPFv3 to support both IPv6 and IPv4 nodes. Junos OS maps each address family to a separate realm as defined in RFC 5838, *Support for Address Families in OSPFv3*. Each realm maintains a separate set of neighbors and link-state database.

When you configure multiple address families for OSPFv3, there is a new instance ID field that allows multiple OSPFv3 protocol instances per link. This allows a single link to belong to multiple areas.

You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.

These are the default import and export routing tables for each of the four address families:

- IPv6 unicast: **inet6.0**

- IPv6 multicast: **inet6.2**
- IPv4 unicast: **inet.0**
- IPv4 multicast: **inet.2**

With the exception of virtual links, all configurations supported for the default IPv6 unicast family are supported for the address families that have to be configured as realms.

Example: Configuring Multiple Address Families for OSPFv3

IN THIS SECTION

- [Requirements | 42](#)
- [Overview | 43](#)
- [Configuration | 44](#)
- [Verification | 47](#)

This example shows how to configure multiple address families for OSPFv3.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62](#).
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

IN THIS SECTION

- [Topology | 44](#)

By default, OSPFv3 supports unicast IPv6 routes, but you can configure OSPFv3 to support multiple address families. To support an address family other than unicast IPv6, you configure a realm that allows OSPFv3 to advertise IPv4 unicast, IPv4 multicast, or IPv6 multicast routes. Junos OS then maps each address family that you configure to a separate realm with its own set of neighbors and link-state database.




NOTE: By default, LDP synchronization is only supported for OSPFv2. If you configure an IPv4 unicast or IPv4 multicast realm, you can also configure LDP synchronization. Since LDP synchronization is only supported for IPv4, this support is only available for OSPFv3 if you configure an IPv4 realm.

When configuring OSPFv3 to support multiple address families, consider the following:

- You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.
- OSPFv3 uses IPv6 link-local addresses as the source of hello packets and next hop calculations. As such, you must enable IPv6 on the link regardless of the additional realm you configure.

[Figure 2 on page 44](#) shows a connection between Routers R0 and R1. In this example, you configure interface `fe-0/1/0` on Router R0 in area 0 to advertise IPv4 unicast routes, in addition to the default unicast IPv6 routes in area 1, by including the `realm ipv4-unicast` statement. Depending on your network requirements, you can also advertise IPv4 multicast routes by including the `realm-ipv4-multicast` statement, and you can advertise IPv6 multicast routes by including the `realm-ipv6-multicast` statement.

Figure 2: IPv4 Unicast Realm



Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 44](#)
- [Procedure | 45](#)
- [Results | 46](#)

CLI Quick Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To quickly configure multiple address families for OSPFv3, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter commit from configuration mode.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 192.0.2.2/24
set interfaces fe-0/1/0 unit 0 family inet6
set protocols ospf3 area 0.0.0.0 interface fe-0/1/0
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

Procedure

Step-by-Step Procedure

To configure multiple address families for OSPFv3:

1. Configure the device interface participating in OSPFv3.

```
[edit]
user@host# set interfaces fe-0/1/0 unit 0 family inet address 192.0.2.2/24
user@host# set interfaces fe-0/1/0 unit 0 family inet6
```

2. Enter OSPFv3 configuration mode.

```
[edit ]
user@host# edit protocols ospf3
```

3. Add the interface you configured to the OSPFv3 area.

```
[edit protocols ospf3 ]
user@host# set area 0.0.0.0 interface fe-0/1/0
```

4. Configure an IPv4 unicast realm. This allows OSPFv3 to support both IPv4 unicast and IPv6 unicast routes.

```
[edit protocols ospf3 ]
user@host# set realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf3 ]  
user@host# commit
```



NOTE: Repeat this entire configuration on the neighboring device that is part of the realm.

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces  
fe-0/1/0 {  
  unit 0 {  
    family inet {  
      address 192.0.2.2/24;  
    }  
    family inet6;  
  }  
}
```

```
user@host# show protocols ospf3  
realm ipv4-unicast {  
  area 0.0.0.0 {  
    interface fe-0/1/0.0;  
  }  
}  
area 0.0.0.0 {  
  interface fe-0/1/0.0;  
}
```

Verification

IN THIS SECTION

- [Verifying the Link-State Database | 47](#)
- [Verifying the Status of OSPFv3 Interfaces with Multiple Address Families | 47](#)

Confirm that the configuration is working properly.

Verifying the Link-State Database

Purpose

Verify the status of the link-state database for the configured realm, or address family.

Action

From operational mode, enter the `show ospf3 database realm ipv4-unicast` command.

Verifying the Status of OSPFv3 Interfaces with Multiple Address Families

Purpose

Verify the status of the interface for the specified OSPFv3 realm, or address family.

Action

From operational mode, enter the `show ospf3 interface realm ipv4-unicast` command.

4

CHAPTER

Configure OSPF Areas

IN THIS CHAPTER

- [Configuring OSPF Areas | 49](#)
-

Configuring OSPF Areas

IN THIS SECTION

- Understanding OSPF Areas | 50
- OSPF Designated Router Overview | 53
- Example: Configuring an OSPF Router Identifier | 54
- Example: Controlling OSPF Designated Router Election | 57
- Understanding OSPF Areas and Backbone Areas | 60
- Example: Configuring a Single-Area OSPF Network | 62
- Example: Configuring a Multiarea OSPF Network | 65
- Understanding Multiarea Adjacency for OSPF | 71
- Example: Configuring Multiarea Adjacency for OSPF | 72
- Understanding Multiarea Adjacencies for OSPFv3 | 79
- Example: Configuring a Multiarea Adjacency for OSPFv3 | 79
- Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas | 89
- Example: Configuring OSPF Stub and Totally Stubby Areas | 91
- Example: Configuring OSPF Not-So-Stubby Areas | 97
- Understanding OSPFv3 Stub and Totally Stubby Areas | 105
- Example: Configuring OSPFv3 Stub and Totally Stubby Areas | 105
- Understanding OSPFv3 Not-So-Stubby Areas | 121
- Example: Configuring OSPFv3 Not-So-Stubby Areas | 121
- Understanding Not-So-Stubby Areas Filtering | 140
- Example: Configuring OSPFv3 Not-So-Stubby Areas with Filtering | 141
- Understanding OSPF Virtual Links for Noncontiguous Areas | 152
- Example: Configuring OSPF Virtual Links to Connect Noncontiguous Areas | 153
- Example: Configuring OSPFv3 Virtual Links | 159

Understanding OSPF Areas

IN THIS SECTION

- Areas | 50
- Area Border Routers | 51
- Backbone Areas | 51
- AS Boundary Routers | 51
- Backbone Router | 51
- Internal Router | 51
- Stub Areas | 52
- Not-So-Stubby Areas | 52
- Transit Areas | 52
- OSPF Area Types and Accepted LSAs | 52

In OSPF, a single autonomous system (AS) can be divided into smaller groups called *areas*. This reduces the number of link-state advertisements (LSAs) and other OSPF overhead traffic sent on the network, and it reduces the size of the topology database that each router must maintain. The routing devices that participate in OSPF routing perform one or more functions based on their location in the network.

This topic describes the following OSPF area types and routing device functions:

Areas

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Also, routing within the area is determined only by the area's topology, providing the area with some protection from bad routing data.

All routing devices within an area have identical topology databases.

Area Border Routers

Routing devices that belong to more than one area and connect one or more OSPF areas to the backbone area are called *area border routers* (ABRs). At least one interface is within the backbone while another interface is in another area. ABRs also maintain a separate topological database for each area to which they are connected.

Backbone Areas

An OSPF *backbone area* consists of all networks in area ID 0.0.0.0, their attached routing devices, and all ABRs. The backbone itself does not have any ABRs. The backbone distributes routing information between areas. The backbone is simply another area, so the terminology and rules of areas apply: a routing device that is directly connected to the backbone is an internal router on the backbone, and the backbone's topology is hidden from the other areas in the AS.

The routing devices that make up the backbone must be physically contiguous. If they are not, you must configure *virtual links* to create the appearance of backbone connectivity. You can create virtual links between any two ABRs that have an interface to a common nonbackbone area. OSPF treats two routing devices joined by a virtual link as if they were connected to an unnumbered point-to-point network.

AS Boundary Routers

Routing devices that exchange routing information with routing devices in non-OSPF networks are called *AS boundary routers*. They advertise externally learned routes throughout the OSPF AS. Depending on the location of the AS boundary router in the network, it can be an ABR, a backbone router, or an internal router (with the exception of stub areas). Internal routers within a stub area cannot be an AS boundary router because stub areas cannot contain any Type 5 LSAs.

Routing devices within the area where the AS boundary router resides know the path to that AS boundary router. Any routing device outside the area only knows the path to the nearest ABR that is in the same area where the AS boundary router resides.

Backbone Router

Backbone routers are routing devices that have one or more interfaces connected to the OSPF backbone area (area ID 0.0.0.0).

Internal Router

Routing devices that connect to only one OSPF area are called *internal routers*. All interfaces on internal routers are directly connected to networks within a single area.

Stub Areas

Stub areas are areas through which or into which AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological databases and therefore the amount of memory required on the internal routers in the stub area.

Routing devices within a stub area rely on the default routes originated by the area's ABR to reach external AS destinations. You must configure the `default-metric` option on the ABR before it advertises a default route. Once configured, the ABR advertises a default route in place of the external routes that are not being advertised within the stub area, so that routing devices in the stub area can reach destinations outside the area.

The following restrictions apply to stub areas: you cannot create a virtual link through a stub area, a stub area cannot contain an AS boundary router, the backbone cannot be a stub area, and you cannot configure an area as both a stub area and a not-so-stubby area.

Not-So-Stubby Areas

An OSPF stub area has no external routes in it, so you cannot redistribute from another protocol into a stub area. A *not-so-stubby area* (NSSA) allows external routes to be flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

The following restriction applies to NSSAs: you cannot configure an area as both a stub area and an NSSA.

Transit Areas

Transit areas are used to pass traffic from one adjacent area to the backbone (or to another area if the backbone is more than two hops away from an area). The traffic does not originate in, nor is it destined for, the transit area.

OSPF Area Types and Accepted LSAs

The following table gives details about OSPF area types and accepted LSAs:

OSPF Area Types and Accepted LSAs						
Area Types	LSA 1	LSA 2	LSA 3	LSA 4	LSA 5	LSA 7
Backbone Area	Yes	Yes	Yes	Yes	Yes	No
Non-Backbone Area	Yes	Yes	Yes	Yes	Yes	No
Stub Area	Yes	Yes	Yes	No	No	No
Totally Stubby Area	Yes	Yes	No	No	No	No
Not-So-Stubby Area	Yes	Yes	Yes	No	No	Yes

8200034

OSPF Designated Router Overview

Large LANs that have many routing devices and therefore many OSPF adjacencies can produce heavy control-packet traffic as link-state advertisements (LSAs) are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers on all multiaccess networks (broadcast and nonbroadcast multiaccess [NBMA] networks types). Rather than broadcasting LSAs to all their OSPF neighbors, the routing devices send their LSAs to the designated router. Each multiaccess network has a designated router, which performs two main functions:

- Originate network link advertisements on behalf of the network.
- Establish adjacencies with all routing devices on the network, thus participating in the synchronizing of the link-state databases.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the routing device with the highest router identifier (defined by the **router-id** configuration value, which is typically the IP address of the routing device, or the loopback address) is elected the designated router. The routing device with the second highest router identifier is elected the backup designated router. If the designated router fails or loses

connectivity, the backup designated router assumes its role and a new backup designated router election takes place between all the routers in the OSPF network.

OSPF uses the router identifier for two main purposes: to elect a designated router, unless you manually specify a priority value, and to identify the routing device from which a packet is originated. At designated router election, the router priorities are evaluated first, and the routing device with the highest priority is elected designated router. If router priorities tie, the routing device with the highest router identifier, which is typically the routing device's IP address, is chosen as the designated router. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

At least one routing device on each logical IP network or subnet must be eligible to be the designated router for OSPFv2. At least one routing device on each logical link must be eligible to be the designated router for OSPFv3.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router. A priority of 255 means the routing device is always the designated router.

Example: Configuring an OSPF Router Identifier

IN THIS SECTION

- Requirements | 54
- Overview | 55
- Configuration | 55
- Verification | 57

This example shows how to configure an OSPF router identifier.

Requirements

Before you begin:

- Identify the interfaces on the routing device that will participate in OSPF. You must enable OSPF on all interfaces within the network on which OSPF traffic is to travel.
- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#)

Overview

The router identifier is used by OSPF to identify the routing device from which a packet originated. Junos OS selects a router identifier according to the following set of rules:

1. By default, Junos OS selects the lowest configured physical IP address of an interface as the router identifier.
2. If a loopback interface is configured, the IP address of the loopback interface becomes the router identifier.
3. If multiple loopback interfaces are configured, the lowest loopback address becomes the router identifier.
4. If a router identifier is explicitly configured using the `router-id address` statement under the `[edit routing-options]` hierarchy level, the above three rules are ignored.



NOTE: 1. The router identifier behavior described here holds good even when configured under `[edit routing-instances routing-instance-name routing-options]` and `[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options]` hierarchy levels.

2. If the router identifier is modified in a network, the link-state advertisements (LSAs) advertised by the previous router identifier are retained in the OSPF database until the LSA retransmit interval has timed out. Hence, it is strongly recommended that you explicitly configure the router identifier under the `[edit routing-options]` hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

In this example, you configure the OSPF router identifier by setting its router ID value to the IP address of the device, which is 192.0.2.24.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 56](#)
- [Procedure | 56](#)
- [Results | 56](#)

CLI Quick Configuration

To quickly configure an OSPF router identifier, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set routing-options router-id 192.0.2.24
```

Procedure

Step-by-Step Procedure

To configure an OSPF router identifier:

1. Configure the OSPF router identifier by entering the [router-id] configuration value.

```
[edit]
user@host# set routing-options router-id 192.0.2.24
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options router-id` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options router-id
router-id 192.0.2.24;
```

Verification

After you configure the router ID and activate OSPF on the routing device, the router ID is referenced by multiple OSPF operational mode commands that you can use to monitor and troubleshoot the OSPF protocol. The router ID fields are clearly marked in the output.

Example: Controlling OSPF Designated Router Election

IN THIS SECTION

- Requirements | 57
- Overview | 57
- Configuration | 58
- Verification | 59

This example shows how to control OSPF designated router election.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

Overview

This example shows how to control OSPF designated router election. Within the example, you set the OSPF interface to **ge-0/0/1** and the device priority to 200. The higher the priority value, the greater likelihood the routing device will become the designated router.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 58](#)
- [Procedure | 58](#)
- [Results | 59](#)

CLI Quick Configuration

To quickly configure an OSPF designated router election, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

Procedure

Step-by-Step Procedure

To control OSPF designated router election:

1. Configure an OSPF interface and specify the device priority.



NOTE: To specify an OSPFv3 interface, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.3 {
  interface ge-0/0/1.0 {
    priority 200;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Designated Router Election | 59](#)

Confirm that the configuration is working properly.

Verifying the Designated Router Election

Purpose

Based on the priority you configured for a specific OSPF interface, you can confirm the address of the area's designated router. The DR ID, DR, or DR-ID field displays the address of the area's designated router. The BDR ID, BDR, or BDR-ID field displays the address of the backup designated router.

Action


From operational mode, enter the `show ospf interface` and the `show ospf neighbor` commands for OSPFv2, and enter the `show ospf3 interface` and the `show ospf3 neighbor` commands for OSPFv3.

Understanding OSPF Areas and Backbone Areas

OSPF networks in an autonomous system (AS) are administratively grouped into areas. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions similar to a network address. Within an area, the topology database contains only information about the area, link-state advertisements (LSAs) are flooded only to nodes within the area, and routes are computed only within the area. The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Subnetworks are divided into other areas, which are connected to form the whole of the main network. Routing devices that are wholly within an area are called internal routers. All interfaces on internal routers are directly connected to networks within the area.

The central area of an AS, called the backbone area, has a special function and is always assigned the area ID 0.0.0.0. (Within a simple, single-area network, this is also the ID of the area.) Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a routing device that has interfaces in more than one area. These connecting routing devices are called border area routers (ABRs). [Figure 3 on page 60](#) shows an OSPF topology of three areas connected by two ABRs.

Figure 3: Multiarea OSPF Topology



Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for


transmitting the traffic through the appropriate ABR to the destination area. The ABRs summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate ABR. For example, in the OSPF areas shown in [Figure 3 on page 60](#), packets sent from Router A to Router C are automatically routed through ABR B.

Junos OS supports active backbone detection. Active backbone detection is implemented to verify that ABRs are connected to the backbone. If the connection to the backbone area is lost, then the routing device's default metric is not advertised, effectively rerouting traffic through another ABR with a valid connection to the backbone. Active backbone detection enables transit through an ABR with no active backbone connection. An ABR advertises to other routing devices that it is an ABR even if the connection to the backbone is down, so that the neighbors can consider it for interarea routes.

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate ABR and on to the remote host within the destination area.

In large networks with many areas, in which direct connectivity between all areas and the backbone area is physically difficult or impossible, you can configure virtual links to connect noncontiguous areas. Virtual links use a transit area that contains two or more ABRs to pass network traffic from one adjacent area to another. For example, [Figure 4 on page 61](#) shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 4: OSPF Topology with a Virtual Link



In the topology shown in [Figure 4 on page 61](#), a virtual link is established between area 0.0.0.3 and the backbone area through area 0.0.0.2. All outbound traffic destined for other areas is routed through area 0.0.0.2 to the backbone area and then to the appropriate ABR. All inbound traffic destined for area 0.0.0.3 is routed to the backbone area and then through area 0.0.0.2.

Example: Configuring a Single-Area OSPF Network

IN THIS SECTION

- [Requirements | 62](#)
- [Overview | 62](#)
- [Configuration | 63](#)
- [Verification | 64](#)

This example shows how to configure a single-area OSPF network.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

Overview

IN THIS SECTION


- [Topology | 63](#)

To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

In an autonomous system (AS), the backbone area is always assigned area ID 0.0.0.0 (within a simple, single-area network, this is also the ID of the area). Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by area border routers that have interfaces in more than one area. You must also create a backbone area if your network consists of multiple areas. In this example, you create the backbone area and add interfaces, such as **ge-0/0/0**, as needed to the OSPF area.

To use OSPF on the device, you must configure at least one OSPF area, such as the one shown in [Figure 5 on page 63](#).

Figure 5: Typical Single-Area OSPF Network Topology



Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 63](#)
- [Procedure | 64](#)
- [Results | 64](#)

CLI Quick Configuration

To quickly configure a single-area OSPF network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

Procedure

Step-by-Step Procedure

To configure a single-area OSPF network:

1. Configure the single-area OSPF network by specifying the area ID and associated interface.



NOTE: For a single-area OSPFv3 network, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Interfaces in the Area | 65](#)

Confirm that the configuration is working properly.

Verifying the Interfaces in the Area

Purpose

Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.

Action

From operational mode, enter the `show ospf interface` command for OSPFv2, and enter the `show ospf3 interface` command for OSPFv3.

Example: Configuring a Multiarea OSPF Network

IN THIS SECTION

- [Requirements | 65](#)
- [Overview | 66](#)
- [Configuration | 67](#)
- [Verification | 71](#)

This example shows how to configure a multiarea OSPF network. To reduce traffic and topology maintenance for the devices in an OSPF autonomous system (AS), you can group the OSPF-enabled routing devices into multiple areas.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62.](#)

Overview

IN THIS SECTION


- [Topology | 67](#)

To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

Each OSPF area consists of routing devices configured with the same area number. In [Figure 6 on page 67](#), Router B resides in the backbone area of the AS. The backbone area is always assigned area ID 0.0.0.0. (All area IDs must be unique within an AS.) All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. In this example, these area border routers are A, C, D, and E. You create an additional area (area 2) and assign it unique area ID 0.0.0.2, and then add interface **ge-0/0/0** to the OSPF area.

To reduce traffic and topology maintenance for the devices in an OSPF AS, you can group them into multiple areas as shown in [Figure 6 on page 67](#). In this example, you create the backbone area, create an additional area (area 2) and assign it unique area ID 0.0.0.2, and you configure Device B as the area border router, where interface **ge-0/0/0** participates in OSPF area 0 and interface **ge-0/0/2** participates in OSPF area 2.

Figure 6: Typical Multiarea OSPF Network Topology



Topology

Configuration

IN THIS SECTION

- Procedure | 67
- Results | 70

Procedure

CLI Quick Configuration

To quickly configure a multiarea OSPF network, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Device A

```
[edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

Device C

```
[edit]  
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

Device B

```
[edit]  
set protocols ospf area 0.0.0.0 interface ge-0/0/0  
set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

Device D

```
[edit]  
set protocols ospf area 0.0.0.2 interface ge-0/0/0  
set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

Device E

```
[edit]  
set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

Step-by-Step Procedure

To configure a multiarea OSPF network:

1. Configure the backbone area.



NOTE: For an OSPFv3 network, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

```
[edit]
user@C# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

```
[edit]
user@B# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. Configure an additional area for your OSPF network.



NOTE: For a multiarea OSPFv3 network, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.2 interface ge-0/0/0
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@E# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
}
```

```
user@C# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
}
```

```
user@B# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
}
area 0.0.0.2 {
    interface ge-0/0/2.0;
}
```

```
user@D# show protocols ospf
area 0.0.0.2 {
    interface ge-0/0/0.0;
    interface ge-0/0/2.0;
}
```

```
user@E# show protocols ospf
area 0.0.0.2 {
    interface ge-0/0/2.0;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Interfaces in the Area | 71](#)

Confirm that the configuration is working properly.

Verifying the Interfaces in the Area

Purpose

Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.

Action

From operational mode, enter the `show ospf interface` command for OSPFv2, and enter the `show ospf3 interface` command for OSPFv3.

Understanding Multiarea Adjacency for OSPF

By default, a single interface can belong to only one OSPF area. However, in some situations, you might want to configure an interface to belong to more than one area. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. For example, you can configure an interface to belong to multiple areas with a high-speed backbone link between two area border routers (ABRs) so you can create multiarea adjacencies that belong to different areas.

In Junos OS Release 9.2 and later, you can configure a *logical interface* to belong to more than one OSPFv2 area. Support for OSPFv3 was introduced in Junos OS Release 9.4. As defined in RFC 5185, *OSPF Multi-Area Adjacency*, the ABRs establish multiple adjacencies belonging to different areas over the same logical interface. Each multiarea adjacency is announced as a point-to-point unnumbered link in the configured area by the routers connected to the link. For each area, one of the logical interfaces is treated as primary, and the remaining interfaces that are configured for the area are designated as secondary.

Any logical interface not configured as a secondary interface for an area is treated as the primary interface for that area. A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.

Example: Configuring Multiarea Adjacency for OSPF

IN THIS SECTION

- Requirements | 72
- Overview | 72
- Configuration | 74
- Verification | 77

This example shows how to configure multiarea adjacency for OSPF.

Requirements

Before you begin, plan your multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview


By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:

- For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.
- Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.
- Since the purpose of a secondary interface is to advertise a topological path through an OSPF area, you cannot configure a secondary interface or a primary interface with one or more secondary

interfaces to be passive. Passive interfaces advertise their address, but do not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).

- Any logical interface not configured as a secondary interface for an area is treated as a primary interface for that area. A logical interface can be configured as the primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.
- You cannot configure the secondary statement with the interface all statement.
- You cannot configure a secondary interface by its IP address.

Figure 7: Multiarea Adjacency in OSPF



In this example, you configure an interface to be in two areas, creating a multiarea adjacency with a link between two ABRs: ABR R1 and ABR R2. On each ABR, area 0.0.0.1 contains the primary interface and is the primary link between the ABRs, and area 0.0.0.2 contains the secondary logical interface, which you configure by including the secondary statement. You configure interface so-0/0/0 on ABR R1 and interface so-1/0/0 on ABR R2.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 74](#)
- [Procedure | 74](#)
- [Results | 76](#)

CLI Quick Configuration

To quickly configure a secondary logical interface for an OSPF area, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Configuration on ABR R1:

```
[edit]
set interfaces so-0/0/0 unit 0 family inet address 192.0.2.45/24
set routing-options router-id 10.255.0.1
set protocols ospf area 0.0.0.1 interface so-0/0/0
set protocols ospf area 0.0.0.2 interface so-0/0/0 secondary
```

Configuration on ABR R2:

```
[edit]
set interfaces so-1/0/0 unit 0 family inet address 192.0.2.37/24
set routing-options router-id 10.255.0.2
set protocols ospf area 0.0.0.1 interface so-1/0/0
set protocols ospf area 0.0.0.2 interface so-1/0/0 secondary
```

Procedure

Step-by-Step Procedure

To configure a secondary logical interface:

1. Configure the device interfaces.



NOTE: For OSPFv3, on each interface specify the **inet6** address family and include the IPv6 address.

```
[edit]
user@R1# set interfaces so-0/0/0 unit 0 family inet address 192.0.2.45/24
```

```
[edit]
user@R2# set interfaces so-1/0/0 unit 0 family inet address 192.0.2.37/24
```

2. Configure the router identifier.

```
[edit]
user@R1# set routing-options router-id 10.255.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.255.0.2
```

3. On each ABR, configure the primary interface for the OSPF area.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.1 interface so-0/0/0
```

```
[edit ]
user@R2# set protocols ospf area 0.0.0.1 interface so-1/0/0
```

4. On each ABR, configure the secondary interface for the OSPF area.

```
[edit ]
user@R1# set protocols ospf area 0.0.0.2 so-0/0/0 secondary
```

```
[edit ]
user@R2# set protocols ospf area 0.0.0.2 so-1/0/0 secondary
```

5. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show routing-options`, and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R1:

```
user@R1# show interfaces
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.45/24;
    }
  }
}
```

```
user@R1# show routing-options
router-id 10.255.0.1;
```

```
user@R1# show protocols ospf
area 0.0.0.1 {
  interface so-0/0/0.0;
```

```
}  
area 0.0.0.2 {  
    interface so-0/0/0.0 {  
        secondary;  
    }  
}
```

Configuration on ABR R2:

```
user@R2# show interfaces  
so-0/0/0 {  
    unit 0 {  
        family inet {  
            address 192.0.2.37/24;  
        }  
    }  
}
```

```
user@R2# show routing-options  
router-id 10.255.0.2;
```

```
user@R2# show protocols ospf  
area 0.0.0.1 {  
    interface so-1/0/0.0;  
}  
area 0.0.0.2 {  
    interface so-1/0/0.0 {  
        secondary;  
    }  
}
```

Verification

IN THIS SECTION

- [Verifying the Secondary Interface | 78](#)
- [Verifying the Interfaces in the Area | 78](#)

- Verifying Neighbor Adjacencies | 78

Confirm that the configuration is working properly.

Verifying the Secondary Interface

Purpose

Verify that the secondary interface appears for the configured area. The Secondary field is displayed if the interface is configured as a secondary interface. The output might also show the same interface listed in multiple areas.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Verifying the Interfaces in the Area

Purpose

Verify the interfaces configured for the specified area.

Action

From operational mode, enter the `show ospf interface area area-id` command for OSPFv2, and enter the `show ospf3 interface area area-id` command for OSPFv3..

Verifying Neighbor Adjacencies

Purpose

Verify the primary and secondary neighbor adjacencies. The Secondary field displays if the neighbor is on a secondary interface.

Action

From operational mode, enter the `show ospf neighbor detail` command for OSPFv2, and enter the `show ospf3 neighbor detail` command for OSPFv3.

Understanding Multiarea Adjacencies for OSPFv3

An area is a set of networks and hosts within an OSPFv3 domain that have been administratively grouped together. By default, a single interface can belong to only one OSPFv3 area. However, in some situations, you might want to configure an interface to belong to more than one area to avoid suboptimal routing. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over higher-cost intra-area links.

In Junos OS Release 9.2 and later, you can configure an interface to belong to more than one OSPFv2 area. Support for OSPFv3 was introduced in Junos OS Release 9.4. As defined in RFC 5185, *OSPF Multi-Area Adjacency*, the ABRs establish multiple adjacencies belonging to different areas over the same *logical interface*. Each multiarea adjacency is announced as a point-to-point unnumbered link in the configured area by the routers connected to the link.

An interface is considered to be primarily in one area. When you configure the same interface in another area, it is considered to be secondarily in the other area. You designate the secondary area by including the `secondary` statement at the `[edit protocols ospf3 area area-number interface interface-name]` hierarchy level.

Example: Configuring a Multiarea Adjacency for OSPFv3

IN THIS SECTION

- [Requirements | 80](#)
- [Overview | 80](#)
- [Configuration | 81](#)
- [Verification | 87](#)

This example shows how to configure a multiarea adjacency for OSPFv3.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

OSPFv3 intra-area paths are preferred over inter-area paths. In this example, Device R1 and Device R2 are area border routers (ABRs) with interfaces in both area 0 and in area 1. The link between Device R1 and R2 is in area 0 and is a high-speed link. The links in area 1 are lower speed.

If you want to forward some of area 1's traffic between Device R1 and Device R2 over the high-speed link, one method to accomplish this goal is to make the high-speed link a multiarea adjacency so that the link is part of both area 0 and area 1.

If the high-speed link between Device R1 and Device R2 remains in area 1 only, Device R1 always routes traffic to Device R4 and Device R5 through area 1 over the lower-speed links. Device R1 also uses the intra-area area 1 path through Device R3 to get to area 1 destinations downstream of Device R2.


Clearly, this scenario results in suboptimal routing.

An OSPF virtual link cannot be used to resolve this issue without moving the link between Device R1 and Device R2 to area 1. You might not want to do this if the physical link belongs to the network's backbone topology.

The OSPF/OSPFv3 protocol extension described in RFC 5185, *OSPF Multi-Area Adjacency* resolves the issue, by allowing the link between Device R1 and Device R2 to be part of both the backbone area and area 1.

To create a multiarea adjacency, you configure an interface to be in two areas, with ge-1/2/0 on Device R1 configured in both area 0 and area 1, and ge-1/2/0 on Device R2 configured in both area 0 and area 1. On both Device R1 and Device R2, area 0 contains the primary interface and is the primary link between the devices. Area 1 contains the secondary logical interface, which you configure by including the secondary statement.

Figure 8: OSPFv3 Multiarea Adjacency



"CLI Quick Configuration" on page 81 shows the configuration for all of the devices in Figure 8 on page 81. The section "No Link Title" on page 83 describes the steps on Device R1 and Device R2.

Configuration

IN THIS SECTION

- Procedure | 81

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R1

```
set interfaces ge-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:2::2/64
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:9009::1/128
```

```

set protocols ospf3 area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.1 interface ge-1/2/0.0 secondary

```

Device R2

```

set interfaces ge-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::1/64
set interfaces fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:6::1/64
set interfaces lo0 unit 0 family inet address 10.2.2.2/32
set interfaces lo0 unit 0 family inet6 address 2001:db9:9001::2/128
set protocols ospf3 area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/2.0
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.1 interface ge-1/2/0.0 secondary

```

Device R3

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:2::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:3::1/64
set interfaces lo0 unit 0 family inet address 10.3.3.3/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:9009::3/128
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0

```

Device R4

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:3::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::1/64
set interfaces fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:5::1/64
set interfaces lo0 unit 0 family inet address 10.4.4.4/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:9009::4/128
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/2.0

```

Device R5

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
set interfaces lo0 unit 0 family inet address 10.5.5.5/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:9009::5/128
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
```

Device R6

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:6::2/64
set interfaces lo0 unit 0 family inet address 10.6.6.6/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:9009::6/128
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces]
user@R1# set ge-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
user@R1# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:2::2/64
user@R1# set lo0 unit 0 family inet address 10.1.1.1/32
user@R1# set lo0 unit 0 family inet6 address 2001:db8:9009::1/128
```

2. Enable OSPFv3 on the interfaces that are in area 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@R1# set interface ge-1/2/0.0
user@R1# set interface lo0.0 passive
```

3. Enable OSPFv3 on the interface that is in area 1.

```
[edit protocols ospf3 area 0.0.0.1]
user@R1# set interface fe-1/2/1.0
user@R1# set interface ge-1/2/0.0 secondary
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces]
user@R2# set ge-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::2/64
user@R2# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::2/64
user@R2# set fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:6::2/64
user@R2# set lo0 unit 0 family inet address 10.2.2.2/32
user@R2# set lo0 unit 0 family inet6 address 2001:db8:9009::2/128
```

2. Enable OSPFv3 on the interfaces that are in area 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@R2# set interface ge-1/2/0.0
user@R2# set interface lo0.0 passive
```

3. Enable OSPFv3 on the interface that is in area 1.

```
[edit protocols ospf3 area 0.0.0.1]
user@R2# set interface fe-1/2/2.0
user@R2# set interface fe-1/2/1.0
user@R2# set interface ge-1/2/0.0 secondary
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces` and `show protocols` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device R1

```
user@R1# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:1::1/64/64;
    }
  }
}
fe-1/2/1 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:2::2/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.1.1.1/32;
    }
    family inet6 {
      address 2001:db8:9009::1/128;
    }
  }
}
```

```
user@R1# show protocols
ospf3 {
  area 0.0.0.0 {
    interface ge-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
```

```
}  
area 0.0.0.1 {  
    interface fe-1/2/1.0;  
    interface ge-1/2/0.0 {  
        secondary;  
    }  
}  
}
```

Device R2

```
user@R2# show interfaces  
ge-1/2/0 {  
    unit 0 {  
        family inet6 {  
            address 2001:db8:9009:1::2/64;  
        }  
    }  
}  
fe-1/2/1 {  
    unit 0 {  
        family inet6 {  
            address 2001:db8:9009:4::1/64;  
        }  
    }  
}  
fe-1/2/2 {  
    unit 0 {  
        family inet6 {  
            address 2001:db8:9009:6::2/64;  
        }  
    }  
}  
lo0 {  
    unit 0 {  
        family inet {  
            address 10.2.2.2/32;  
        }  
        family inet6 {  
            address 2001:db8:9009::2/128;  
        }  
    }  
}
```

```
}  
}
```

```
user@R2# show protocols  
ospf3 {  
  area 0.0.0.0 {  
    interface ge-1/2/0.0;  
    interface lo0.0 {  
      passive;  
    }  
  }  
  area 0.0.0.1 {  
    interface fe-1/2/2.0;  
    interface fe-1/2/1.0;  
    interface ge-1/2/0.0 {  
      secondary;  
    }  
  }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Flow of Traffic | 87](#)
- [Verifying That the Traffic Flow Changes When You Remove the Multiarea Adjacency | 88](#)

Confirm that the configuration is working properly.

Verifying the Flow of Traffic

Purpose

Verify that traffic uses the high-speed link between Device R1 and Device R2 to reach destinations in area 1.

Action

From operational mode on Device R1, use the traceroute command check the traffic flow to Device R5 and Device R6.

```
user@R1> traceroute 2001:db8:9009::6
traceroute6 to 2001:db8:9009::6 (2001:db8:9009::6 ) from 2001:db8:9009:1::1 , 64 hops max, 12
byte packets
 1 2001:db8:9009:1::2 (2001:db8:9009:1::2 ) 1.361 ms 1.166 ms 1.117 ms
 2 2001:db8:9009::6 (2001:db8:9009::6 ) 1.578 ms 1.484 ms 1.488 ms
```

```
user@R1> traceroute 2001:db8:9009::5
traceroute6 to 2001:db8:9009::5 (2001:db8:9009::5) from 2001:db8:9009:1::1, 64 hops max, 12 byte
packets
 1 2001:db8:9009:1::2 (2001:db8:9009:1::2) 1.312 ms 1.472 ms 1.132 ms
 2 2001:db8:9009:4::1 (2001:db8:9009:4::1) 1.137 ms 1.174 ms 1.126 ms
 3 2001:db8:9009::5 (5::5) 1.591 ms 1.445 ms 1.441 ms
```

Meaning

The traceroute output shows that traffic uses the 9009:1:: link between Device R1 and Device R2.

Verifying That the Traffic Flow Changes When You Remove the Multiarea Adjacency

Purpose

Verify the results without the multiarea adjacency configured.

Action

1. Deactivate the backbone link interfaces in area 1 on both R1 and R2.

```
user@R1# deactivate protocols ospf3 area 0.0.0.1 interface ge-1/2/0.0
user@R1# commit
user@R2# deactivate protocols ospf3 area 0.0.0.1 interface ge-1/2/0.0
user@R2# commit
```

- From operational mode on Device R1, use the traceroute command check the traffic flow to Device R5 and Device R6.

```

user@R1> traceroute 2001:db8:9009::6
traceroute6 to 2001:db8:9009::6 (2001:db8:9009::6) from 2001:db8:9009:2::2, 64 hops max, 12
byte packets
 1 2001:db8:9009:2::1 (2001:db8:9009:2::1) 1.314 ms 8.523 ms 8.310 ms
 2 2001:db8:9009:3::2 (2001:db8:9009:3::2) 1.166 ms 1.162 ms 1.172 ms
 3 2001:db8:9009:4::1 (2001:db8:9009:4::1) 1.386 ms 1.182 ms 1.138 ms
 4 2001:db8:9009::6 (2001:db8:9009::6) 1.605 ms 1.469 ms 1.438 ms

```

```

user@R1> traceroute 2001:db8:9009::5
traceroute6 to 2001:db8:9009::5 (2001:db8:9009::5) from 2001:db8:9009:2::2, 64 hops max, 12
byte packets
 1 2001:db8:9009:2::1 (2001:db8:9009:2::1) 1.365 ms 1.174 ms 1.133 ms
 2 2001:db8:9009:3::2 (2001:db8:9009:2::1) 1.157 ms 1.198 ms 1.138 ms
 3 2001:db8:9009:5::5 (2001:db8:9009:5::5) 1.584 ms 1.461 ms 1.443 ms

```


Meaning

Without the multiarea adjacency, the output shows suboptimal routing with traffic taking the path through the area 1 low-speed-links.

Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas

Figure 9 on page 90 shows an autonomous system (AS) across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

Figure 9: OSPF AS Network with Stub Areas and NSSAs



To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router (ABR) interface to the area as a stub interface, you suppress external route advertisements through the ABR. Instead, the ABR advertises a default route (through itself) in place of the external routes and generates network summary (Type 3) link-state advertisements (LSAs). Packets destined for external routes are automatically sent to the ABR, which acts as a gateway for outbound traffic and routes the traffic appropriately.



NOTE: You must explicitly configure the ABR to generate a default route when attached to a stub or not-so-stubby-area (NSSA). To inject a default route with a specified metric value into the area, you must configure the `default-metric` option and specify a metric value.

For example, area 0.0.0.3 in [Figure 9 on page 90](#) is not directly connected to the outside network. All outbound traffic is routed through the ABR to the backbone and then to the destination addresses. By designating area 0.0.0.3 as a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

A stub area that only allows routes internal to the area and restricts Type 3 LSAs from entering the stub area is often called a totally stubby area. You can convert area 0.0.0.3 to a totally stubby area by configuring the ABR to only advertise and allow the default route to enter into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area.



NOTE: If you incorrectly configure a totally stubby area, you might encounter network connectivity issues. You should have advanced knowledge of OSPF and understand your network environment before configuring totally stubby areas.

Similar to area 0.0.0.3 in [Figure 9 on page 90](#), area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route

advertisements to the area and advertise the static customer routes by designating the area an NSSA. In an NSSA, the AS boundary router generates NSSA external (Type 7) LSAs and floods them into the NSSA, where they are contained. Type 7 LSAs allow an NSSA to support the presence of AS boundary routers and their corresponding external routing information. The ABR converts Type 7 LSAs into AS external (Type 5) LSAs and leaks them to the other areas, but external routes from other areas are not advertised within the NSSA.

Example: Configuring OSPF Stub and Totally Stubby Areas

IN THIS SECTION

- Requirements | 91
- Overview | 92
- Configuration | 93
- Verification | 96

This example shows how to configure an OSPF stub area and a totally stubby area to control the advertisement of external routes into an area.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 93](#)

The backbone area, which is 0 in [Figure 10 on page 93](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an autonomous system (AS). All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by area border routers (ABRs) that have interfaces in more than one area.

Stub areas are areas through which or into which OSPF does not flood AS external link-state advertisements (Type 5 LSAs). You might create stub areas when much of the topology database consists of AS external advertisements and you want to minimize the size of the topology databases on the internal routers in the stub area.

The following restrictions apply to stub areas:

- You cannot create a virtual link through a stub area.
- A stub area cannot contain an AS boundary router.
- You cannot configure the backbone as a stub area.
- You cannot configure an area as both a stub area and an not-so-stubby area (NSSA).

In this example, you configure each routing device in area 7 (area ID 0.0.0.7) as a stub router and some additional settings on the ABR:

- `stub`—Specifies that this area become a stub area and not be flooded with Type 5 LSAs. You must include the `stub` statement on all routing devices that are in area 7 because this area has no external connections.
- `default-metric`—Configures the ABR to generate a default route with a specified metric into the stub area. This default route enables packet forwarding from the stub area to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to a stub. You must explicitly configure this option to generate a default route.
- `no-summaries`—(Optional) Prevents the ABR from advertising summary routes into the stub area by converting the stub area into a totally stubby area. If configured in combination with the `default-metric` statement, a totally stubby area only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized


or allowed into a totally stubby area. Only the ABR requires this additional configuration because it is the only routing device within the totally stubby area that creates Type 3 LSAs used to receive and send traffic from outside of the area.



NOTE: In Junos OS Release 8.5 and later, the following applies:

- A router-identifier interface that is not configured to run OSPF is no longer advertised as a stub network in OSPF LSAs.
- OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also advertises the direct route with the configured mask length, as in earlier releases.

Figure 10: OSPF Network Topology with Stub Areas and NSSAs



Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 94](#)
- [Procedure | 94](#)
- [Results | 95](#)

CLI Quick Configuration

- To quickly configure an OSPF stub area, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the stub area.

```
[edit]
set protocols ospf area 07 stub
```

- To quickly configure the ABR to inject a default route into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

```
[edit]
set protocols ospf area 07 stub default-metric 10
```

- (Optional) To quickly configure the ABR to restrict all summary advertisements and allow only internal routes and default route advertisements into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

```
[edit]
set protocols ospf area 0.0.0.7 stub no-summaries
```

Procedure

Step-by-Step Procedure

To configure OSPF stub areas:

1. On all routing devices in the area, configure an OSPF stub area.



NOTE: To specify an OSPFv3 stub area, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub
```

2. On the ABR, inject a default route into the area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub default-metric 10
```

3. (Optional) On the ABR, restrict summary LSAs from entering the area. This step converts the stub area into a totally stubby area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub no-summaries
```

4. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices:

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub;
}
```

Configuration on the ABR (the output also includes the optional setting):

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub default-metric 10 no-summaries;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Interfaces in the Area | 96](#)
- [Verifying the Type of OSPF Area | 96](#)

Confirm that the configuration is working properly.

Verifying the Interfaces in the Area

Purpose

Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub as the type of OSPF area.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Verifying the Type of OSPF Area

Purpose

Verify that the OSPF area is a stub area. Confirm that the output displays Normal Stub as the Stub type.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Example: Configuring OSPF Not-So-Stubby Areas

IN THIS SECTION

- [Requirements | 97](#)
- [Overview | 97](#)
- [Configuration | 99](#)
- [Verification | 103](#)

This example shows how to configure an OSPF not-so-stubby area (NSSA) to control the advertisement of external routes into an area.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 99](#)

The backbone area, which is 0 in [Figure 11 on page 99](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by ABRs that have interfaces in more than one area.

An OSPF stub area has no external routes, so you cannot redistribute routes from another protocol into a stub area. OSPF NSSAs allow external routes to be flooded within the area.

In addition, you might have a situation when exporting Type 7 LSAs into the NSSA is unnecessary. When an AS boundary router is also an ABR with an NSSA attached, Type 7 LSAs are exported into the NSSA by default. If the ABR is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. During route redistribution, this routing device generates both Type 5 LSAs and Type 7 LSAs. You can disable exporting Type 7 LSAs into the NSSA.



NOTE: The following restriction applies to NSSAs: You cannot configure an area as both a stub area and an NSSA.

You configure each routing device in area 9 (area ID 0.0.0.9) with the following setting:

- `nssa`—Specifies an OSPF NSSA. You must include the `nssa` statement on all routing devices in area 9 because this area only has external connections to static routes.


You also configure the ABR in area 9 with the following additional settings:

- `no-summaries`—Prevents the ABR from advertising summary routes into the NSSA. If configured in combination with the `default-metric` statement, the NSSA only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into the NSSA. Only the ABR requires this additional configuration because it is the only routing device within the NSSA that creates Type 3 LSAs used to receive and send traffic from outside the area.
- `default-lsa`—Configures the ABR to generate a default route into the NSSA. In this example, you configure the following:
 - `default-metric`—Specifies that the ABR generate a default route with a specified metric into the NSSA. This default route enables packet forwarding from the NSSA to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to an NSSA. You must explicitly configure this option for the ABR to generate a default route.
 - `metric-type`—(Optional) Specifies the external metric type for the default LSA, which can be either Type 1 or Type 2. When OSPF exports route information from external ASs, it includes a cost, or external metric, in the route. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. Type 2 external metrics use only the external cost assigned by the AS boundary router. By default, OSPF uses the Type 2 external metric.
 - `type-7`—(Optional) Floods Type 7 default LSAs into the NSSA if the `no-summaries` statement is configured. By default, when the `no-summaries` statement is configured, a Type 3 LSA is injected into

NSSAs for Junos OS release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the type-7 statement.

The second example also shows the optional configuration required to disable exporting Type 7 LSAs into the NSSA by including the no-nssa-abr statement on the routing device that performs the functions of both an ABR and an AS boundary router.

Figure 11: OSPF Network Topology with Stub Areas and NSSAs



Topology

Configuration

IN THIS SECTION

- [Configuring Routing Devices to Participate in a Not-So-Stubby-Area | 100](#)
- [Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas | 102](#)

Configuring Routing Devices to Participate in a Not-So-Stubby-Area

CLI Quick Configuration

To quickly configure an OSPF NSSA, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the NSSA.

```
[edit]
set protocols ospf area 0.0.0.9 nssa
```

To quickly configure an ABR that participates in an OSPF NSSA, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.9 nssa default-lsa default-metric 10
set protocols ospf area 0.0.0.9 nssa default-lsa metric-type 1
set protocols ospf area 0.0.0.9 nssa default-lsa type-7
set protocols ospf area 0.0.0.9 nssa no-summaries
```

Step-by-Step Procedure

To configure OSPF NSSAs:

1. On all routing devices in the area, configure an OSPF NSSA.



NOTE: To specify an OSPFv3 NSSA area, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.9 nssa
```

2. On the ABR, enter OSPF configuration mode and specify the NSSA area 0.0.0.9 that you already created.

```
[edit ]
user@host# edit protocols ospf area 0.0.0.9 nssa
```

3. On the ABR, inject a default route into the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa default-metric 10
```

4. (Optional) On the ABR, specify the external metric type for the default route.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa metric-type 1
```

5. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa type-7
```

6. On the ABR, restrict summary LSAs from entering the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set no-summaries
```

7. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices in the area:

```
user@host# show protocols ospf
area 0.0.0.9 {
    nssa;
}
```

Configuration on the ABR. The output also includes the optional `metric-type` and `type-7` statements.

```
user@host# show protocols ospf
  area 0.0.0.9 {
    nssa {
      default-lsa {
        default-metric 10;
        metric-type 1;
        type-7;
      }
      no-summaries;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas

CLI Quick Configuration

To quickly disable exporting Type 7 LSAs into the NSSA, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode. You configure this setting on an AS boundary router that is also an ABR with an NSSA area attached.

```
[edit]
set protocols ospf no-nssa-abr
```

Step-by-Step Procedure

You can configure this setting if you have an AS boundary router that is also an ABR with an NSSA area attached.

1. Disable exporting Type 7 LSAs into the NSSA.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf no-nssa-abr
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-nssa-abr;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Interfaces in the Area | 104](#)
- [Verifying the Type of OSPF Area | 104](#)
- [Verifying the Type of LSAs | 104](#)

Confirm that the configuration is working properly.

Verifying the Interfaces in the Area

Purpose

Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub NSSA as the type of OSPF area.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Verifying the Type of OSPF Area

Purpose

Verify that the OSPF area is a stub area. Confirm that the output displays Not so Stubby Stub as the Stub type.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Verifying the Type of LSAs

Purpose

Verify the type of LSAs that are in the area. If you disabled exporting Type 7 LSAs into an NSSA, confirm that the Type field does not include NSSA as a type of LSA.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Understanding OSPFv3 Stub and Totally Stubby Areas

Junos OS OSPFv3 configuration for IPv6 networks is identical to OSPFv2 configuration. You configure the protocol with `set ospf3` commands instead of `set ospf` commands and use `show ospf3` commands instead of `show ospf` commands to check the OSPF status. Also, make sure to set IPv6 addresses on the interfaces running OSPFv3.

Stub areas are areas through which or into which OSPF does not flood AS external link-state advertisements (Type 5 LSAs). You might create stub areas when much of the topology database consists of AS external advertisements and you want to minimize the size of the topology databases on the internal routers in the stub area.

The following restrictions apply to stub areas:

- You cannot create a virtual link through a stub area.
- A stub area cannot contain an AS boundary router.
- You cannot configure the backbone as a stub area.
- You cannot configure an area as both a stub area and an not-so-stubby area (NSSA).

Example: Configuring OSPFv3 Stub and Totally Stubby Areas

IN THIS SECTION

- [Requirements | 105](#)
- [Overview | 106](#)
- [Configuration | 107](#)
- [Verification | 117](#)

This example shows how to configure an OSPFv3 stub area and a totally stubby area to control the advertisement of external routes into an area.


Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Figure 12 on page 106 shows the topology used in this example.

Figure 12: OSPFv3 Network Topology with Stub Areas



In this example, you configure each routing device in area 7 (area ID 0.0.0.7) as a stub router and some additional settings on the ABR:

- **stub**—Specifies that this area become a stub area and not be flooded with Type 5 LSAs. You must include the `stub` statement on all routing devices that are in area 7 because this area has no external connections.
- **default-metric**—Configures the ABR to generate a default route with a specified metric into the stub area. This default route enables packet forwarding from the stub area to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to a stub. You must explicitly configure this option to generate a default route.
- **no-summaries**—(Optional) Prevents the ABR from advertising summary routes into the stub area by converting the stub area into a totally stubby area. If configured in combination with the `default-metric` statement, a totally stubby area only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized

or allowed into a totally stubby area. Only the ABR requires this additional configuration because it is the only routing device within the totally stubby area that creates Type 3 LSAs used to receive and send traffic from outside of the area.



NOTE: In Junos OS Release 8.5 and later, the following applies:

- A router-identifier interface that is not configured to run OSPF is no longer advertised as a stub network in OSPF LSAs.
- OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also advertises the direct route with the configured mask length, as in earlier releases.

"CLI Quick Configuration" on page 107 shows the configuration for all of the devices in Figure 12 on page 106. The section "No Link Title" on page 109 describes the steps on Device 2, Device 6, Device 7, and Device 8.

Configuration

IN THIS SECTION

- Procedure | 107

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device 1

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:2::1/64
set interfaces fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:3::1/64
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface fe-1/2/1.0
```

```
set protocols ospf3 area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

Device 2

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:2::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::1/64
set interfaces lo0 unit 0 family inet address 10.2.2.2/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.7 stub default-metric 10
set protocols ospf3 area 0.0.0.7 stub no-summaries
set protocols ospf3 area 0.0.0.7 interface fe-1/2/1.0
```

Device 3

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:3::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:5::1/64
set interfaces lo0 unit 0 family inet address 10.3.3.3/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.9 interface fe-1/2/1.0
```

Device 4

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:6::1/64
set interfaces lo0 unit 0 family inet address 10.4.4.4/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.3 interface fe-1/2/1.0
```

Device 5

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:6::2/64
set interfaces lo0 unit 0 family inet address 10.5.5.5/32
set protocols ospf3 area 0.0.0.3 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.3 interface lo0.0 passive
```

Device 6

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:4::2/64
set interfaces lo0 unit 0 family inet address 10.6.6.6/32
set protocols ospf3 area 0.0.0.7 stub
set protocols ospf3 area 0.0.0.7 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.7 interface lo0.0 passive
```

Device 7

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:7::1/64
set interfaces lo0 unit 0 family inet address 10.7.7.7/32
set protocols ospf3 export static-to-ospf
set protocols ospf3 area 0.0.0.9 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.9 interface lo0.0 passive
set policy-options policy-statement static-to-ospf term 1 from protocol static
set policy-options policy-statement static-to-ospf term 1 then accept
set routing-options rib inet6.0 static route 2001:db8:1010::1/128 next-hop 2001:db8:9009:7::2
set routing-options rib inet6.0 static route 2001:db8:2020::1/128 next-hop 2001:db8:9009:7::2
```

Device 8

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:7::2/64
set interfaces lo0 unit 0 family inet address 10.8.8.8/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:1010::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:2020::1/128
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 2:

1. Configure the interfaces.

```
[edit interfaces]
user@2# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:2::2/64
```

```
user@2# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::1/64
user@2# set lo0 unit 0 family inet address 10.2.2.2/32
```

2. Enable OSPFv3 on the interfaces that are in area 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@2# set interface fe-1/2/0.0
user@2# set interface lo0.0 passive
```

3. Enable OSPFv3 on the interface that is in area 7.

```
[edit protocols ospf3 area 0.0.0.7]
user@2# set interface fe-1/2/1.0
```

4. Specify area 7 as an OSPFv3 stub area.

The stub statement is required on all routing devices in the area.

```
[edit protocols ospf3 area 0.0.0.7]
user@2# set stub
```

5. On the ABR, inject a default route into the area.

```
[edit protocols ospf3 area 0.0.0.7]
user@2# set stub default-metric 10
```

6. (Optional) On the ABR, restrict summary LSAs from entering the area.

This step converts the stub area into a totally stubby area.

```
[edit protocols ospf3 area 0.0.0.7]
user@2# set stub no-summaries
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 6:

1. Configure the interfaces.

```
[edit interfaces]
user@6# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:4::2/64
user@6# set lo0 unit 0 family inet address 10.6.6.6/32
```

2. Enable OSPFv3 on the interface that is in area 7.

```
[edit protocols ospf3 area 0.0.0.7]
user@6# set interface fe-1/2/0.0
user@6# set interface lo0.0 passive
```

3. Specify area 7 as an OSPFv3 stub area.

The stub statement is required on all routing devices in the area.

```
[edit protocols ospf3 area 0.0.0.7]
user@6# set stub
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 7:

1. Configure the interfaces.

```
[edit interfaces]
user@7# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
user@7# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:7::1/64
user@7# set lo0 unit 0 family inet address 10.7.7.7/32
```

2. Enable OSPFv3 on the interface that is in area 9.

```
[edit protocols ospf3 area 0.0.0.9]
user@7# set interface fe-1/2/0.0
user@7# set interface lo0.0 passive
```

3. Configure static routes that enable connectivity to the customer routes.

```
[edit routing-options rib inet6.0 static]
user@7# set route 1010::1/128 next-hop 2001:db8:9009:7::2
user@7# set route 2020::1/128 next-hop 2001:db8:9009:7::2
```

4. Configure a routing policy to redistribute the static routes.

```
[edit policy-options policy-statement static-to-ospf term 1]
user@7# set from protocol static
user@7# set then accept
```

5. Apply the routing policy to the OSPFv3 instance.

```
[edit protocols ospf3]
user@7# set export static-to-ospf
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 8:

1. Configure the interfaces.

```
[edit interfaces]
user@8# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:7::2/64
user@8# set lo0 unit 0 family inet address 10.8.8.8/32
```

2. Configure two loopback interface addresses to simulate customer routes.

```
[edit interfaces lo0 unit 0 family inet6]
user@8# set address 2001:db8:1010::1/128
user@8# set address 2001:db8:2020::1/128
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device 2

```
user@2# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:2::2/64;
    }
  }
}
fe-1/2/1 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:4::1/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.2.2.2/32;
    }
  }
}
```

```
user@2# show protocols
ospf3 {
  area 0.0.0.0 {
```

```
interface fe-1/2/0.0;
interface lo0.0 {
    passive;
}
}
area 0.0.0.7 {
    stub default-metric 10 no-summaries;
    interface fe-1/2/1.0;
}
}
```

Device 6

```
user@6# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet6 {
            address 2001:db8:9009:4::2/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.6.6.6/32;
        }
    }
}
```

```
user@6# show protocols
ospf3 {
    area 0.0.0.7 {
        stub;
        interface fe-1/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
```

Device 7

```
user@7# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:5::2/64;
    }
  }
}
fe-1/2/1 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:7::1/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.7.7.7/32;
    }
  }
}
```

```
user@7# show protocols
ospf3 {
  export static-to-ospf;
  area 0.0.0.9 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
```

```
user@7# show policy-options
policy-statement static-to-ospf {
  term 1 {
```

```

        from protocol static;
        then accept;
    }
}

```

```

user@7# show routing-options
rib inet6.0 {
    static {
        route 2001:db8:1010::1/128 next-hop 2001:db8:9009:7::2;
        route 2001:db8:2020::1/128 next-hop 2001:db8:9009:7::2;
    }
}

```

Device 8

```

user@8# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet6 {
            address 2001:db8:9009:7::2/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.8.8.8/32;
        }
        family inet6 {
            address 2001:db8:1010::1/128;
            address 2001:db8:2020::1/128;
        }
    }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Type of OSPFv3 Area | 117](#)
- [Verifying the Routes in the OSPFv3 Stub Area | 118](#)

Confirm that the configuration is working properly.

Verifying the Type of OSPFv3 Area

Purpose

Verify that the OSPFv3 area is a stub area. Confirm that the output displays Stub as the Stub type.

Action

From operational mode on Device 2 and on Device 6, enter the `show ospf3 overview` command.

```
user@2> show ospf3 overview
Instance: master
  Router ID: 10.2.2.2
  Route table index: 51
  Area border router
  LSA refresh time: 50 minutes
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 2, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Area: 0.0.0.7
    Stub type: Stub, Stub cost: 10
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 24
```

```
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

```
user@6> show ospf3 overview
Instance: master
Router ID: 10.6.6.6
Route table index: 46
LSA refresh time: 50 minutes
Area: 0.0.0.7
  Stub type: Stub
  Area border routers: 1, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 17
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

Meaning

On Device 2, the stub type of area 0 is Not Stub. The stub type of area 7 is Stub. The stub default metric is 10.

On Device 6, the stub type of area 7 is Stub.

Verifying the Routes in the OSPFv3 Stub Area

Purpose

Make sure that the expected routes are present in the routing tables.

Action

From operational mode on Device 6 and Device 2, enter the `show route` command.

```
user@6> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

10.6.6.6/32      *[Direct/0] 1d 01:57:12
                 > via lo0.0

inet6.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0           *[OSPF3/10] 00:10:52, metric 11
                > via fe-1/2/0.0
2001:db8:9009:4::/64      *[Direct/0] 1d 01:56:31
                > via fe-1/2/0.0
                [OSPF3/10] 1d 01:56:31, metric 1
                > via fe-1/2/0.0
2001:db8:9009:4::2/128    *[Local/0] 1d 01:56:53
                Local via fe-1/2/0.0
fe80::/64      *[Direct/0] 1d 01:56:31
                > via fe-1/2/0.0
fe80::2a0:a514:0:a4c/128
                *[Local/0] 1d 01:56:53
                Local via fe-1/2/0.0
ff02::5/128    *[OSPF3/10] 1d 01:58:22, metric 1
                MultiRecv

```

```

user@2> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.2.2.2/32    *[Direct/0] 1d 02:16:13
                > via lo0.0

inet6.0: 14 destinations, 17 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:1010::1/128      *[OSPF3/150] 00:30:15, metric 0, tag 0
                > via fe-1/2/0.0
2001:db8:2020::1/128      *[OSPF3/150] 00:30:15, metric 0, tag 0
                > via fe-1/2/0.0
2001:db8:9009:1::/64      *[OSPF3/10] 1d 02:15:54, metric 2
                > via fe-1/2/0.0
2001:db8:9009:2::/64      *[Direct/0] 1d 02:15:54
                > via fe-1/2/0.0
                [OSPF3/10] 1d 02:15:54, metric 1

```

```

                > via fe-1/2/0.0
2001:db8:9009:2::2/128      *[Local/0] 1d 02:15:54
                        Local via fe-1/2/0.0
2001:db8:9009:3::/64      *[OSPF3/10] 1d 02:15:54, metric 2
                > via fe-1/2/0.0
2001:db8:9009:4::/64      *[Direct/0] 1d 02:15:54
                > via fe-1/2/1.0
                        [OSPF3/10] 05:38:05, metric 1
                > via fe-1/2/1.0
2001:db8:9009:4::1/128    *[Local/0] 1d 02:15:54
                        Local via fe-1/2/1.0
2001:db8:9009:5::/64      *[OSPF3/10] 1d 02:15:54, metric 3
                > via fe-1/2/0.0
2001:db8:9009:6::/64      *[OSPF3/10] 1d 01:33:10, metric 3
                > via fe-1/2/0.0
fe80::/64                 *[Direct/0] 1d 02:15:54
                > via fe-1/2/0.0
                        [Direct/0] 1d 02:15:54
                > via fe-1/2/1.0
fe80::2a0:a514:0:64c/128
                        *[Local/0] 1d 02:15:54
                        Local via fe-1/2/0.0
fe80::2a0:a514:0:94c/128
                        *[Local/0] 1d 02:15:54
                        Local via fe-1/2/1.0
ff02::5/128              *[OSPF3/10] 1d 02:17:45, metric 1
                        MultiRecv

```

Meaning

On Device 6, the default route has been learned because of the `default-metric` statement on the ABR, Device 2. Otherwise, the only OSPFv3 routes in Device 6's routing table are the network address `2001:db8:9009:4::/64` and the OSPFv3 multicast address `ff02::5/128` for all SPF link-state routers, also known as AllSPFRouters.

On Device 2, all of the OSPFv3 routes have been learned, including the external customer routes, `2001:db8:1010::1/128` and `2001:db8:2020::1/128`.

Understanding OSPFv3 Not-So-Stubby Areas

Like an OSPF stub area, an OSPFv3 stub area has no external routes, so you cannot redistribute routes from another protocol into a stub area. Not-so-stubby-areas (NSSAs) allow external routes to be flooded within the area. Routers in an NSSA do not receive external link-state advertisements (LSAs) from area border routers (ABRs), but are allowed to send external routing information for redistribution. They use type 7 LSAs to tell the ABRs about these external routes, which the ABR then translates to type 5 external LSAs and floods as normal to the rest of the OSPF network.

Example: Configuring OSPFv3 Not-So-Stubby Areas

IN THIS SECTION

- [Requirements | 121](#)
- [Overview | 121](#)
- [Configuration | 123](#)
- [Verification | 134](#)

This example shows how to configure an OSPFv3 not-so-stubby area (NSSA) to control the advertisement of external routes into the area.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device 7 redistributes static Customer 1 routes into OSPFv3. Device 7 is in area 9, which is configured as an NSSA. Device 3 is the ABR attached to the NSSA. An NSSA is a type of stub area that can import autonomous system external routes and send them to other areas, but still cannot receive AS-external routes from other areas. Because area 9 is defined as an NSSA, Device 7 uses type 7 LSAs to tell the ABR (Device 3) about these external routes. Device 3 then translates the type 7 routes to type 5 external LSAs and floods them as normal to the rest of the OSPF network.

In area 3, Device 5 redistributes static Customer 2 routes into OSPFv3. These routes are learned on Device 3, but not on Device 7 or 10. Device 3 injects a default static route into area 9 so that Device 7 and 10 can still reach the Customer 2 routes.


You configure each routing device in area 9 (area ID 0.0.0.9) with the following setting:

- `nssa`—Specifies an OSPFv3 NSSA. You must include the `nssa` statement on all routing devices in area 9.

You also configure the ABR in area 9 with the following additional settings:

- `no-summaries`—Prevents the ABR from advertising summary routes into the NSSA. If configured in combination with the `default-metric` statement, the NSSA only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into the NSSA. Only the ABR requires this additional configuration because it is the only routing device within the NSSA that creates Type 3 summary LSAs used to receive and send traffic from outside the area.
- `default-lsa`—Configures the ABR to generate a default route into the NSSA. In this example, you configure the following:
 - `default-metric`—Specifies that the ABR generate a default route with a specified metric into the NSSA. This default route enables packet forwarding from the NSSA to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to an NSSA. You must explicitly configure this option for the ABR to generate a default route.
 - `metric-type`—(Optional) Specifies the external metric type for the default LSA, which can be either Type 1 or Type 2. When OSPFv3 exports route information from external ASs, it includes a cost, or external metric, in the route. The difference between the two metrics is how OSPFv3 calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. Type 2 external metrics use only the external cost assigned by the AS boundary router. By default, OSPFv3 uses the Type 2 external metric.
 - `type-7`—(Optional) Floods Type 7 default LSAs into the NSSA if the `no-summaries` statement is configured. By default, when the `no-summaries` statement is configured, a Type 3 LSA is injected into NSSAs for Junos OS release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the `type-7` statement.

Figure 13: OSPFv3 Network Topology with an NSSA



"CLI Quick Configuration" on page 124 shows the configuration for all of the devices in Figure 13 on page 123. The section "No Link Title" on page 126 describes the steps on Device 3, Device 7, and Device 9.

Configuration

IN THIS SECTION

- Procedure | 124

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device 1

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:3::1/64
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf3 area 0.0.0.0 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

Device 3

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:3::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:5::1/64
set interfaces lo0 unit 0 family inet address 10.3.3.3/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.9 nssa default-lsa default-metric 10
set protocols ospf3 area 0.0.0.9 nssa default-lsa metric-type 1
set protocols ospf3 area 0.0.0.9 nssa default-lsa type-7
set protocols ospf3 area 0.0.0.9 nssa no-summaries
set protocols ospf3 area 0.0.0.9 interface fe-1/2/1.0
```

Device 4

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:6::1/64
set interfaces lo0 unit 0 family inet address 10.4.4.4/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.3 interface fe-1/2/1.0
```

Device 5

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:6::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:7::1/64
set interfaces lo0 unit 0 family inet address 10.5.5.5/32
set protocols ospf3 export static-to-ospf
set protocols ospf3 area 0.0.0.3 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.3 interface lo0.0 passive
set policy-options policy-statement static-to-ospf term 1 from protocol static
set policy-options policy-statement static-to-ospf term 1 then accept
set routing-options rib inet6.0 static route 2001:db8:1010::1/128 next-hop 2001:db8:9009:7::2
set routing-options rib inet6.0 static route 2001:db8:2020::1/128 next-hop 2001:db8:9009:7::2

```

Device 7

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:8::1/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:9::1/64
set interfaces lo0 unit 0 family inet address 10.7.7.7/32
set protocols ospf3 export static2-to-ospf
set protocols ospf3 area 0.0.0.9 nssa
set protocols ospf3 area 0.0.0.9 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.9 interface lo0.0 passive
set policy-options policy-statement static2-to-ospf term 1 from protocol static
set policy-options policy-statement static2-to-ospf term 1 then accept
set routing-options rib inet6.0 static route 2001:db8:3030::1/128 next-hop 2001:db8:9009:8::2
set routing-options rib inet6.0 static route 2001:db8:4040::1/128 next-hop 2001:db8:9009:8::2

```

Device 8

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:7::2/64
set interfaces lo0 unit 0 family inet address 10.8.8.8/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:1010::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:2020::1/128

```

Device 9

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:8::2/64
set interfaces lo0 unit 0 family inet address 10.9.9.9/32

```

```
set interfaces lo0 unit 0 family inet6 address 2001:db8:3030::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:4040::1/128
```

Device 10

```
set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:9::2/64
set interfaces lo0 unit 0 family inet address 10.10.10.10/32
set protocols ospf3 area 0.0.0.9 nssa
set protocols ospf3 area 0.0.0.9 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.9 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.9 interface lo0.0 passive
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 3:

1. Configure the interfaces.

```
[edit interfaces]
user@3# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:3::2/64
user@3# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:5::1/64
user@3# set lo0 unit 0 family inet address 10.3.3.3/32
```

2. Enable OSPFv3 on the interfaces that are in area 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@3# set interface fe-1/2/0.0
user@3# set interface lo0.0 passive
```

3. Enable OSPFv3 on the interface that is in area 9.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set interface fe-1/2/1.0
```

4. Configure an OSPFv3 NSSA.

The nssa statement is required on all routing devices in the area.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set nssa
```

5. On the ABR, inject a default route into the area.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set default-lsa default-metric 10
```

6. (Optional) On the ABR, specify the external metric type for the default route.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set nssa default-lsa metric-type 1
```

7. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set nssa default-lsa type-7
```

8. On the ABR, restrict summary LSAs from entering the area.

```
[edit protocols ospf3 area 0.0.0.9]
user@3# set nssa no-summaries
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 5:

1. Configure the interfaces.

```
[edit interfaces]
user@5# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:6::2/64
user@5# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:7::1/64
user@5# set lo0 unit 0 family inet address 10.5.5.5/32
```

2. Enable OSPFv3 on the interface that is in area 3.

```
[edit protocols ospf3 area 0.0.0.3]
user@5# set interface fe-1/2/0.0
user@5# set interface lo0.0 passive
```

3. Configure static routes that enable connectivity to the customer routes.

```
[edit routing-options rib inet6.0 static]
user@5# set route 1010::1/128 next-hop 2001:db8:9009:7::2
user@5# set route 2020::1/128 next-hop 2001:db8:9009:7::2
```

4. Configure a routing policy to redistribute the static routes.

```
[edit policy-options policy-statement static-to-ospf term 1]
user@5# set from protocol static
user@5# set then accept
```

5. Apply the routing policy to the OSPFv3 instance.

```
[edit protocols ospf3]
user@5# set export static-to-ospf
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 7:

1. Configure the interfaces.

```
[edit interfaces]
user@7# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
user@7# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:7::1/64
user@7# set lo0 unit 0 family inet address 10.7.7.7/32
```

2. Enable OSPFv3 on the interface that is in area 9.

```
[edit protocols ospf3 area 0.0.0.9]
user@7# set interface fe-1/2/0.0
user@7# set interface lo0.0 passive
```

3. Configure an OSPFv3 NSSA.

The nssa statement is required on all routing devices in the area.

```
[edit protocols ospf3 area 0.0.0.9]
user@7# set nssa
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 8:

1. Configure the interfaces.

```
[edit interfaces]
user@8# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:7::2/64
user@8# set lo0 unit 0 family inet address 10.8.8.8/32
```

2. Configure two loopback interface addresses to simulate customer routes.

```
[edit interfaces lo0 unit 0 family inet6]
user@8# set address 2001:db8:1010::1/128
user@8# set address 2001:db8:2020::1/128
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device 3

```
user@3# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:3::2/64;
    }
  }
}
fe-1/2/1 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:5::1/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.3.3.3/32;
    }
  }
}
}
```

```
user@3# show protocols
ospf3 {
```

```
area 0.0.0.0 {
    interface fe-1/2/0.0;
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.9 {
    nssa {
        default-lsa {
            default-metric 10;
            metric-type 1;
            type-7;
        }
        no-summaries;
    }
    interface fe-1/2/1.0;
}
}
```

Device 5

```
user@5# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet6 {
            address 2001:db8:9009:6::2/64;
        }
    }
}
fe-1/2/1 {
    unit 0 {
        family inet6 {
            address 2001:db8:9009:7::1/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.5.5.5/32;
        }
    }
}
```

```

    }
}

```

```

user@5# show protocols
ospf3 {
  export static-to-ospf;
  area 0.0.0.3 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

```

```

user@5# show policy-options
policy-statement static-to-ospf {
  term 1 {
    from protocol static;
    then accept;
  }
}

```

```

user@5# show routing-options
rib inet6.0 {
  static {
    route 2001:db8:1010::1/128 next-hop 2001:db8:9009:7::2;
    route 2001:db8:2020::1/128 next-hop 2001:db8:9009:7::2;
  }
}

```

Device 7

```

user@7# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:5::2/64;
    }
  }
}

```

```

    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.7.7.7/32;
        }
    }
}
}

```

```

user@7# show protocols
ospf3 {
    area 0.0.0.9 {
        nssa;
        interface fe-1/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
}

```

Device 8

```

user@8# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet6 {
            address 2001:db8:9009:7::2/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.8.8.8/32;
        }
        family inet6 {
            address 2001:db8:1010::1/128;
            address 2001:db8:2020::1/128;
        }
    }
}

```

```
}
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Type of OSPFv3 Area | 134](#)
- [Verifying the Routes in the OSPFv3 Stub Area | 136](#)
- [Verifying the Type of LSAs | 140](#)

Confirm that the configuration is working properly.

Verifying the Type of OSPFv3 Area

Purpose

Verify that the OSPFv3 area is an NSSA area. Confirm that the output displays `Stub NSSA` as the Stub type.

Action

From operational mode on Device 3, Device 7, and Device 10 enter the `show ospf3 overview` command.

```
user@3> show ospf3 overview
Instance: master
Router ID: 10.3.3.3
Route table index: 36
Area border router, AS boundary router, NSSA router
LSA refresh time: 50 minutes
Area: 0.0.0.0
  Stub type: Not Stub
Area border routers: 2, AS boundary routers: 0
Neighbors
  Up (in full state): 1
Area: 0.0.0.9
  Stub type: Stub NSSA, Stub cost: 10
```

```
Area border routers: 0, AS boundary routers: 1
Neighbors
  Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 22
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

```
user@7> show ospf3 overview
```

```
Instance: master
Router ID: 10.7.7.7
Route table index: 44
AS boundary router, NSSA router
LSA refresh time: 50 minutes
Area: 0.0.0.9
  Stub type: Stub NSSA
  Area border routers: 1, AS boundary routers: 1
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 2
Full SPF runs: 11
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

```
user@10> show ospf3 overview
```

```
Instance: master
Router ID: 10.10.10.10
Route table index: 55
NSSA router
LSA refresh time: 50 minutes
Area: 0.0.0.9
  Stub type: Stub NSSA
  Area border routers: 1, AS boundary routers: 2
  Neighbors
    Up (in full state): 2
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 6
```

```
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

Meaning

On Device 3, the stub type of area 0 is Not Stub. The stub type of area 9 is Stub NSSA. The stub default metric is 10.

On Device 7 and Device 10, the stub type of area 9 is Stub NSSA.

Verifying the Routes in the OSPFv3 Stub Area

Purpose

Make sure that the expected routes are present in the routing tables.

Action

From operational mode on Device 7 and Device 3, enter the `show route` command.

```
user@7> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.7.7.7/32          *[Direct/0] 3d 03:00:23
                    > via lo0.0

inet6.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                *[OSPF3/150] 01:01:31, metric 12, tag 0
                    > via fe-1/2/1.0
2001:db8:3030::1/128  *[Static/5] 01:01:43
                    > to 9009:8::2 via fe-1/2/0.0
2001:db8:4040::1/128  *[Static/5] 01:01:43
                    > to 9009:8::2 via fe-1/2/0.0
2001:db8:9009:5::/64  *[OSPF3/10] 01:01:33, metric 2
                    > via fe-1/2/1.0
2001:db8:9009:8::/64  *[Direct/0] 01:01:43
                    > via fe-1/2/0.0
2001:db8:9009:8::1/128 *[Local/0] 01:02:01
```

```

                Local via fe-1/2/0.0
2001:db8:9009:9::/64      *[Direct/0] 01:01:45
                        > via fe-1/2/1.0
                        [OSPF3/10] 01:01:44, metric 1
                        > via fe-1/2/1.0
2001:db8:9009:9::1/128  *[Local/0] 01:02:01
                        Local via fe-1/2/1.0
fe80::/64                *[Direct/0] 01:01:45
                        > via fe-1/2/1.0
                        [Direct/0] 01:01:43
                        > via fe-1/2/0.0
fe80::2a0:a514:0:f4c/128
                        *[Local/0] 01:02:01
                        Local via fe-1/2/0.0
fe80::2a0:a514:0:114c/128
                        *[Local/0] 01:02:01
                        Local via fe-1/2/1.0
ff02::5/128             *[OSPF3/10] 3d 03:01:25, metric 1
                        MultiRecv

```

```

user@10> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.10/32      *[Direct/0] 01:01:59
                    > via lo0.0

inet6.0: 11 destinations, 14 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                *[OSPF3/150] 01:01:35, metric 11, tag 0
                    > via fe-1/2/0.0
2001:db8:3030::1/128  *[OSPF3/150] 01:01:35, metric 0, tag 0
                    > via fe-1/2/1.0
2001:db8:4040::1/128  *[OSPF3/150] 01:01:35, metric 0, tag 0
                    > via fe-1/2/1.0
2001:db8:9009:5::/64  *[Direct/0] 01:01:50
                    > via fe-1/2/0.0
                    [OSPF3/10] 01:01:50, metric 1
                    > via fe-1/2/0.0
2001:db8:9009:5::2/128  *[Local/0] 01:01:50

```

```

                Local via fe-1/2/0.0
2001:db8:9009:9::/64      *[Direct/0] 01:01:50
                        > via fe-1/2/1.0
                        [OSPF3/10] 01:01:40, metric 1
                        > via fe-1/2/1.0
2001:db8:9009:9::2/128   *[Local/0] 01:01:50
                        Local via fe-1/2/1.0
fe80::/64                *[Direct/0] 01:01:50
                        > via fe-1/2/0.0
                        [Direct/0] 01:01:50
                        > via fe-1/2/1.0
fe80::2a0:a514:0:c4c/128
                        *[Local/0] 01:01:50
                        Local via fe-1/2/0.0
fe80::2a0:a514:0:124c/128
                        *[Local/0] 01:01:50
                        Local via fe-1/2/1.0
ff02::5/128              *[OSPF3/10] 01:02:16, metric 1
                        MultiRecv

```

```

user@3> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.3.3.3/32              *[Direct/0] 3d 03:03:10
                        > via lo0.0

inet6.0: 15 destinations, 18 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:1010::1/128     *[OSPF3/150] 01:04:21, metric 0, tag 0
                        > via fe-1/2/0.0
2001:db8:2020::1/128     *[OSPF3/150] 01:04:21, metric 0, tag 0
                        > via fe-1/2/0.0
2001:db8:3030::1/128     *[OSPF3/150] 01:03:57, metric 0, tag 0
                        > via fe-1/2/1.0
2001:db8:4040::1/128     *[OSPF3/150] 01:03:57, metric 0, tag 0
                        > via fe-1/2/1.0
2001:db8:9009:1::/64     *[OSPF3/10] 3d 03:02:06, metric 2
                        > via fe-1/2/0.0
2001:db8:9009:3::/64     *[Direct/0] 3d 03:02:55

```

```

> via fe-1/2/0.0
[OSPF3/10] 3d 03:02:54, metric 1
> via fe-1/2/0.0
2001:db8:9009:3::2/128    *[Local/0] 3d 03:02:55
    Local via fe-1/2/0.02001:db8:9009:5::/64    *[Direct/0] 01:04:09
> via fe-1/2/1.0
[OSPF3/10] 01:04:09, metric 1
> via fe-1/2/1.0
2001:db8:9009:5::1/128    *[Local/0] 3d 03:02:54
    Local via fe-1/2/1.0
2001:db8:9009:6::/64      *[OSPF3/10] 3d 02:19:14, metric 3
> via fe-1/2/0.0
2001:db8:9009:9::/64      *[OSPF3/10] 01:04:02, metric 2
> via fe-1/2/1.0
fe80::/64                *[Direct/0] 3d 03:02:55
> via fe-1/2/0.0
[Direct/0] 01:04:09
> via fe-1/2/1.0
fe80::2a0:a514:0:84c/128
    *[Local/0] 3d 03:02:55
    Local via fe-1/2/0.0
fe80::2a0:a514:0:b4c/128
    *[Local/0] 3d 03:02:54
    Local via fe-1/2/1.0
ff02::5/128              *[OSPF3/10] 3d 03:03:50, metric 1
    MultiRecv

```

Meaning

On Device 7, the default route has been learned because of the `default-metric` statement on the ABR, Device 3. Otherwise, the only OSPFv3 routes in Device 7's routing table are those local to area 9 and the OSPFv3 multicast address `ff02::5/128` for all SPF link-state routers, also known as AllSPFRouters.

Device 10 has the default route injected by Device 3 and also the OSPF external routes injected by Device 7.

Neither Device 7 nor Device 10 has the external customer routes that were injected into OSPFv3 by Device 5.

On Device 3, all of the OSPFv3 routes have been learned, including the external customer routes, `2001:db8:1010::1/128` and `2001:db8:2020::1/128`.

Verifying the Type of LSAs

Purpose

Verify the type of LSAs that are in the area.

Action

From operational mode on Device 7, enter the `show ospf3 database nssa detail` command.

```
user@7> show ospf3 database nssa detail
Area 0.0.0.9
Type      ID           Adv Rtr      Seq          Age  Cksum  Len
NSSA      0.0.0.1      10.3.3.3     0x8000002a  1462 0xf406  28
  Prefix ::/0
  Prefix-options 0x0, Metric 10, Type 1,
NSSA      *0.0.0.1     10.7.7.7     0x80000003  1625 0x88df  60
  Prefix 2001:db8:3030::1/128
  Prefix-options 0x8, Metric 0, Type 2,
  Fwd addr 2001:db8:9009:9::1,
NSSA      *0.0.0.2     10.7.7.7     0x80000003  1025 0xef57  60
  Prefix 2001:db8:4040::1/128
  Prefix-options 0x8, Metric 0, Type 2,
  Fwd addr 2001:db8:9009:9::1,
```

Meaning

On Device 7, the NSSA LSAs are the type 1 external default route, learned from Device 3, and the type 2 external static routes to the Customer 1 network.

Understanding Not-So-Stubby Areas Filtering

You might have a situation when exporting Type 7 LSAs into a not-so-stubby area (NSSA) is unnecessary. When an autonomous system boundary router (ASBR) is also an area border router (ABR) with an NSSA attached, Type 7 LSAs are exported into the NSSA by default.

Also, when the ASBR (also an ABR) is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. During route redistribution, this routing device generates both Type 5 LSAs and Type 7 LSAs. Hence, to avoid the same route getting redistributed twice (from Type 5 LSAs and Type 7 LSAs).

7 LSAs), you can disable exporting Type 7 LSAs into the NSSA by including the `no-nssa-abr` statement on the routing device.

Example: Configuring OSPFv3 Not-So-Stubby Areas with Filtering

IN THIS SECTION

- [Requirements | 141](#)
- [Overview | 141](#)
- [Configuration | 142](#)
- [Verification | 148](#)

This example shows how to configure an OSPFv3 not-so-stubby area (NSSA) when there is no need to inject external routes into the NSSA as Type 7 link-state advertisements (LSAs).

Requirements


No special configuration beyond device initialization is required before configuring this example.

Overview

When an autonomous system border router (ASBR) is also an NSSA area border router (ABR), the routing device generates Type 5 as well as Type 7 LSAs. You can prevent the router from creating Type 7 LSAs for the NSSA with the `no-nssa-abr` statement.

In this example, Device 5 and Device 3 are in customer networks. Device 4 and Device 2 are both injecting the customer routes into OSPFv3. Area 1 is an NSSA. Because Device 4 is both an NSSA ABR and an ASBR, it generates both type 7 and type 5 LSAs and injects type 7 LSAs into area 1 and type 5 LSAs into area 0. To stop type 7 LSAs from being injected into area 1, the `no-nssa-abr` statement is included in the Device 4 configuration.

Figure 14: OSPFv3 Network Topology with an NSSA ABR That Is Also an ASBR



"CLI Quick Configuration" on page 142 shows the configuration for all of the devices in Figure 14 on page 142. The section "No Link Title" on page 144 describes the steps on Device 4.

Configuration

IN THIS SECTION

- Procedure | 142

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device 1

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:6::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:5::1/64
set interfaces lo0 unit 0 family inet address 0.1.1.1/32
set protocols ospf3 area 0.0.0.1 nssa
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive

```

Device 2

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:5::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:4::1/64
set interfaces lo0 unit 0 family inet address 10.2.2.2/32
set protocols ospf3 export static2-to-ospf
set protocols ospf3 area 0.0.0.1 nssa
set protocols ospf3 area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement static2-to-ospf term 1 from protocol static
set policy-options policy-statement static2-to-ospf term 1 then accept
set routing-options rib inet6.0 static route 3030::1/128 next-hop 2001:db8:9009:4::2
set routing-options rib inet6.0 static route 4040::1/128 next-hop 2001:db8:9009:4::2

```

Device 3

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:4::2/64
set interfaces lo0 unit 0 family inet address 10.3.3.3/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:3030::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:4040::1/128
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:9009:4::1

```

Device 4

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::2/64
set interfaces fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:6::1/64
set interfaces fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:3::1/64
set interfaces lo0 unit 0 family inet address 10.4.4.4/32
set protocols ospf3 export static-to-ospf
set protocols ospf3 no-nssa-abr

```

```

set protocols ospf3 area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 nssa default-lsa default-metric 10
set protocols ospf3 area 0.0.0.1 nssa default-lsa metric-type 1
set protocols ospf3 area 0.0.0.1 nssa default-lsa type-7
set protocols ospf3 area 0.0.0.1 nssa no-summaries
set protocols ospf3 area 0.0.0.1 interface fe-1/2/1.0
set policy-options policy-statement static-to-ospf term 1 from protocol static
set policy-options policy-statement static-to-ospf term 1 then accept
set routing-options rib inet6.0 static route 2001:db8:1010::1/128 next-hop 2001:db8:9009:1::1
set routing-options rib inet6.0 static route 2001:db8:2020::1/128 next-hop 2001:db8:9009:1::1

```

Device 5

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::1/64
set interfaces lo0 unit 0 family inet address 10.5.5.5/32
set interfaces lo0 unit 0 family inet6 address 2001:db8:1010::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:2020::1/128
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:9009:1::2

```

Device 6

```

set interfaces fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:3::2/64
set interfaces lo0 unit 0 family inet address 10.6.6.6/32
set protocols ospf3 area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see “Using the CLI Editor in Configuration Mode” in the [CLI User Guide](#).

To configure Device 4:

1. Configure the interfaces.

```

[edit interfaces]
user@4# set fe-1/2/0 unit 0 family inet6 address 2001:db8:9009:1::2/64
user@4# set fe-1/2/1 unit 0 family inet6 address 2001:db8:9009:6::1/64

```

```
user@4# set fe-1/2/2 unit 0 family inet6 address 2001:db8:9009:3::1/64
user@4# set lo0 unit 0 family inet address 10.4.4.4/32
```

2. Enable OSPFv3 on the interfaces that are in area 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@4# set interface fe-1/2/2.0
user@4# set interface lo0.0 passive
```

3. Enable OSPFv3 on the interface that is in area 1.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set interface fe-1/2/1.0
```

4. Configure an OSPFv3 NSSA.

The nssa statement is required on all routing devices in the area.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set nssa
```

5. On the ABR, inject a default route into the area.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set nssa default-lsa default-metric 10
```

6. (Optional) On the ABR, specify the external metric type for the default route.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set nssa default-lsa metric-type 1
```

7. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set nssa default-lsa type-7
```

8. On the ABR, restrict summary LSAs from entering the area.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set nssa no-summaries
```

9. Disable exporting Type 7 LSAs into the NSSA.

This setting is useful if you have an AS boundary router that is also an ABR with an NSSA area attached.

```
[edit protocols ospf3]
user@4# set no-nssa-abr
```

10. Configure static routes to the customer network.

```
[edit routing-options rib inet6.0 static]
user@4# set route 2001:db8:1010::1/128 next-hop 2001:db8:9009:1::1
user@4# set route 2001:db8:2020::1/128 next-hop 2001:db8:9009:1::1
```

11. Configure a policy to inject the static routes into OSPFv3.

```
[edit policy-options policy-statement static-to-ospf term 1]
user@4# set from protocol static
user@4# set then accept
```

12. Apply the policy to OSPFv3.

```
[edit protocols ospf3]
user@4# set export static-to-ospf
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device 4

```
user@4# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet6 {
      address 2001:db8:9009:1::2/64;
    }
  }
  unit 0 {
    family inet6 {
      address 2001:db8:9009:6::1/64;
    }
  }
  unit 0 {
    family inet6 {
      address 2001:db8:9009:3::1/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.4.4.4/32;
    }
  }
}
```

```
user@4# show protocols
ospf3 {
  export static-to-ospf;
  no-nssa-abr;
  area 0.0.0.0 {
    interface fe-1/2/2.0;
    interface lo0.0 {
      passive;
    }
  }
  area 0.0.0.1 {
    nssa {
      default-lsa {
```

```

        default-metric 10;
        metric-type 1;
        type-7;
    }
    no-summaries;
}
interface fe-1/2/1.0;
}
}

```

```

user@4# show policy-options
policy-statement static-to-ospf {
  term 1 {
    from protocol static;
    then accept;
  }
}

```

```

user@4# show routing-options
rib inet6.0 {
  static {
    route 2001:db8:1010::1/128 next-hop 2001:db8:9009:1::1;
    route 2001:db8:2020::1/128 next-hop 2001:db8:9009:1::1;
  }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Routes in the OSPFv3 Stub Area | 149](#)
- [Verifying the Type of LSAs | 151](#)

Confirm that the configuration is working properly.

Verifying the Routes in the OSPFv3 Stub Area

Purpose

Make sure that the expected routes are present in the routing tables.

Action

From operational mode on Device 1 and Device 6, enter the `show route` command.

```

user@1> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.1.1.1/32      *[Direct/0] 03:25:44
                > via lo0.0

inet6.0: 11 destinations, 14 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0           *[OSPF3/150] 01:52:58, metric 11, tag 0
                > via fe-1/2/0.0
2001:db8:3030::1/128  *[OSPF3/150] 02:44:02, metric 0, tag 0
                > via fe-1/2/1.0
2001:db8:4040::1/128  *[OSPF3/150] 02:44:02, metric 0, tag 0
                > via fe-1/2/1.0
2001:db8:9009:5::/64  *[Direct/0] 03:25:34
                > via fe-1/2/1.0
                [OSPF3/10] 03:25:24, metric 1
                > via fe-1/2/1.0
2001:db8:9009:5::1/128  *[Local/0] 03:25:34
                Local via fe-1/2/1.0
2001:db8:9009:6::/64  *[Direct/0] 03:25:34
                > via fe-1/2/0.0
                [OSPF3/10] 03:25:34, metric 1
                > via fe-1/2/0.0
2001:db8:9009:6::2/128  *[Local/0] 03:25:34
                Local via fe-1/2/0.0
fe80::/64      *[Direct/0] 03:25:34
                > via fe-1/2/0.0
                [Direct/0] 03:25:34
                > via fe-1/2/1.0

```

```

fe80::2a0:a514:0:44c/128
    *[Local/0] 03:25:34
    Local via fe-1/2/0.0
fe80::2a0:a514:0:74c/128
    *[Local/0] 03:25:34
    Local via fe-1/2/1.0
ff02::5/128
    *[OSPF3/10] 03:27:00, metric 1
    MultiRecv

```

```

user@6> show route
inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.6.6.6/32
    *[Direct/0] 03:26:57
    > via lo0.0

inet6.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:1010::1/128
    *[OSPF3/150] 03:16:59, metric 0, tag 0
    > via fe-1/2/0.0
2001:db8:2020::1/128
    *[OSPF3/150] 03:16:59, metric 0, tag 0
    > via fe-1/2/0.0
2001:db8:3030::1/128
    *[OSPF3/150] 02:44:34, metric 0, tag 0
    > via fe-1/2/0.0
2001:db8:4040::1/128
    *[OSPF3/150] 02:44:34, metric 0, tag 0
    > via fe-1/2/0.0
2001:db8:9009:3::/64
    *[Direct/0] 03:26:29
    > via fe-1/2/0.0
    [OSPF3/10] 03:26:29, metric 1
    > via fe-1/2/0.0
2001:db8:9009:3::2/128
    *[Local/0] 03:26:29
    Local via fe-1/2/0.0
2001:db8:9009:5::/64
    *[OSPF3/10] 02:44:34, metric 3
    > via fe-1/2/0.0
2001:db8:9009:6::/64
    *[OSPF3/10] 03:16:59, metric 2
    > via fe-1/2/0.0
fe80::/64
    *[Direct/0] 03:26:29
    > via fe-1/2/0.0
fe80::2a0:a514:0:64c/128
    *[Local/0] 03:26:29

```

```

Local via fe-1/2/0.0
ff02::5/128      *[OSPF3/10] 03:27:37, metric 1
                  MultiRecv

```

Meaning

On Device 1, the default route (::/0) has been learned because of the default-metric statement on the ABR, Device 4. The customer routes 2001:db8:3030::1 and 2001:db8:4040::1 have been learned from Device 2. The 2001:db8:1010::1 and 2001:db8:2020::1 routes have been suppressed. They are not needed because the default route can be used instead.

On Device 6 in area 0, all of the customer routes have been learned.

Verifying the Type of LSAs

Purpose

Verify the type of LSAs that are in the area.

Action

From operational mode on Device 1, enter the show ospf3 database nssa detail command.

```

user@4> show ospf3 database nssa detail
Area 0.0.0.1
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
  NSSA      0.0.0.1     10.2.2.2     0x80000004   2063 0xceaf 60
    Prefix 3030::1/128
    Prefix-options 0x8, Metric 0, Type 2,
    Fwd addr 2001:db8:9009:5::2,
  NSSA      0.0.0.2     10.2.2.2     0x80000004   1463 0x3627 60
    Prefix 4040::1/128
    Prefix-options 0x8, Metric 0, Type 2,
    Fwd addr 2001:db8:9009:5::2,
  NSSA      *0.0.0.1    10.4.4.4     0x80000003   35   0x25f8 28
    Prefix ::/0
    Prefix-options 0x0, Metric 10, Type 1,

```


Meaning

Device 4 is not sending Type 7 (NSSA) LSAs for customer routes 2001:db8:1010::1/128 and 2001:db8:2020::1/128. If you were to delete or deactivate the `no-nssa-abr` statement and then rerun the `show ospf3 database nssa detail` command, you would see that Device 4 is sending Type 7 LSAs for 2001:db8:1010::1/128 and 2001:db8:2020::1/128.

Understanding OSPF Virtual Links for Noncontiguous Areas

OSPF requires that all areas in an autonomous system (AS) must be physically connected to the backbone area (area 0). In large networks with many areas, in which direct connectivity between all areas and the backbone area is physically difficult or impossible, you can configure virtual links to connect noncontiguous areas. Virtual links use a transit area that contains two or more area border routers (ABRs) to pass network traffic from one adjacent area to another. The transit area must have full routing information and it cannot be a stub area. For example, [Figure 15 on page 152](#) shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 15: OSPF Topology with a Virtual Link



In the topology shown in [Figure 15 on page 152](#), a virtual link is established between area 0.0.0.3 and the backbone area through area 0.0.0.2. The virtual link transits area 0.0.0.2. All outbound traffic destined for other areas is routed through area 0.0.0.2 to the backbone area and then to the appropriate ABR. All inbound traffic destined for area 0.0.0.3 is routed to the backbone area and then through area 0.0.0.2.

Example: Configuring OSPF Virtual Links to Connect Noncontiguous Areas

IN THIS SECTION

- [Requirements | 153](#)
- [Overview | 153](#)
- [Configuration | 154](#)
- [Verification | 158](#)

This example shows how to configure an OSPF virtual link to connect noncontiguous areas.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on [page 62](#).
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on [page 65](#).

Overview

IN THIS SECTION

- [Topology | 154](#)

If any routing device on the backbone is not physically connected to the backbone, you must establish a virtual connection between that routing device and the backbone to connect the noncontiguous areas.


To configure an OSPF virtual link through an area, you specify the router ID (IP address) of the routing devices at each end of the virtual link. These routing devices must be area border routers (ABRs), with one that is physically connected to the backbone. You cannot configure virtual links through stub areas. You must also specify the number of the area through which the virtual link transits (also known as the

transit area). You apply these settings to the backbone area (defined by the area 0.0.0.0) configuration on the ABRs that are part of the virtual link.

In this example, Device R1 and Device R2 are the routing devices at each end of the virtual link, with Device R1 physically connected to the backbone, as shown in [Figure 16 on page 154](#). You configure the following virtual link settings:

- **neighbor-id**—Specifies the IP address of the routing device at the other end of the virtual link. In this example, Device R1 has a router ID of 192.0.2.5, and Device R2 has a router ID of 192.0.2.3.
- **transit-area**—Specifies the area identifier through which the virtual link transits. In this example, area 0.0.0.3 is not connected to the backbone, so you configure a virtual link session between area 0.0.0.3 and the backbone area through area 0.0.0.2. Area 0.0.0.2 is the transit area.

Figure 16: OSPF Virtual Link



Topology

Configuration

IN THIS SECTION

- CLI Quick Configuration | 155
- Procedure | 155
- Results | 157

CLI Quick Configuration

- To quickly configure an OSPF virtual link on the local routing device (Device R1), copy the following commands and paste them into the CLI.



NOTE: You must configure both routing devices that are part of the virtual link and specify the applicable neighbor ID on each routing device.

```
[edit]
set routing-options router-id 192.0.2.5
set protocols ospf area 0.0.0.0 virtual-link neighbor-id 192.0.2.3 transit-area 0.0.0.2
```

- To quickly configure an OSPF virtual link on the remote routing device (Device R2), copy the following commands and paste them into the CLI.

```
[edit]
set routing-options router-id 192.0.2.3
set protocols ospf area 0.0.0.0 virtual-link neighbor-id 192.0.2.5 transit-area 0.0.0.2
```

Procedure

Step-by-Step Procedure

To configure an OSPF virtual link on the local routing device (Device R1):

1. Configure the router ID.

```
[edit]
user@R1# set routing-options router-id 192.0.2.5
```

2. Enter OSPF configuration mode and specify OSPF area 0.0.0.0.



NOTE: For an OSPFv3 virtual link, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# edit protocols ospf area 0.0.0.0
```

3. Configure an OSPF virtual link and specify the transit area 0.0.0.2.
This routing device must be an ABR that is physically connected to the backbone.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set virtual-link neighbor-id 192.0.2.3 transit-area 0.0.0.2
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@R1# commit
```

Step-by-Step Procedure

To configure an OSPF virtual link on the remote ABR (Device R2, the routing device at the other end of the link):

1. Configure the router ID.

```
[edit]
user@R2# set routing-options router-id 192.0.2.3
```

2. Enter OSPF configuration mode and specify OSPF area 0.0.0.0.



NOTE: For an OSPFv3 virtual link, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R2# edit protocols ospf area 0.0.0.0
```

3. Configure an OSPF virtual link on the remote ABR and specify the transit area 0.0.0.2.
This routing device is not physically connected to the backbone.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set virtual-link neighbor-id 192.0.2.5 transit-area 0.0.0.2
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@R2# commit
```

Results

Confirm your configuration by entering the show **routing-options** and the show protocols ospf commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on the local routing device (Device R1):

```
user@R1#: show routing-options
router-id 192.0.2.5;
```

```
user@R1# show protocols ospf
area 0.0.0.0 {
    virtual-link neighbor-id 192.0.2.3 transit-area 0.0.0.2;
}
```

Configuration on the remote ABR (Device R2):

```
user@R2#: show routing-options
router-id 192.0.2.3;
```

```
user@R2# show protocols ospf
area 0.0.0.0 {
    virtual-link neighbor-id 192.0.2.5 transit-area 0.0.0.2;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying Entries in the Link-State Database | 158](#)
- [Verifying OSPF Interface Status and Configuration | 158](#)

Confirm that the configuration is working properly.

Verifying Entries in the Link-State Database

Purpose

Verify that the entries in the OSPFv2 or OSPFv3 link-state database display. The Router field in the OSPFv2 output displays LSA information, including the type of link. If configured as a virtual link, the Type is Virtual. For each router link, the Type field in the OSPFv3 output displays the type of interface. If configured as a virtual link, the Type is Virtual.

Action

From operational mode, enter the `show ospf database detail` command for OSPFv2, and enter the `show ospf3 database detail` command for OSPFv3.

Verifying OSPF Interface Status and Configuration

Purpose

Verify that the OSPFv2 or OSPFv3 interface is configured and status displays. The Type field displays the type of interface. If the interface is configured as part of a virtual link, the Type is Virtual.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Example: Configuring OSPFv3 Virtual Links

IN THIS SECTION

- [Requirements | 159](#)
- [Overview | 159](#)
- [Configuration | 160](#)
- [Verification | 174](#)

This example shows how to configure OSPF version 3 (OSPFv3) with some areas that do not have a direct adjacency to the backbone area (area 0). When an area lacks an adjacency with area 0, a virtual link is required to connect to the backbone through a non-backbone area. The area through which you configure the virtual link, known as a transit area, must have full routing information. The transit area cannot be a stub area.


Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

[Figure 17 on page 160](#) shows the topology used in this example.

Figure 17: OSPFv3 with Virtual Links



Device 0, Device 1, Device 2, and Device 3 are connected to the OSPFv3 backbone Area 0. Device 2, Device 3, and Device 4 connect to each other across Area 1, and Area 2 is located between Device 4 and Device 5. Because Device 5 does not have a direct adjacency to Area 0, a virtual link is required across Area 1 between Device 3 and Device 4. Similarly, because Device 0 and Device 1 have two separate Area 0 backbone sections, you need to configure a second virtual link across Area 1 between Device 2 and Device 3.

Configuration

IN THIS SECTION

- Procedure | 161

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device 0

```
set logical-systems 0 interfaces so-0/3/2 unit 0 family inet6 address 9009:1::1/64
set logical-systems 0 interfaces lo0 unit 0 family inet address 192.168.0.1/32
set logical-systems 0 interfaces lo0 unit 0 family inet6 address feee::10:255:71:4/128
set logical-systems 0 protocols ospf3 area 0.0.0.0 interface so-0/3/2.0
set logical-systems 0 protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set logical-systems 0 routing-options router-id 192.168.0.1
```

Device 1

```
set logical-systems 1 interfaces at-2/0/0 atm-options vpi 0
set logical-systems 1 interfaces at-2/0/0 unit 0 family inet6 address 9009:2::1/64
set logical-systems 1 interfaces at-2/0/0 unit 0 vci 0.77
set logical-systems 1 interfaces lo0 unit 0 family inet address 192.168.1.1/32
set logical-systems 1 interfaces lo0 unit 0 family inet6 address feee::10:255:71:1/128
set logical-systems 1 protocols ospf3 area 0.0.0.0 interface at-2/0/0.0
set logical-systems 1 protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set logical-systems 1 routing-options router-id 192.168.1.1
```

Device 2

```
set logical-systems 2 interfaces so-0/2/0 unit 0 family inet6 address 9009:3::1/64
set logical-systems 2 interfaces fe-1/1/0 unit 0 family inet6 address 9009:4::1/64
set logical-systems 2 interfaces at-0/3/1 atm-options vpi 0 maximum-vc 1200
set logical-systems 2 interfaces at-0/3/1 unit 0 family inet6 address 9009:2::2/64
set logical-systems 2 interfaces at-0/3/1 unit 0 vci 0.77
set logical-systems 2 interfaces lo0 unit 0 family inet address 192.168.2.1/32
set logical-systems 2 interfaces lo0 unit 0 family inet6 address feee::10:255:71:11/128
set logical-systems 2 protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.3.1 transit-
area 0.0.0.1
set logical-systems 2 protocols ospf3 area 0.0.0.0 interface at-0/3/1.0
set logical-systems 2 protocols ospf3 area 0.0.0.1 interface fe-1/1/0.0
```

```

set logical-systems 2 protocols ospf3 area 0.0.0.1 interface so-0/2/0.0
set logical-systems 2 protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set logical-systems 2 routing-options router-id 192.168.2.1

```

Device 3

```

set logical-systems 3 interfaces so-0/3/2 unit 0 family inet6 address 9009:1::2/64
set logical-systems 3 interfaces t1-0/2/1 unit 0 family inet6 address 9009:5::1/64
set logical-systems 3 interfaces so-0/3/0 unit 0 family inet6 address 9009:3::2/64
set logical-systems 3 interfaces lo0 unit 0 family inet address 192.168.3.1/32
set logical-systems 3 interfaces lo0 unit 0 family inet6 address feee::10:255:71:3/128
set logical-systems 3 protocols ospf3 area 0.0.0.1 interface so-0/3/0.0
set logical-systems 3 protocols ospf3 area 0.0.0.1 interface t1-0/2/1.0
set logical-systems 3 protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set logical-systems 3 protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.2.1 transit-
area 0.0.0.1
set logical-systems 3 protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.4.1 transit-
area 0.0.0.1
set logical-systems 3 protocols ospf3 area 0.0.0.0 interface so-0/3/2.0
set logical-systems 3 routing-options router-id 192.168.3.1

```

Device 4

```

set logical-systems 4 interfaces t1-0/2/1 unit 0 family inet6 address 9009:5::2/64
set logical-systems 4 interfaces fe-0/0/0 unit 0 family inet6 address 9009:6::1/64
set logical-systems 4 interfaces fe-1/1/0 unit 0 family inet6 address 9009:4::2/64
set logical-systems 4 interfaces lo0 unit 0 family inet address 192.168.4.1/32
set logical-systems 4 interfaces lo0 unit 0 family inet6 address feee::10:255:71:5/128
set logical-systems 4 protocols ospf3 area 0.0.0.1 interface fe-1/1/0.0
set logical-systems 4 protocols ospf3 area 0.0.0.1 interface t1-0/2/1.0
set logical-systems 4 protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set logical-systems 4 protocols ospf3 area 0.0.0.2 interface fe-0/0/0.0
set logical-systems 4 protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.3.1 transit-
area 0.0.0.1
set logical-systems 4 routing-options router-id 192.168.4.1

```

Device 5

```

set logical-systems 5 interfaces fe-0/0/0 unit 0 family inet6 address 9009:6::2/64
set logical-systems 5 interfaces lo0 unit 0 family inet address 192.168.5.1/32
set logical-systems 5 interfaces lo0 unit 0 family inet6 address feee::10:255:71:6/128

```

```

set logical-systems 5 protocols ospf3 area 0.0.0.2 interface fe-0/0/0.0
set logical-systems 5 protocols ospf3 area 0.0.0.2 interface lo0.0 passive
set logical-systems 5 routing-options router-id 192.168.5.1

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 0:

1. Configure the interfaces.

```

[edit interfaces]
user@0# set so-0/3/2 unit 0 family inet6 address 9009::1/64
user@0# set lo0 unit 0 family inet address 192.168.0.1/32
user@0# set lo0 unit 0 family inet6 address feee::10:255:71:4/128

```

2. Add the interfaces into Area 0 of the OSPFv3 process.

```

[edit protocols ospf3 area 0.0.0.0]
user@0# set interface so-0/3/2.0
user@0# set interface lo0.0 passive

```

3. Configure the router ID.

```

[edit routing-options]
user@0# set router-id 192.168.0.1

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 1:

1. Configure the interfaces.

```
[edit interfaces]
user@1# set at-2/0/0 atm-options vpi 0
user@1# set at-2/0/0 unit 0 family inet6 address 9009:2::1/64
user@1# set at-2/0/0 unit 0 vci 0.77
user@1# set lo0 unit 0 family inet address 192.168.1.1/32
user@1# set lo0 unit 0 family inet6 address feee::10:255:71:1/128
```

2. Add the interfaces into Area 0 of the OSPFv3 process.

```
[edit protocols ospf3 area 0.0.0.0]
user@1# set interface at-2/0/0.0
user@1# set interface lo0.0 passive
```

3. Configure the router ID.

```
[edit routing-options]
user@1# set router-id 192.168.1.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 2:

1. Configure the interfaces.

```
[edit interfaces]
user@2# set so-0/2/0 unit 0 family inet6 address 9009:3::1/64
user@2# set fe-1/1/0 unit 0 family inet6 address 9009:4::1/64
user@2# set at-0/3/1 atm-options vpi 0 maximum-vcs 1200
user@2# set at-0/3/1 unit 0 family inet6 address 9009:2::2/64
user@2# set at-0/3/1 unit 0 vci 0.77
user@2# set lo0 unit 0 family inet address 192.168.2.1/32
user@2# set lo0 unit 0 family inet6 address feee::10:255:71:11/128
```

2. Add the interfaces connected to Device 1, Device 3, and Device 4 into the OSPFv3 process.

```
[edit protocols ospf3 area 0.0.0.0]
user@2# set interface at-0/3/1.0
[edit protocols ospf3 area 0.0.0.1]
user@2# set interface fe-1/1/0.0
user@2# set interface so-0/2/0.0
user@2# set interface lo0.0 passive
```

3. Configure the virtual link to Device 3 through Area 1 so that Device 1 can access the discontinuous portion of the OSPF backbone found on Device 0.

```
[edit protocols ospf3 area 0.0.0.0]
user@2# set virtual-link neighbor-id 192.168.3.1 transit-area 0.0.0.1
```

4. Configure the router ID.

```
[edit routing-options]
user@2# set router-id 192.168.2.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 3:

1. Configure the interfaces.

```
[edit interfaces]
user@3# set so-0/3/2 unit 0 family inet6 address 9009:1::2/64
user@3# set t1-0/2/1 unit 0 family inet6 address 9009:5::1/64
user@3# set so-0/3/0 unit 0 family inet6 address 9009:3::2/64
user@3# set lo0 unit 0 family inet address 192.168.3.1/32
user@3# set lo0 unit 0 family inet6 address feee::10:255:71:3/128
```

2. For the OSPFv3 process on Device 3, configure the interfaces connected to Device 2 and Device 4 into Area 1 and the interface connected to Device 0 into Area 0.

```
[edit protocols ospf3 area 0.0.0.1]
user@3# set interface so-0/3/0.0
user@3# set interface t1-0/2/1.0
user@3# set interface lo0.0 passive
[edit protocols ospf3 area 0.0.0.0]
user@3# set interface so-0/3/2.0
```

3. Configure two virtual links through Area 1—one connecting to Device 2 and the second connecting to Device 4.

The virtual links allow Device 5 to access the OSPF backbone, and connect the discontinuous sections of Area 0 located at Device 0 and Device 1.

```
[edit protocols ospf3 area 0.0.0.0]
user@3# set virtual-link neighbor-id 192.168.2.1 transit-area 0.0.0.1
user@3# set virtual-link neighbor-id 192.168.4.1 transit-area 0.0.0.1
```

4. Configure the router ID.

```
[edit routing-options]
user@3# set router-id 192.168.3.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 4:

1. Configure the interfaces.

```
[edit interfaces]
user@4# set t1-0/2/1 unit 0 family inet6 address 9009:5::2/64
user@4# set fe-0/0/0 unit 0 family inet6 address 9009:6::1/64
user@4# set fe-1/1/0 unit 0 family inet6 address 9009:4::2/64
```

```
user@4# set lo0 unit 0 family inet address 192.168.4.1/32
user@4# set lo0 unit 0 family inet6 address feee::10:255:71:5/128
```

2. On Device 4, add the connected interfaces into the OSPFv3 process.

```
[edit protocols ospf3 area 0.0.0.1]
user@4# set interface fe-1/1/0.0
user@4# set interface t1-0/2/1.0
user@4# set interface lo0.0 passive
[edit protocols ospf3 area 0.0.0.2]
user@4# set interface fe-0/0/0.0
```

3. Configure the virtual link to Device 3 through Area 1 so that Device 5 can access the OSPF backbone.

```
[edit protocols ospf3 area 0.0.0.0]
user@4# set virtual-link neighbor-id 192.168.3.1 transit-area 0.0.0.1
```

4. Configure the router ID.

```
[edit routing-options]
user@4# set router-id 192.168.4.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device 5:

1. Configure the interfaces.

```
[edit interfaces]
user@5# set fe-0/0/0 unit 0 family inet6 address 9009:6::2/64
user@5# set lo0 unit 0 family inet address 192.168.5.1/32
user@5# set lo0 unit 0 family inet6 address feee::10:255:71:6/128
```

2. Add the interfaces into the OSPFv3 process.

```
[edit protocols ospf3 area 0.0.0.2]
user@5# set interface fe-0/0/0.0
user@5# set interface lo0.0 passive
```

3. Configure the router ID.

```
[edit routing-options]
user@5# set router-id 192.168.5.1
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device 0

```
user@0# show interfaces
so-0/3/2 {
  unit 0 {
    family inet6 {
      address 9009:1::1/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family inet6 {
      address feee::10:255:71:4/128;
    }
  }
}
user@0# show protocols
ospf3 {
  area 0.0.0.0 {
```

```

    interface so-0/3/2.0;
    interface lo0.0 {
        passive;
    }
}
user@0# show routing-options
router-id 192.168.0.1;

```

Device 1

```

user@1# show interfaces
at-2/0/0 {
    atm-options {
        vpi 0;
    }
    unit 0 {
        family inet6 {
            address 9009:2::1/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.1.1/32;
        }
        family inet6 {
            address feee::10:255:71:1/128;
        }
    }
}
user@1# show protocols
ospf3 {
    area 0.0.0.0 {
        interface at-2/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
}

```

```
user@1# show routing-options
router-id 192.168.1.1;
```

Device 2

```
user@2# show interfaces
so-0/2/0 {
  unit 0 {
    family inet6 {
      address 9009:3::1/64;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet6 {
      address 9009:4::1/64;
    }
  }
}
at-0/3/1 {
  atm-options {
    vpi 0 {
      maximum-vcs 1200;
    }
  }
  unit 0 {
    vci 0.77;
    family inet6 {
      address 9009:2::2/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.2.1/32;
    }
    family inet6 {
      address feee::10:255:71:11/128;
    }
  }
}
```

```
}
user@2# show protocols
ospf3 {
  area 0.0.0.0 {
    virtual-link neighbor-id 192.168.3.1 transit-area 0.0.0.1;
    interface at-0/3/1.0;
  }
  area 0.0.0.1 {
    interface fe-1/1/0.0;
    interface so-0/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
user@2# show routing-options
router-id 192.168.2.1;
```

Device 3

```
user@3# show interfaces
so-0/3/2 {
  unit 0 {
    family inet6 {
      address 9009:1::2/64;
    }
  }
}
t1-0/2/1 {
  unit 0 {
    family inet6 {
      address 9009:5::1/64;
    }
  }
}
so-0/3/0 {
  unit 0 {
    family inet6 {
      address 9009:3::2/64;
    }
  }
}
```

```
lo0 {
  unit 0 {
    family inet {
      address 192.168.3.1/32;
    }
    family inet6 {
      address feee::10:255:71:3/128;
    }
  }
}
user@3# show protocols
ospf3 {
  area 0.0.0.1 {
    interface so-0/3/0.0;
    interface t1-0/2/1.0;
    interface lo0.0 {
      passive;
    }
  }
  area 0.0.0.0 {
    virtual-link neighbor-id 192.168.2.1 transit-area 0.0.0.1;
    virtual-link neighbor-id 192.168.4.1 transit-area 0.0.0.1;
    interface so-0/3/2.0;
  }
}
user@3# show routing-options
router-id 192.168.3.1;
```

Device 4

```
user@4# show interfaces
t1-0/2/1 {
  unit 0 {
    family inet6 {
      address 9009:5::2/64;
    }
  }
}
fe-0/0/0 {
  unit 0 {
    family inet6 {
      address 9009:6::1/64;
```

```
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet6 {
      address 9009:4::2/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.4.1/32;
    }
    family inet6 {
      address feee::10:255:71:5/128;
    }
  }
}
}
user@4# show protocols
ospf3 {
  area 0.0.0.1 {
    interface fe-1/1/0.0;
    interface t1-0/2/1.0;
    interface lo0.0 {
      passive;
    }
  }
  area 0.0.0.2 {
    interface fe-0/0/0.0;
  }
  area 0.0.0.0 {
    virtual-link neighbor-id 192.168.3.1 transit-area 0.0.0.1;
  }
}
user@4# show routing-options
router-id 192.168.4.1;
```

Device 5

```
user@5# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet6 {
      address 9009:6::2/64;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.5.1/32;
    }
    family inet6 {
      address feee::10:255:71:6/128;
    }
  }
}
user@5# show protocols
ospf3 {
  area 0.0.0.2 {
    interface fe-0/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
user@5# show routing-options
router-id 192.168.5.1;
```

If you are done configuring the devices, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Device 0 Status | 175](#)
- [Device 1 Status | 178](#)

- Device 2 Status | 181
- Device 3 Status | 184
- Device 4 Status | 188
- Device 5 Status | 192

Confirm that the configuration is working properly.

To verify proper operation of OSPFv3 for IPv6, use the following commands:

- `show ospf3 interface`
- `show ospf3 neighbor`
- `show ospf3 database`
- `show ospf3 route`
- `show interfaces terse` (to see the IPv6 link local address assigned to the **lo0** interface)



NOTE: To view prefix information, you must use the **extensive** option with the `show ospf3 database` command.

Device 0 Status

Purpose

Verify that Device 0 has learned the expected routes and has established the expected neighbor adjacencies.

In the `show ospf3 database` sample output, the stars indicate the “best” routes. These routes are the routes that are installed in the routing table.

Action

```
user@0> show ospf3 database
```

```
Area 0.0.0.0
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	*0.0.0.0	192.168.0.1	0x8000008f	1858	0x6e21	40

```

Router      0.0.0.0          192.168.1.1      0x8000008f  1861  0x523d  40
Router      0.0.0.0          192.168.2.1      0x80000090  1918  0x9e62  56
Router      0.0.0.0          192.168.3.1      0x80000092  2104  0x46d   72
Router      0.0.0.0          192.168.4.1      0x8000008f  2012  0x7016  40
InterArPfx  0.0.0.1          192.168.2.1      0x80000093   231  0xfc5c  36
InterArPfx  0.0.0.2          192.168.2.1      0x80000093   43  0x156   36
InterArPfx  0.0.0.3          192.168.2.1      0x80000092  1731  0x31a4  44
InterArPfx  0.0.0.4          192.168.2.1      0x8000008f  2668  0xc51f  44
InterArPfx  0.0.0.5          192.168.2.1      0x80000091  2856  0xfa59  36
InterArPfx  0.0.0.6          192.168.2.1      0x80000090  2481  0xe3fb  44
InterArPfx  0.0.0.1          192.168.3.1      0x80000093   417  0xf562  36
InterArPfx  0.0.0.2          192.168.3.1      0x80000093  2854  0x84d   36
InterArPfx  0.0.0.3          192.168.3.1      0x80000092  1729  0xbc26  44
InterArPfx  0.0.0.4          192.168.3.1      0x8000008f  2667  0x2ca9  44
InterArPfx  0.0.0.5          192.168.3.1      0x80000091   229  0xe56e  36
InterArPfx  0.0.0.6          192.168.3.1      0x8000008f  2292  0xde01  44
InterArPfx  0.0.0.2          192.168.4.1      0x80000092   794  0xf461  36
InterArPfx  0.0.0.3          192.168.4.1      0x80000092   606  0xf85b  36
InterArPfx  0.0.0.4          192.168.4.1      0x80000091   419  0xfe54  36
InterArPfx  0.0.0.5          192.168.4.1      0x80000090  1825  0xd906  44
InterArPfx  0.0.0.6          192.168.4.1      0x8000008f  2669  0xf1eb  44
InterArPfx  0.0.0.7          192.168.4.1      0x80000091   981  0xbc95  36
InterArPfx  0.0.0.8          192.168.4.1      0x8000008f  2481  0x8f4f  44
InterArPfx  0.0.0.9          192.168.4.1      0x80000090  2294  0xf0dd  44
InterArPfx  0.0.0.10         192.168.4.1      0x8000008f   231  0xac5a  44
IntraArPfx *0.0.0.1         192.168.0.1      0x80000094  2858  0xbf9f  64
IntraArPfx  0.0.0.1          192.168.1.1      0x80000095  2861  0x87d6  64
IntraArPfx  0.0.0.1          192.168.2.1      0x80000096   793  0xc7bd  64
IntraArPfx  0.0.0.1          192.168.3.1      0x80000097  1167  0x93f0  64

```

```
interface so-0/3/2.0 Area 0.0.0.0
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	*0.0.0.2	192.168.0.1	0x80000091	858	0xc0c7	56
Link	0.0.0.8	192.168.3.1	0x80000091	1354	0x84f9	56

```
user@0> show ospf3 interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DRother	0.0.0.0	0.0.0.0	0.0.0.0	0
so-0/3/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

```
user@0> show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
192.168.3.1	so-0/3/2.0	Full	128	33

Neighbor-address fe80::2a0:a514:0:24c

user@0> **show ospf3 route**

Prefix	Path Type	Route Type	NH Type	Metric
192.168.1.1	Intra	Router	IP	3
NH-interface so-0/3/2.0				
192.168.2.1	Intra	Area BR	IP	2
NH-interface so-0/3/2.0				
192.168.3.1	Intra	Area BR	IP	1
NH-interface so-0/3/2.0				
192.168.4.1	Intra	Area BR	IP	2
NH-interface so-0/3/2.0				
9009:1::/64	Intra	Network	IP	1
NH-interface so-0/3/2.0				
9009:1::2/128	Intra	Network	IP	1
NH-interface so-0/3/2.0				
9009:2::/64	Intra	Network	IP	3
NH-interface so-0/3/2.0				
9009:2::2/128	Intra	Network	IP	2
NH-interface so-0/3/2.0				
9009:3::/64	Inter	Network	IP	2
NH-interface so-0/3/2.0				
9009:4::/64	Inter	Network	IP	3
NH-interface so-0/3/2.0				
9009:5::/64	Inter	Network	IP	2
NH-interface so-0/3/2.0				
9009:6::/64	Inter	Network	IP	3
NH-interface so-0/3/2.0				
9009:6::1/128	Inter	Network	IP	2
NH-interface so-0/3/2.0				
fee::10:255:71:1/128	Intra	Network	IP	3
NH-interface so-0/3/2.0				
fee::10:255:71:3/128	Inter	Network	IP	1
NH-interface so-0/3/2.0				
fee::10:255:71:4/128	Intra	Network	IP	0
NH-interface lo0.0				
fee::10:255:71:5/128	Inter	Network	IP	2
NH-interface so-0/3/2.0				
fee::10:255:71:6/128	Inter	Network	IP	3
NH-interface so-0/3/2.0				
fee::10:255:71:11/128	Inter	Network	IP	2
NH-interface so-0/3/2.0				

```

user@0> show interfaces terse
Interface          Admin Link Proto  Local          Remote
lt-1/2/0
so-0/3/2.0         up    up    inet6  9009:1::1/64
                  fe80::2a0:a514:0:14c/64
lo0
lo0.0              up    up    inet   192.168.0.1    --> 0/0
                  inet6  fe80::2a0:a50f:fc56:14c
                  feee::10:255:71:4
...

```

Device 1 Status

Purpose

Verify that Device 1 has learned the expected routes and has established the expected neighbor adjacencies.

Action

```

user@1> show ospf3 interface
Interface          State Area          DR ID          BDR ID          Nbrs
lo0.0              DRother 0.0.0.0        0.0.0.0        0.0.0.0        0
at-2/0/0.0         PtToPt 0.0.0.0        0.0.0.0        0.0.0.0        1
user@1> show ospf3 neighbor
ID                Interface          State Pri Dead
192.168.2.1       at-2/0/0.0        Full 128 37
Neighbor-address fe80::2a0:a514:0:c4c

user@1> show ospf3 database
Area 0.0.0.0
Type ID            Adv Rtr          Seq            Age Cksum Len
Router 0.0.0.0          192.168.0.1     0x8000008f    2334 0x6e21 40
Router *0.0.0.0     192.168.1.1     0x8000008f    2331 0x523d 40
Router 0.0.0.0          192.168.2.1     0x80000090    2390 0x9e62 56
Router 0.0.0.0          192.168.3.1     0x80000092    2578 0x46d 72
Router 0.0.0.0          192.168.4.1     0x8000008f    2486 0x7016 40
InterArPfx 0.0.0.1          192.168.2.1     0x80000093    703 0xfc5c 36
InterArPfx 0.0.0.2          192.168.2.1     0x80000093    515 0x156 36
InterArPfx 0.0.0.3          192.168.2.1     0x80000092    2203 0x31a4 44

```

```

InterArPfx 0.0.0.4      192.168.2.1      0x80000090      140 0xc320 44
InterArPfx 0.0.0.5      192.168.2.1      0x80000092      328 0xf85a 36
InterArPfx 0.0.0.6      192.168.2.1      0x80000090      2953 0xe3fb 44
InterArPfx 0.0.0.1      192.168.3.1      0x80000093      891 0xf562 36
InterArPfx 0.0.0.2      192.168.3.1      0x80000094      328 0x64e 36
InterArPfx 0.0.0.3      192.168.3.1      0x80000092      2203 0xbc26 44
InterArPfx 0.0.0.4      192.168.3.1      0x80000090      141 0x2aaa 44
InterArPfx 0.0.0.5      192.168.3.1      0x80000091      703 0xe56e 36
InterArPfx 0.0.0.6      192.168.3.1      0x8000008f      2766 0xde01 44
InterArPfx 0.0.0.2      192.168.4.1      0x80000092      1268 0xf461 36
InterArPfx 0.0.0.3      192.168.4.1      0x80000092      1080 0xf85b 36
InterArPfx 0.0.0.4      192.168.4.1      0x80000091      893 0xfe54 36
InterArPfx 0.0.0.5      192.168.4.1      0x80000090      2299 0xd906 44
InterArPfx 0.0.0.6      192.168.4.1      0x80000090      143 0xefec 44
InterArPfx 0.0.0.7      192.168.4.1      0x80000091      1455 0xbc95 36
InterArPfx 0.0.0.8      192.168.4.1      0x8000008f      2955 0x8f4f 44
InterArPfx 0.0.0.9      192.168.4.1      0x80000090      2768 0xf0dd 44
InterArPfx 0.0.0.10     192.168.4.1      0x8000008f      705 0xac5a 44
IntraArPfx 0.0.0.1      192.168.0.1      0x80000095      334 0xbda0 64
IntraArPfx *0.0.0.1      192.168.1.1      0x80000096      331 0x85d7 64
IntraArPfx 0.0.0.1      192.168.2.1      0x80000096      1265 0xc7bd 64
IntraArPfx 0.0.0.1      192.168.3.1      0x80000097      1641 0x93f0 64

```

```
interface at-2/0/0.0 Area 0.0.0.0
```

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      *0.0.0.2    192.168.1.1 0x80000091  1331 0xaecd 56
Link      0.0.0.8     192.168.2.1 0x80000091  1453 0x80f3 56

```

```
user@1> show ospf3 route
```

```

Prefix                                         Path Route      NH  Metric
Type Type                                     Type
192.168.0.1                                  Intra Router  IP  3
  NH-interface at-2/0/0.0
192.168.2.1                                  Intra Area BR  IP  1
  NH-interface at-2/0/0.0
192.168.3.1                                  Intra Area BR  IP  2
  NH-interface at-2/0/0.0
192.168.4.1                                  Intra Area BR  IP  3
  NH-interface at-2/0/0.0
9009:1::/64                                   Intra Network  IP  3
  NH-interface at-2/0/0.0
9009:1::/128                                  Intra Network  IP  2
  NH-interface at-2/0/0.0

```

```

9009:2::/64                               Intra Network   IP   1
  NH-interface at-2/0/0.0
9009:2::2/128                             Intra Network   IP   1
  NH-interface at-2/0/0.0
9009:3::/64                               Inter Network   IP   2
  NH-interface at-2/0/0.0
9009:4::/64                               Inter Network   IP   2
  NH-interface at-2/0/0.0
9009:5::/64                               Inter Network   IP   3
  NH-interface at-2/0/0.0
9009:6::/64                               Inter Network   IP   4
  NH-interface at-2/0/0.0
9009:6::1/128                             Inter Network   IP   3
  NH-interface at-2/0/0.0
feee::10:255:71:1/128                   Intra Network   IP   0
  NH-interface lo0.0
feee::10:255:71:3/128                   Inter Network   IP   2
  NH-interface at-2/0/0.0
feee::10:255:71:4/128                   Intra Network   IP   3
  NH-interface at-2/0/0.0
feee::10:255:71:5/128                   Inter Network   IP   2
  NH-interface at-2/0/0.0
feee::10:255:71:6/128                   Inter Network   IP   4
  NH-interface at-2/0/0.0
feee::10:255:71:11/128                  Inter Network   IP   1
  NH-interface at-2/0/0.0

```

```
user@1> show interfaces terse
```

```

Interface          Admin Link Proto  Local                Remote
lt-1/2/0
at-2/0/0.0         up    up    inet6  9009:2::1/64
                  fe80::2a0:a514:0:b4c/64
lo0
lo0.0              up    up    inet   192.168.1.1         --> 0/0
                  inet6  fe80::2a0:a50f:fc56:14c
                  feee::10:255:71:1
...

```

Device 2 Status

Purpose

Verify that Device 2 has learned the expected routes and has established the expected neighbor adjacencies.

Action

```
user@2> show ospf3 interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
at-0/3/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
v1-192.168.3.1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
lo0.0	DRother	0.0.0.1	0.0.0.0	0.0.0.0	0
so-0/2/0.0	PtToPt	0.0.0.1	0.0.0.0	0.0.0.0	1
fe-1/1/0.0	PtToPt	0.0.0.1	0.0.0.0	0.0.0.0	1

```
user@2> show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
192.168.1.1	at-0/3/1.0	Full	128	32
	Neighbor-address fe80::2a0:a514:0:b4c			
192.168.3.1	v1-192.168.3.1	Full	0	35
	Neighbor-address 9009:3::2			
192.168.3.1	so-0/2/0.0	Full	128	38
	Neighbor-address fe80::2a0:a514:0:74c			
192.168.4.1	fe-1/1/0.0	Full	128	30
	Neighbor-address fe80::2a0:a514:0:a4c			

```
user@2> show ospf3 database
```

```
Area 0.0.0.0
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	0.0.0.0	192.168.0.1	0x8000008f	2771	0x6e21	40
Router	0.0.0.0	192.168.1.1	0x8000008f	2770	0x523d	40
Router	*0.0.0.0	192.168.2.1	0x80000090	2827	0x9e62	56
Router	0.0.0.0	192.168.3.1	0x80000093	15	0x26e	72
Router	0.0.0.0	192.168.4.1	0x8000008f	2923	0x7016	40
InterArPfx	*0.0.0.1	192.168.2.1	0x80000093	1140	0xfc5c	36
InterArPfx	*0.0.0.2	192.168.2.1	0x80000093	952	0x156	36
InterArPfx	*0.0.0.3	192.168.2.1	0x80000092	2640	0x31a4	44
InterArPfx	*0.0.0.4	192.168.2.1	0x80000090	577	0xc320	44
InterArPfx	*0.0.0.5	192.168.2.1	0x80000092	765	0xf85a	36
InterArPfx	*0.0.0.6	192.168.2.1	0x80000091	390	0xe1fc	44

InterArPfx	0.0.0.1	192.168.3.1	0x80000093	1328	0xf562	36
InterArPfx	0.0.0.2	192.168.3.1	0x80000094	765	0x64e	36
InterArPfx	0.0.0.3	192.168.3.1	0x80000092	2640	0xbc26	44
InterArPfx	0.0.0.4	192.168.3.1	0x80000090	578	0x2aaa	44
InterArPfx	0.0.0.5	192.168.3.1	0x80000091	1140	0xe56e	36
InterArPfx	0.0.0.6	192.168.3.1	0x80000090	203	0xdc02	44
InterArPfx	0.0.0.2	192.168.4.1	0x80000092	1705	0xf461	36
InterArPfx	0.0.0.3	192.168.4.1	0x80000092	1517	0xf85b	36
InterArPfx	0.0.0.4	192.168.4.1	0x80000091	1330	0xfe54	36
InterArPfx	0.0.0.5	192.168.4.1	0x80000090	2736	0xd906	44
InterArPfx	0.0.0.6	192.168.4.1	0x80000090	580	0xefec	44
InterArPfx	0.0.0.7	192.168.4.1	0x80000091	1892	0xbc95	36
InterArPfx	0.0.0.8	192.168.4.1	0x80000090	392	0x8d50	44
InterArPfx	0.0.0.9	192.168.4.1	0x80000091	205	0xeede	44
InterArPfx	0.0.0.10	192.168.4.1	0x8000008f	1142	0xac5a	44
IntraArPfx	0.0.0.1	192.168.0.1	0x80000095	771	0xbda0	64
IntraArPfx	0.0.0.1	192.168.1.1	0x80000096	770	0x85d7	64
IntraArPfx	*0.0.0.1	192.168.2.1	0x80000096	1702	0xc7bd	64
IntraArPfx	0.0.0.1	192.168.3.1	0x80000097	2078	0x93f0	64

Area 0.0.0.1

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	*0.0.0.0	192.168.2.1	0x80000093	15	0x8f62	56
Router	0.0.0.0	192.168.3.1	0x80000093	2828	0x39b7	56
Router	0.0.0.0	192.168.4.1	0x80000092	16	0x8768	56
InterArPfx	*0.0.0.1	192.168.2.1	0x80000094	1515	0xec6c	36
InterArPfx	*0.0.0.3	192.168.2.1	0x80000090	202	0x994d	44
InterArPfx	*0.0.0.4	192.168.2.1	0x8000008f	1327	0xd839	44
InterArPfx	0.0.0.1	192.168.3.1	0x80000094	1703	0xd781	36
InterArPfx	0.0.0.3	192.168.3.1	0x80000090	390	0xe002	44
InterArPfx	0.0.0.4	192.168.3.1	0x8000008f	1515	0xc34e	44
InterArPfx	0.0.0.1	192.168.4.1	0x80000093	1422	0x193b	36
InterArPfx	0.0.0.3	192.168.4.1	0x80000090	672	0xed1	44
InterArPfx	0.0.0.4	192.168.4.1	0x8000008f	1235	0xe824	44
IntraArPfx	*0.0.0.1	192.168.2.1	0x80000097	2265	0x6bf1	76
IntraArPfx	0.0.0.1	192.168.3.1	0x80000099	953	0xad8	76
IntraArPfx	0.0.0.1	192.168.4.1	0x80000098	2079	0x3c26	76

interface at-0/3/1.0 Area 0.0.0.0

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	0.0.0.2	192.168.1.1	0x80000091	1770	0xaecd	56
Link	*0.0.0.8	192.168.2.1	0x80000091	1890	0x80f3	56

```
interface so-0/2/0.0 Area 0.0.0.1
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	*0.0.0.6	192.168.2.1	0x80000092	2452	0x6018	56
Link	0.0.0.7	192.168.3.1	0x80000092	2453	0x3a3d	56

```
interface fe-1/1/0.0 Area 0.0.0.1
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	*0.0.0.7	192.168.2.1	0x80000092	2077	0x8de7	56
Link	0.0.0.8	192.168.4.1	0x80000091	2172	0x8ce5	56

```
user@2> show ospf3 route
```

Prefix	Path	Route Type	NH Type	Metric
192.168.0.1		Intra Router	IP	2
NH-interface (null), NH-addr feee::10:255:71:3				
192.168.1.1		Intra Router	IP	1
NH-interface at-0/3/1.0				
192.168.3.1		Intra Area BR	IP	1
NH-interface so-0/2/0.0				
192.168.4.1		Intra Area BR	IP	1
NH-interface fe-1/1/0.0				
9009:1::/64		Intra Network	IP	2
NH-interface so-0/2/0.0				
9009:1::2/128		Intra Network	IP	1
NH-interface so-0/2/0.0				
9009:2::/64		Intra Network	IP	1
NH-interface at-0/3/1.0				
9009:2::2/128		Intra Network	IP	0
NH-interface at-0/3/1.0				
9009:3::/64		Intra Network	IP	1
NH-interface so-0/2/0.0				
9009:4::/64		Intra Network	IP	1
NH-interface fe-1/1/0.0				
9009:5::/64		Intra Network	IP	2
NH-interface so-0/2/0.0				
NH-interface fe-1/1/0.0				
9009:6::/64		Inter Network	IP	2
NH-interface fe-1/1/0.0				
9009:6::1/128		Inter Network	IP	1
NH-interface fe-1/1/0.0				
feee::10:255:71:1/128		Intra Network	IP	1
NH-interface at-0/3/1.0				
feee::10:255:71:3/128		Intra Network	IP	1

```

NH-interface so-0/2/0.0
feee::10:255:71:4/128          Intra Network   IP   2
NH-interface so-0/2/0.0
feee::10:255:71:5/128          Intra Network   IP   1
NH-interface fe-1/1/0.0
feee::10:255:71:6/128          Inter Network   IP   2
NH-interface fe-1/1/0.0
feee::10:255:71:11/128         Intra Network   IP   0
NH-interface lo0.0

```

```
user@2> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
lt-1/2/0					
so-0/2/0.0	up	up	inet6	9009:3::1/64 fe80::2a0:a514:0:84c/64	
fe-1/1/0.0	up	up	inet6	9009:4::1/64 fe80::2a0:a514:0:94c/64	
at-0/3/1.0	up	up	inet6	9009:2::2/64 fe80::2a0:a514:0:c4c/64	
lo0					
lo0.0	up	up	inet inet6	192.168.2.1 fe80::2a0:a50f:fc56:14c feee::10:255:71:11	--> 0/0
...					

Device 3 Status

Purpose

Verify that Device 3 has learned the expected routes and has established the expected neighbor adjacencies.

Action

```
user@3> show ospf3 interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/3/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
v1-192.168.2.1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
v1-192.168.4.1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
lo0.0	DRother	0.0.0.1	0.0.0.0	0.0.0.0	0
t1-0/2/1.0	PtToPt	0.0.0.1	0.0.0.0	0.0.0.0	1

```
so-0/3/0.0      PtToPt 0.0.0.1      0.0.0.0      0.0.0.0      1
```

```
user@3> show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
192.168.0.1	so-0/3/2.0	Full	128	31
	Neighbor-address fe80::2a0:a514:0:14c			
192.168.2.1	v1-192.168.2.1	Full	0	33
	Neighbor-address 9009:3::1			
192.168.4.1	v1-192.168.4.1	Full	0	38
	Neighbor-address 9009:5::2			
192.168.4.1	t1-0/2/1.0	Full	128	35
	Neighbor-address fe80::2a0:a514:0:44c			
192.168.2.1	so-0/3/0.0	Full	128	37
	Neighbor-address fe80::2a0:a514:0:84c			

```
user@3> show ospf3 database
```

```
Area 0.0.0.0
```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	0.0.0.0	192.168.0.1	0x80000090	11	0x6c22	40
Router	0.0.0.0	192.168.1.1	0x80000090	12	0x503e	40
Router	0.0.0.0	192.168.2.1	0x80000091	69	0x9c63	56
Router	*0.0.0.0	192.168.3.1	0x80000093	255	0x26e	72
Router	0.0.0.0	192.168.4.1	0x80000090	163	0x6e17	40
InterArPfx	0.0.0.1	192.168.2.1	0x80000093	1382	0xfc5c	36
InterArPfx	0.0.0.2	192.168.2.1	0x80000093	1194	0x156	36
InterArPfx	0.0.0.3	192.168.2.1	0x80000092	2882	0x31a4	44
InterArPfx	0.0.0.4	192.168.2.1	0x80000090	819	0xc320	44
InterArPfx	0.0.0.5	192.168.2.1	0x80000092	1007	0xf85a	36
InterArPfx	0.0.0.6	192.168.2.1	0x80000091	632	0xe1fc	44
InterArPfx	*0.0.0.1	192.168.3.1	0x80000093	1568	0xf562	36
InterArPfx	*0.0.0.2	192.168.3.1	0x80000094	1005	0x64e	36
InterArPfx	*0.0.0.3	192.168.3.1	0x80000092	2880	0xbc26	44
InterArPfx	*0.0.0.4	192.168.3.1	0x80000090	818	0x2aaa	44
InterArPfx	*0.0.0.5	192.168.3.1	0x80000091	1380	0xe56e	36
InterArPfx	*0.0.0.6	192.168.3.1	0x80000090	443	0xdc02	44
InterArPfx	0.0.0.2	192.168.4.1	0x80000092	1945	0xf461	36
InterArPfx	0.0.0.3	192.168.4.1	0x80000092	1757	0xf85b	36
InterArPfx	0.0.0.4	192.168.4.1	0x80000091	1570	0xfe54	36
InterArPfx	0.0.0.5	192.168.4.1	0x80000090	2976	0xd906	44
InterArPfx	0.0.0.6	192.168.4.1	0x80000090	820	0xefec	44
InterArPfx	0.0.0.7	192.168.4.1	0x80000091	2132	0xbc95	36
InterArPfx	0.0.0.8	192.168.4.1	0x80000090	632	0x8d50	44
InterArPfx	0.0.0.9	192.168.4.1	0x80000091	445	0xeede	44

```

InterArPfx 0.0.0.10      192.168.4.1      0x8000008f  1382 0xac5a 44
IntraArPfx 0.0.0.1       192.168.0.1      0x80000095  1011 0xbda0 64
IntraArPfx 0.0.0.1       192.168.1.1      0x80000096  1012 0x85d7 64
IntraArPfx 0.0.0.1       192.168.2.1      0x80000096  1944 0xc7bd 64
IntraArPfx *0.0.0.1     192.168.3.1      0x80000097  2318 0x93f0 64

```

Area 0.0.0.1

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Router    0.0.0.0     192.168.2.1 0x80000093  257 0x8f62 56
Router    *0.0.0.0    192.168.3.1 0x80000094   68 0x37b8 56
Router    0.0.0.0     192.168.4.1 0x80000092  257 0x8768 56
InterArPfx 0.0.0.1     192.168.2.1 0x80000094  1757 0xec6c 36
InterArPfx 0.0.0.3     192.168.2.1 0x80000090   444 0x994d 44
InterArPfx 0.0.0.4     192.168.2.1 0x8000008f  1569 0xd839 44
InterArPfx *0.0.0.1     192.168.3.1 0x80000094  1943 0xd781 36
InterArPfx *0.0.0.3     192.168.3.1 0x80000090   630 0xe002 44
InterArPfx *0.0.0.4     192.168.3.1 0x8000008f  1755 0xc34e 44
InterArPfx 0.0.0.1     192.168.4.1 0x80000093  1663 0x193b 36
InterArPfx 0.0.0.3     192.168.4.1 0x80000090   913 0xed1 44
InterArPfx 0.0.0.4     192.168.4.1 0x8000008f  1476 0xe824 44
IntraArPfx 0.0.0.1     192.168.2.1 0x80000097  2507 0x6bf1 76
IntraArPfx *0.0.0.1     192.168.3.1 0x80000099  1193 0xad8 76
IntraArPfx 0.0.0.1     192.168.4.1 0x80000098  2320 0x3c26 76

```

interface so-0/3/2.0 Area 0.0.0.0

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.2     192.168.0.1 0x80000091  2011 0xc0c7 56
Link      *0.0.0.8    192.168.3.1 0x80000091  2505 0x84f9 56

```

interface t1-0/2/1.0 Area 0.0.0.1

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      *0.0.0.9    192.168.3.1 0x80000092  2130 0x1661 56
Link      0.0.0.7     192.168.4.1 0x80000092  2507 0x383f 56

```

interface so-0/3/0.0 Area 0.0.0.1

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.6     192.168.2.1 0x80000092  2694 0x6018 56
Link      *0.0.0.7    192.168.3.1 0x80000092  2693 0x3a3d 56

```

user@3> **show ospf3 route**

```

Prefix                                     Path Route      NH  Metric
                                           Type Type      Type
192.168.0.1                               Intra Router  IP  1

```

```

NH-interface so-0/3/2.0
192.168.1.1          Intra Router   IP   2
  NH-interface (null), NH-addr feee::10:255:71:11
192.168.2.1          Intra Area BR  IP   1
  NH-interface so-0/3/0.0
192.168.4.1          Intra Area BR  IP   1
  NH-interface t1-0/2/1.0
9009:1::/64          Intra Network  IP   1
  NH-interface so-0/3/2.0
9009:1::2/128        Intra Network  IP   0
  NH-interface so-0/3/2.0
9009:2::/64          Intra Network  IP   2
  NH-interface so-0/3/0.0
9009:2::2/128        Intra Network  IP   1
  NH-interface so-0/3/0.0
9009:3::/64          Intra Network  IP   1
  NH-interface so-0/3/0.0
9009:4::/64          Intra Network  IP   2
  NH-interface so-0/3/0.0
  NH-interface t1-0/2/1.0
9009:5::/64          Intra Network  IP   1
  NH-interface t1-0/2/1.0
9009:6::/64          Inter Network  IP   2
  NH-interface t1-0/2/1.0
9009:6::1/128        Inter Network  IP   1
  NH-interface t1-0/2/1.0
feee::10:255:71:1/128 Intra Network  IP   2
  NH-interface so-0/3/0.0
feee::10:255:71:3/128 Intra Network  IP   0
  NH-interface lo0.0
feee::10:255:71:4/128 Intra Network  IP   1
  NH-interface so-0/3/2.0
feee::10:255:71:5/128 Intra Network  IP   1
  NH-interface t1-0/2/1.0
feee::10:255:71:6/128 Inter Network  IP   2
  NH-interface t1-0/2/1.0
feee::10:255:71:11/128 Intra Network  IP   1
  NH-interface so-0/3/0.0

```

```
user@3> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
lt-1/2/0					
so-0/3/2.0	up	up	inet6	9009:1::2/64	

```

t1-0/2/1.0          up    up    inet6  fe80::2a0:a514:0:24c/64
                  up    up    inet6  9009:5::1/64
                  up    up    inet6  fe80::2a0:a514:0:34c/64
so-0/3/0.0          up    up    inet6  9009:3::2/64
                  up    up    inet6  fe80::2a0:a514:0:74c/64
lo0
lo0.0               up    up    inet   192.168.3.1      --> 0/0
                  up    up    inet6  fe80::2a0:a50f:fc56:14c
                  up    up    inet6  feee::10:255:71:3
...

```

Device 4 Status

Purpose

Verify that Device 4 has learned the expected routes and has established the expected neighbor adjacencies.

Action

```

user@4> show ospf3 interface
Interface          State Area          DR ID          BDR ID          Nbrs
lo0.0              DROther 0.0.0.1        0.0.0.0         0.0.0.0         0
fe-1/1/0.0         PtToPt  0.0.0.1        0.0.0.0         0.0.0.0         1
t1-0/2/1.0         PtToPt  0.0.0.1        0.0.0.0         0.0.0.0         1
fe-0/0/0.0         PtToPt  0.0.0.2        0.0.0.0         0.0.0.0         1
vl-192.168.3.1    PtToPt  0.0.0.0        0.0.0.0         0.0.0.0         1

user@4> show ospf3 neighbor
ID                Interface          State  Pri  Dead
192.168.2.1       fe-1/1/0.0        Full   128  35
  Neighbor-address fe80::2a0:a514:0:94c
192.168.3.1       t1-0/2/1.0        Full   128  34
  Neighbor-address fe80::2a0:a514:0:34c
192.168.5.1       fe-0/0/0.0        Full   128  39
  Neighbor-address fe80::2a0:a514:0:64c
192.168.3.1       vl-192.168.3.1    Full    0   33
  Neighbor-address 9009:5::1

user@4> show ospf3 database
Area 0.0.0.0

```

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	0.0.0.0	192.168.0.1	0x80000090	270	0x6c22	40
Router	0.0.0.0	192.168.1.1	0x80000090	271	0x503e	40
Router	0.0.0.0	192.168.2.1	0x80000091	328	0x9c63	56
Router	0.0.0.0	192.168.3.1	0x80000093	514	0x26e	72
Router	*0.0.0.0	192.168.4.1	0x80000090	420	0x6e17	40
InterArPfx	0.0.0.1	192.168.2.1	0x80000093	1641	0xfc5c	36
InterArPfx	0.0.0.2	192.168.2.1	0x80000093	1453	0x156	36
InterArPfx	0.0.0.3	192.168.2.1	0x80000093	141	0x2fa5	44
InterArPfx	0.0.0.4	192.168.2.1	0x80000090	1078	0xc320	44
InterArPfx	0.0.0.5	192.168.2.1	0x80000092	1266	0xf85a	36
InterArPfx	0.0.0.6	192.168.2.1	0x80000091	891	0xe1fc	44
InterArPfx	0.0.0.1	192.168.3.1	0x80000093	1827	0xf562	36
InterArPfx	0.0.0.2	192.168.3.1	0x80000094	1264	0x64e	36
InterArPfx	0.0.0.3	192.168.3.1	0x80000093	139	0xba27	44
InterArPfx	0.0.0.4	192.168.3.1	0x80000090	1077	0x2aaa	44
InterArPfx	0.0.0.5	192.168.3.1	0x80000091	1639	0xe56e	36
InterArPfx	0.0.0.6	192.168.3.1	0x80000090	702	0xdc02	44
InterArPfx	*0.0.0.2	192.168.4.1	0x80000092	2202	0xf461	36
InterArPfx	*0.0.0.3	192.168.4.1	0x80000092	2014	0xf85b	36
InterArPfx	*0.0.0.4	192.168.4.1	0x80000091	1827	0xfe54	36
InterArPfx	*0.0.0.5	192.168.4.1	0x80000091	233	0xd707	44
InterArPfx	*0.0.0.6	192.168.4.1	0x80000090	1077	0xefec	44
InterArPfx	*0.0.0.7	192.168.4.1	0x80000091	2389	0xbc95	36
InterArPfx	*0.0.0.8	192.168.4.1	0x80000090	889	0x8d50	44
InterArPfx	*0.0.0.9	192.168.4.1	0x80000091	702	0xeede	44
InterArPfx	*0.0.0.10	192.168.4.1	0x8000008f	1639	0xac5a	44
IntraArPfx	0.0.0.1	192.168.0.1	0x80000095	1270	0xbda0	64
IntraArPfx	0.0.0.1	192.168.1.1	0x80000096	1271	0x85d7	64
IntraArPfx	0.0.0.1	192.168.2.1	0x80000096	2203	0xc7bd	64
IntraArPfx	0.0.0.1	192.168.3.1	0x80000097	2577	0x93f0	64

Area 0.0.0.1

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Router	0.0.0.0	192.168.2.1	0x80000093	515	0x8f62	56
Router	0.0.0.0	192.168.3.1	0x80000094	327	0x37b8	56
Router	*0.0.0.0	192.168.4.1	0x80000092	514	0x8768	56
InterArPfx	0.0.0.1	192.168.2.1	0x80000094	2015	0xec6c	36
InterArPfx	0.0.0.3	192.168.2.1	0x80000090	702	0x994d	44
InterArPfx	0.0.0.4	192.168.2.1	0x8000008f	1827	0xd839	44
InterArPfx	0.0.0.1	192.168.3.1	0x80000094	2202	0xd781	36
InterArPfx	0.0.0.3	192.168.3.1	0x80000090	889	0xe002	44
InterArPfx	0.0.0.4	192.168.3.1	0x8000008f	2014	0xc34e	44

	Type	Type	Type	
192.168.0.1	Intra	Router	IP	2
NH-interface (null), NH-addr feee::10:255:71:3				
192.168.1.1	Intra	Router	IP	3
NH-interface (null), NH-addr feee::10:255:71:3				
192.168.2.1	Intra	Area BR	IP	1
NH-interface fe-1/1/0.0				
192.168.3.1	Intra	Area BR	IP	1
NH-interface t1-0/2/1.0				
192.168.5.1	Intra	Router	IP	1
NH-interface fe-0/0/0.0				
9009:1::/64	Intra	Network	IP	2
NH-interface t1-0/2/1.0				
9009:1::2/128	Intra	Network	IP	1
NH-interface t1-0/2/1.0				
9009:2::/64	Intra	Network	IP	2
NH-interface fe-1/1/0.0				
9009:2::2/128	Intra	Network	IP	1
NH-interface fe-1/1/0.0				
9009:3::/64	Intra	Network	IP	2
NH-interface t1-0/2/1.0				
NH-interface fe-1/1/0.0				
9009:4::/64	Intra	Network	IP	1
NH-interface fe-1/1/0.0				
9009:5::/64	Intra	Network	IP	1
NH-interface t1-0/2/1.0				
9009:6::/64	Intra	Network	IP	1
NH-interface fe-0/0/0.0				
9009:6::1/128	Intra	Network	IP	0
NH-interface fe-0/0/0.0				
feee::10:255:71:1/128	Intra	Network	IP	2
NH-interface fe-1/1/0.0				
feee::10:255:71:3/128	Intra	Network	IP	1
NH-interface t1-0/2/1.0				
feee::10:255:71:4/128	Intra	Network	IP	2
NH-interface t1-0/2/1.0				
feee::10:255:71:5/128	Intra	Network	IP	0
NH-interface lo0.0				
feee::10:255:71:6/128	Intra	Network	IP	1
NH-interface fe-0/0/0.0				
feee::10:255:71:11/128	Intra	Network	IP	1
NH-interface fe-1/1/0.0				

```

user@4> show interfaces terse
Interface          Admin Link Proto  Local          Remote
lt-1/2/0
t1-0/2/1.0        up    up    inet6  9009:5::2/64
                  fe80::2a0:a514:0:44c/64
fe-0/0/0.0        up    up    inet6  9009:6::1/64
                  fe80::2a0:a514:0:54c/64
fe-1/1/0.0        up    up    inet6  9009:4::2/64
                  fe80::2a0:a514:0:a4c/64
lo0
lo0.0             up    up    inet   192.168.4.1    --> 0/0
                  inet6  fe80::2a0:a50f:fc56:14c
                  feee::10:255:71:5
...

```

Device 5 Status

Purpose

Verify that Device 5 has learned the expected routes and has established the expected neighbor adjacencies.

Action

```

user@5> show ospf3 interface
Interface      State Area          DR ID          BDR ID          Nbrs
lo0.0          DRother 0.0.0.2        0.0.0.0        0.0.0.0         0
fe-0/0/0.0     PtToPt 0.0.0.2        0.0.0.0        0.0.0.0         1
user@5> show ospf3 neighbor
ID             Interface          State   Pri   Dead
192.168.4.1    fe-0/0/0.0        Full   128   34
  Neighbor-address fe80::2a0:a514:0:54c

user@5> show ospf3 database
Area 0.0.0.2
Type  ID             Adv Rtr          Seq           Age  Cksum  Len
Router  0.0.0.0        192.168.4.1     0x80000091    509  0x4741  40
Router  *0.0.0.0       192.168.5.1     0x80000090    732  0x3a50  40
InterArPfx 0.0.0.1        192.168.4.1     0x80000094    2759 0xfa5a  36
InterArPfx 0.0.0.2        192.168.4.1     0x80000094    2572 0xfe54  36
InterArPfx 0.0.0.3        192.168.4.1     0x80000093    603  0xe7f6  44

```

```

InterArPfx 0.0.0.4      192.168.4.1      0x80000091 2947 0xda7a 36
InterArPfx 0.0.0.5      192.168.4.1      0x80000090 1447 0xab35 44
InterArPfx 0.0.0.6      192.168.4.1      0x80000091 1259 0xdc3 44
InterArPfx 0.0.0.7      192.168.4.1      0x80000090 2009 0xa2b2 36
InterArPfx 0.0.0.9      192.168.4.1      0x80000090 1822 0x9cb5 36
InterArPfx 0.0.0.11     192.168.4.1      0x80000090 1072 0x8f49 44
InterArPfx 0.0.0.12     192.168.4.1      0x80000090 791 0x37a3 44
InterArPfx 0.0.0.13     192.168.4.1      0x8000008f 1916 0x689e 44
InterArPfx 0.0.0.14     192.168.4.1      0x8000008f 1728 0x6c98 44
IntraArPfx 0.0.0.1      192.168.4.1      0x80000099 322 0x80f6 64
IntraArPfx *0.0.0.1     192.168.5.1      0x80000095 1732 0xf25a 64

```

```
interface fe-0/0/0.0 Area 0.0.0.2
```

```

Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.6     192.168.4.1 0x80000093  416  0x77fd 56
Link      *0.0.0.2    192.168.5.1 0x80000091  2732 0xb1c7 56

```

```
user@5> show interfaces terse
```

```

Interface      Admin Link Proto  Local          Remote
lt-1/2/0
fe-0/0/0.0     up    up    inet6  9009::6::2/64
                fe80::2a0:a514:0:64c/64
lo0
lo0.0          up    up    inet   192.168.5.1    --> 0/0
                inet6  fe80::2a0:a50f:fc56:14c
                feee::10:255:71:6
...

```

RELATED DOCUMENTATION

[OSPF Overview | 2](#)

[OSPF Packets Overview | 7](#)

[Understanding OSPF Configurations | 14](#)

5

CHAPTER

Configure OSPF Route Control

IN THIS CHAPTER

- [Configuring OSPF Route Control | 195](#)
-

Configuring OSPF Route Control

IN THIS SECTION

- [Understanding OSPF Route Summarization | 195](#)
- [Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements Sent into the Backbone Area | 196](#)
- [Example: Limiting the Number of Prefixes Exported to OSPF | 205](#)
- [Understanding OSPF Traffic Control | 208](#)
- [Example: Controlling the Cost of Individual OSPF Network Segments | 211](#)
- [Understanding Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors | 217](#)
- [Example: Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors | 218](#)
- [Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth | 231](#)
- [Example: Controlling OSPF Route Preferences | 235](#)
- [Understanding OSPF Overload Function | 238](#)
- [Example: Configuring OSPF to Make Routing Devices Appear Overloaded | 240](#)
- [Understanding the SPF Algorithm Options for OSPF | 245](#)
- [Example: Configuring SPF Algorithm Options for OSPF | 246](#)
- [Configuring OSPF Refresh and Flooding Reduction in Stable Topologies | 250](#)
- [Understanding Synchronization Between LDP and IGP | 252](#)
- [Example: Configuring Synchronization Between LDP and OSPF | 252](#)
- [OSPFv2 Compatibility with RFC 1583 Overview | 257](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 | 258](#)

Understanding OSPF Route Summarization

Area border routers (ABRs) send summary link advertisements to describe the routes to other areas. Depending on the number of destinations, an area can get flooded with a large number of link-state records, which can utilize routing device resources. To minimize the number of advertisements that are flooded into an area, you can configure the ABR to coalesce, or summarize, a range of IP addresses and send reachability information about these addresses in a single link-state advertisement (LSA). You can

summarize one or more ranges of IP addresses, where all routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

For an OSPF area, you can summarize and filter intra-area prefixes. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place. For an OSPF not-so-stubby area (NSSA), you can only coalesce or filter NSSA external (Type 7) LSAs before they are translated into AS external (Type 5) LSAs and enter the backbone area. All external routes learned within the area that do not fall into the range of one of the prefixes are advertised individually to other areas.

In addition, you can also limit the number of prefixes (routes) that are exported into OSPF. By setting a user-defined maximum number of prefixes, you prevent the routing device from flooding an excessive number of routes into an area.

Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements Sent into the Backbone Area

IN THIS SECTION

- Requirements | 196
- Overview | 197
- Configuration | 199
- Verification | 204

This example shows how to summarize routes sent into the backbone area.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a static route. See [Examples: Configuring Static Routes](#) in the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview


IN THIS SECTION

- [Topology | 198](#)

You can summarize a range of IP addresses to minimize the size of the backbone router's link-state database. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

[Figure 18 on page 198](#) shows the topology used in this example. R5 is the ABR between area 0.0.0.4 and the backbone. The networks in area 0.0.0.4 are 10.0.8.4/30, 10.0.8.0/30, and 10.0.8.8/30, which can be summarized as 10.0.8.0/28. R3 is the ABR between NSSA area 0.0.0.3 and the backbone. The networks in area 0.0.0.3 are 10.0.4.4/30, 10.0.4.0/30, and 10.0.4.12/30, which can be summarized as 10.0.4.0/28. Area 0.0.0.3 also contains external static route 3.0.0.8, which will be flooded throughout the network.

Figure 18: Summarizing Ranges of Routes in OSPF



In this example, you configure the ABRs for route summarization by including the following settings:

- **area-range**—For an area, summarizes a range of IP addresses when sending summary intra-area link advertisements. For an NSSA, summarizes a range of IP addresses when sending NSSA link-state advertisements (Type 7 LSAs). The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas.
- **network/mask-length**—Indicates the summarized IP address range and the number of significant bits in the network mask.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 199](#)
- [Procedure | 200](#)
- [Results | 202](#)

CLI Quick Configuration

- To quickly configure route summarization for an OSPF area, copy the following commands and paste them into the CLI. The following is the configuration on ABR R5:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3/30
set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4/30
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3/30
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5/30
set protocols ospf area 0.0.0.4 stub
set protocols ospf area 0.0.0.4 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28
```

- To quickly configure route summarization for an OSPF NSSA, copy the following commands and paste them into the CLI. The following is the configuration on ABR R3:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10/30
set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1/30
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1/30
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7/30
set protocols ospf area 0.0.0.3 interface fe-0/0/1
set protocols ospf area 0.0.0.3 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
```

```
set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
set protocols ospf area 0.0.0.3 nssa
set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8
```

Procedure

Step-by-Step Procedure

To summarize routes sent to the backbone area:

1. Configure the interfaces.



NOTE: For OSPFv3, include IPv6 addresses.

```
[edit]
user@R5# set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3/30
user@R5# set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4/30
user@R5# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3/30
user@R5# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5/30
```

```
[edit]
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10/30
user@R3# set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1/30
user@R3# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1/30
user@R3# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7/30
```

2. Configure the type of OSPF area.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 stub
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa
```

3. Assign the interfaces to the OSPF areas.

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/1
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/2
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/4
```

```
user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/1
user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/2
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/4
```

4. Summarize the routes that are flooded into the backbone.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
```

5. On ABR R3, restrict the external static route from leaving area 0.0.0.3.

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8
```

6. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R5:

```
user@R5# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.2.3/32;
    }
  }
}
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.8.3/32;
    }
  }
}
fe-0/0/2 {
  unit 0 {
    family inet {
      address 10.0.8.4/32;
    }
  }
}
fe-0/0/4 {
  unit 0 {
    family inet {
      address 10.0.2.5/32;
    }
  }
}
```

```
}  
}
```

```
user@R5# show protocols ospf  
area 0.0.0.0 {  
    interface fe-0/0/0.0;  
    interface fe-0/0/4.0;  
}  
area 0.0.0.4 {  
    stub;  
    area-range 10.0.8.0/28;  
    interface fe-0/0/1.0;  
    interface fe-0/0/2.0;  
}
```

Configuration on ABR R3:

```
user@R3# show interfaces  
fe-0/0/0 {  
    unit 0 {  
        family inet {  
            address 10.0.2.1/32;  
        }  
    }  
}  
fe-0/0/1 {  
    unit 0 {  
        family inet {  
            address 10.0.4.10/32;  
        }  
    }  
}  
fe-0/0/2 {  
    unit 0 {  
        family inet {  
            address 10.0.4.1/32;  
        }  
    }  
}  
fe-0/0/4 {  
    unit 0 {
```

```
    family inet {
        address 10.0.2.7/32;
    }
}
```

```
user@R3t# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/0.0;
    interface fe-0/0/4.0;
}
area 0.0.0.3 {
    nssa {
        area-range 3.0.0.0/8 ;
    }
    area-range 10.0.4.0/28;
    interface fe-0/0/1.0;
    interface fe-0/0/2.0;
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces` and `show protocols ospf3` commands.

Verification

IN THIS SECTION

- [Verifying the Summarized Route | 204](#)

Confirm that the configuration is working properly.

Verifying the Summarized Route

Purpose

Verify that the routes you configured for route summarization are being aggregated by the ABRs before the routes enter the backbone area. Confirm route summarization by checking the entries of the OSPF link-state database for the routing devices in the backbone.

Action

From operational mode, enter the `show ospf database` command for OSPFv2, and enter the `show ospf3 database` command for OSPFv3.

Example: Limiting the Number of Prefixes Exported to OSPF

IN THIS SECTION

- [Requirements | 205](#)
- [Overview | 206](#)
- [Configuration | 206](#)
- [Verification | 208](#)

This example shows how to limit the number of prefixes exported to OSPF.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 206](#)

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area.

You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem. If the routing device exceeds the configured prefix export value, the routing device purges the external prefixes and enters into an overload state. This state ensures that the routing device is not overwhelmed as it attempts to process routing information. The prefix export limit number can be a value from 0 through 4,294,967,295.

In this example, you configure a prefix export limit of 100,000 by including the `prefix-export-limit` statement.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 206](#)
- [Procedure | 207](#)
- [Results | 207](#)

CLI Quick Configuration

To quickly limit the number of prefixes exported to OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter commit from configuration mode.

```
[edit]
set protocols ospf prefix-export-limit 100000
```

Procedure

Step-by-Step Procedure

To limit the number of prefixes exported to OSPF:

1. Configure the prefix export limit value.



NOTE: For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf prefix-export-limit 100000
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
prefix-export-limit 100000;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Prefix Export Limit | 208](#)

Confirm that the configuration is working properly.

Verifying the Prefix Export Limit

Purpose

Verify the prefix export counter that displays the number of routes exported into OSPF.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Understanding OSPF Traffic Control

IN THIS SECTION

- [Controlling the Cost of Individual OSPF Network Segments | 209](#)
- [Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth | 210](#)
- [Controlling OSPF Route Preferences | 210](#)

Once a topology is shared across the network, OSPF uses the topology to route packets between network nodes. Each path between neighbors is assigned a cost based on interface throughput. The default algorithm computes the interface metric based on a reference bandwidth of 100 Mbps using the formula $\text{cost} = \text{reference-bandwidth} / \text{interface bandwidth}$. The result is any interface operating at 100 Mbps or faster is assigned the same metric value of 1. You can manually assign the OSPF interface metric to override the default value. Alternatively, given current Juniper platforms support interfaces that operate

at 400 Gbps, it's often a good idea to configure a larger reference-bandwidth value. Configuring a reference bandwidth value that is based on a multiple of the highest speed interface in your network automatically optimizes network paths based on interface speed and provides room for growth in network speeds.

The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest-path-first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion. Routes with lower total path metrics are preferred over those with higher path metrics.

You can use the following methods to control OSPF traffic:

- Control the cost of individual OSPF network segments
- Dynamically adjust OSPF interface metrics based on bandwidth
- Control OSPF route selection

Controlling the Cost of Individual OSPF Network Segments

OSPF uses the following formula to determine the cost of a route:

$$\text{cost} = \text{reference-bandwidth} / \text{interface bandwidth}$$

You can modify the reference-bandwidth value, which is used to calculate the default interface cost. The interface bandwidth value is not user-configurable and refers to the actual bandwidth of the physical interface.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface.

To control the flow of packets across the network, OSPF allows you to manually assign a cost (or metric) to a particular path segment. When you specify a metric for a specific OSPF interface, that value is used to determine the cost of routes advertised from that interface. For example, if all routers in the OSPF network use default metric values, and you increase the metric on one interface to 5, all paths through that interface have a calculated metric higher than the default and are not preferred.



NOTE: Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface.

When there are multiple equal-cost routes to the same destination in a routing table, an equal-cost multipath (ECMP) set is formed. If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose one of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. Define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the [edit **policy-options**] hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table.

Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes (for example, if the lag loses an interface member or if the interface speed is administratively changed), Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. Junos OS uses the smallest configured bandwidth threshold value that is equal to or greater than the actual interface bandwidth to determine the metric value. If the interface bandwidth is greater than any of the configured bandwidth threshold values, the metric value configured for the interface is used instead of any of the bandwidth-based metric values configured. The ability to recalculate the metric for an interface when its bandwidth changes is especially useful for aggregate interfaces.



NOTE: You must also configure a metric for the interface when you enable bandwidth-based metrics.

Controlling OSPF Route Preferences

You can control the flow of packets through the network using route preferences. Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. Although the default settings are appropriate for most environments, you might want to modify the default settings if all of the routing devices in your OSPF network use the default preference values, or if you are planning to migrate from OSPF to a different interior gateway protocol (IGP). If all of the devices use the default *route preference* values, you can change the route preferences to ensure that the path through a particular device is selected for the forwarding table any time multiple equal-cost paths to a destination exist. When migrating from OSPF to a different IGP, modifying the route preferences allows you to perform the migration in a controlled manner.

SEE ALSO

[OSPF Overview | 2](#)

[Example: Controlling OSPF Route Preferences | 235](#)

[Example: Configuring ECMP Flow-Based Forwarding](#)

Example: Controlling the Cost of Individual OSPF Network Segments

IN THIS SECTION

- [Requirements | 211](#)
- [Overview | 211](#)
- [Configuration | 213](#)
- [Verification | 216](#)

This example shows how to control the cost of individual OSPF network segments.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.

Overview

IN THIS SECTION

- [Topology | 213](#)

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred to those with higher path metrics. In this example, we explore how to control the cost of OSPF network segments.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface. This means that all interfaces faster than 100 Mbps have the same default cost metric of 1. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.


Having the same default metric might not be a problem if all of the interfaces are running at the same speed. If the interfaces operate at different speeds, you might notice that traffic is not routed over the fastest interface because OSPF equally routes packets across the different interfaces. For example, if your routing device has Fast Ethernet and Gigabit Ethernet interfaces running OSPF, each of these interfaces have a default cost metric of 1.

In the first example, you set the reference bandwidth to 10g (10 Gbps, as denoted by 10,000,000,000 bits) by including the **reference-bandwidth** statement. With this configuration, OSPF assigns the Fast Ethernet interface a default metric of 100, and the Gigabit Ethernet interface a metric of 10. Since the Gigabit Ethernet interface has the lowest metric, OSPF selects it when routing packets. The range is 9600 through 1,000,000,000,000 bits.

[Figure 19 on page 213](#) shows three routing devices in area 0.0.0.0 and assumes that the link between Device R2 and Device R3 is congested with other traffic. You can also control the flow of packets across the network by manually assigning a metric to a particular path segment. Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface. To prevent the traffic from Device R3 going directly to Device R2, you adjust the metric on the interface on Device R3 that connects with Device R1 so that all traffic goes through Device R1.

In the second example, you set the metric to 5 on interface **fe-1/0/1** on Device R3 that connects with Device R1 by including the **metric** statement. The range is 1 through 65,535.

Figure 19: OSPF Metric Configuration



9040888

Topology

Configuration

IN THIS SECTION

- [Configuring the Reference Bandwidth | 213](#)
- [Configuring a Metric for a Specific OSPF Interface | 215](#)

Configuring the Reference Bandwidth

CLI Quick Configuration

To quickly configure the reference bandwidth, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and

paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf reference-bandwidth 10g
```

Step-by-Step Procedure

To configure the reference bandwidth:

1. Configure the reference bandwidth to calculate the default interface cost.

 **NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf reference-bandwidth 10g
```

 **TIP:** As a shortcut in this example, you enter `10g` to specify 10 Gbps reference bandwidth. Whether you enter `10g` or `10000000000`, the output of `show protocols ospf` command displays 10 Gbps as `10g`, not `10000000000`.

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

 **NOTE:** Repeat this entire configuration on all routing devices in a shared network.

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
reference-bandwidth 10g;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Configuring a Metric for a Specific OSPF Interface

CLI Quick Configuration

To quickly configure a metric for a specific OSPF interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-1/0/1 metric 5
```

Step-by-Step Procedure

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-1/0/1 metric 5
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-1/0/1.0 {
    metric 5;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

IN THIS SECTION

- [Verifying the Configured Metric | 217](#)
- [Verifying the Route | 217](#)

Confirm that the configuration is working properly.

Verifying the Configured Metric

Purpose

Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

Action

From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Verifying the Route

Purpose

When choosing paths to a destination, OSPF uses the path with the lowest total cost. Confirm that OSPF is using the appropriate path.

Action

From operational mode, enter the **show route** command.

SEE ALSO

[Understanding OSPF Traffic Control | 208](#)

[Example: Controlling OSPF Route Preferences | 235](#)

Understanding Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors

Equal-cost multipath (ECMP) is a popular technique to load balance traffic across multiple paths. With ECMP enabled, if paths to a remote destination have the same cost, then traffic is distributed between them in equal proportion. Equal distribution of traffic across multiple paths is not desirable if the local links to adjacent routers towards the ultimate destination have unequal capacity. Typically the traffic distribution between two links is equal and the link utilization is the same. However, if the capacity of an aggregated Ethernet bundle changes, equal traffic distribution results in imbalance of link utilization. In

this case, weighted ECMP enables load balancing of traffic between equal cost paths in proportion to the capacity of the local links.

Taking as an example, there are two devices interconnected with an aggregated Ethernet bundle with four links and a single link of the same cost. Under normal conditions, both the AE bundles and the single link is utilized evenly to distribute traffic. However, if a link in the AE bundle goes down, there is a change in the link capacity that results in uneven link utilization. Weighted ECMP load balances traffic between the equal cost paths in proportion to the capacity of the local links. In this case, traffic is distributed in 30/40 proportion between the AE bundle and the single link.



NOTE: This feature provides weighted ECMP routing to OSPFv2 neighbors that are one hop away. The operating system supports this feature on immediately connected routers only and does not support weighted ECMP on multihop routers, that is, on routers that are more than one hop away.

To enable weighted ECMP traffic distribution on directly connected OSPFv2 neighbors, configure weighted one-hop statement at the `[edit protocols ospf spf-options multipath]` hierarchy level.



NOTE: You must configure per-packet load balancing policy before configuring this feature. WECMP will be operational if per-packet load balancing policy is in place.



NOTE: For logical interfaces, you must configure interface bandwidth to distribute traffic across equal cost multipaths based on the underlying physical interface bandwidth. If you do not configure the logical bandwidth for each logical interface, the operating system assumes that the entire bandwidth of the physical interface is available for each logical interface.

Example: Weighted ECMP Traffic Distribution on One-Hop OSPFv2 Neighbors

SUMMARY

Use this example to configure weighted equal cost multipath (ECMP) routing for distributing traffic to OSPFv2 neighbors that are one hop away to ensure optimal load balancing.

IN THIS SECTION

- [Example Prerequisites | 219](#)
- [Before You Begin | 219](#)

- [Functional Overview | 220](#)
- [Topology Overview | 220](#)
- [Topology Illustration | 221](#)
- [RO Configuration Steps | 221](#)
- [Verification | 224](#)
- [Appendix 1: Set Commands on All Devices | 229](#)



NOTE: Our content testing team has validated and updated this example.



TIP:
Table 2: Readability Score and Time Estimates

Reading Time	30 minutes
Configuration Time	20 minutes

Example Prerequisites

Hardware requirements	Two MX Series routers.
Software requirements	Junos OS Release 24.2R1 or later running on all devices.

Before You Begin

Benefits	Weighted ECMP routing distributes traffic unequally over multiple paths for better load balancing. It is more efficient than equal distribution of traffic during per-packet load balancing.
Know more	Understanding Weighted ECMP Traffic Distribution on One-Hop OSPF Neighbors

Functional Overview

Technologies used	<ul style="list-style-type: none"> • Ethernet Services: Aggregated Ethernet, VLAN • Routing Protocols: OSPF • Policy: Per-packet load balancing
Primary verification tasks	<ol style="list-style-type: none"> 1. Verify that OSPFv2 distributes traffic, by achieving weighted ECMP when one of the aggregated link is down during per-packet load balancing depending on the available bandwidth. 2. Verify that OSPFv2 distributes traffic, by achieving weighted ECMP on logical interfaces based on the configured logical bandwidth.


Topology Overview

This configuration example depicts three aggregated Ethernet bundles ae0, ae1, and ae2 with two links each configured between Router R0 and Router R1. The Packet Forwarding Engine distributes traffic unequally between the three Ethernet bundles when one of the links goes down, depending on the available bandwidth.

Hostname	Role	Function
R0	The device on which the WECMP is configured.	R0 sends traffic to R1.
R1	The device that is directly connected to R0.	R1 receives traffic from R0.

Topology Illustration

Figure 20: Weighted ECMP Traffic Distribution on One Hop OSPFv2 Neighbors



R0 Configuration Steps

For complete sample configurations on R0, see: "[Appendix 1: Set Commands on All Devices](#)" on page 229

This section highlights the main configuration tasks needed to configure the R0 device for this example. The first step is common to configuring the aggregated Ethernet interfaces. The following set of steps are specific to configuring OSPF on the AE bundles and configuring weighted ECMP.

1. a. Configure the two member links of the ae0, ae1, and ae2 aggregated Ethernet bundles.
- b. Configure IP address and the Link Aggregation Control Protocol (LACP) for the ae0, ae1, and ae2 aggregated Ethernet interfaces.
- c. Configure the aggregated Ethernet interfaces (ae0, ae1, and ae2) for vlan tagging.
- d. Configure the loopback interface address.
- e. Configure the OSPF router identifier by entering the [router-id] configuration value.
- f. Configure logical interfaces with appropriate bandwidth based on the underlying physical bandwidth.



NOTE: For logical interfaces, configure interface bandwidth to distribute traffic across equal-cost multipaths based on the underlying operational interface bandwidth. When you configure multiple logical interfaces on a single interface, configure appropriate logical bandwidth for each logical interface to see the desired traffic distribution over the logical interfaces.

- g. Configure a tunnel interface and specify the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine of RO.

```
[edit]
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae1
set interfaces ge-0/0/4 gigether-options 802.3ad ae1
set interfaces ge-0/0/5 gigether-options 802.3ad ae2
set interfaces ge-0/0/6 gigether-options 802.3ad ae2
```

```
[edit]
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 aggregated-ether-options lacp active
```

```
[edit]
set interfaces ae0 vlan-tagging
set interfaces ae1 vlan-tagging
set interfaces ae2 vlan-tagging
```

```
[edit]
set interfaces lo0 unit 0 family inet address 10.1.255.1/32
```

```
[edit]
set routing-options router-id 10.1.255.1
```

```
[edit]
set interfaces ae0 unit 0 vlan-id 6
set interfaces ae0 unit 0 family inet address 10.0.0.1/24
set interfaces ae1 unit 0 vlan-id 16
set interfaces ae1 unit 0 family inet address 10.0.1.1/24
```

```
set interfaces ae2 unit 0 vlan-id 26
set interfaces ae2 unit 0 family inet address 10.2.2.1/24
```

2. Specify the maximum number of weighted ECMP interfaces that you want to configure. Enable graceful switchover and specify the number of aggregated Ethernet interfaces to be created.

```
[edit]
set chassis maximum-ecmp 64
set chassis aggregated-devices ethernet device-count 3
```

3. Configure OSPF on all the interfaces and on the AE bundles.

```
[edit]
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

4. Configure per-packet load balancing.

```
[edit]
set policy-options policy-statement pplb then load-balance per-packet
```

5. Apply per-packet load balancing policy.

```
[edit]
set routing-options forwarding-table export ppl
```

6. Enable weighted ECMP traffic distribution on directly connected OSPFv2 neighbors.

```
[edit]
set protocols ospf spf-options multipath weighted one-hop
```

Verification

IN THIS SECTION

- [Verifying Equal Distribution of Traffic Over Equal-Cost Multiple Paths | 224](#)
- [Verifying Unequal Traffic Distribution Over Available Bandwidth | 228](#)

Command	Verification Task
show route extensive	Verify equal distribution of traffic over equal-cost multiple paths.
show route extensive	Verify unequal traffic distribution over available bandwidth.
show interfaces extensive	Verify unequal traffic distribution over available bandwidth.

Verifying Equal Distribution of Traffic Over Equal-Cost Multiple Paths

Purpose

To verify that traffic is equally distributed over the aggregated Ethernet bundles.

Action

From operational mode, enter the show route 10.1.255.2 extensive command.

```

user@R0> show route 10.1.255.2 extensive

inet.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
10.1.255.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.255.2/32 -> {list:10.0.0.2, 10.0.1.2, 10.2.2.2}
  *OSPF   Preference: 10
          Next hop type: Router, Next hop index: 0
          Address: 0x819a814
          Next-hop reference count: 2, Next-hop session id: 0
          Kernel Table Id: 0
          Next hop: 10.0.0.2 via ae0.0 weight 0x1 balance 33%
          Session Id: 0

```

```

Next hop: 10.0.1.2 via ae1.0 weight 0x1 balance 33%, selected
Session Id: 0
Next hop: 10.2.2.2 via ae2.0 weight 0x1 balance 33%
Session Id: 0
State: <Active Int>
Age: 4d 17:55:37      Metric: 1
Validation State: unverified
Area: 0.0.0.0
Task: OSPF
Announcement bits (1): 0-KRT
AS path: I
Thread: junos-main

```

```
user@R0> show interfaces ae0.0 extensive
```

```
Logical interface ae0.0 (Index 337) (SNMP ifIndex 578) (Generation 173)
```

```
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.6 ] Encapsulation: ENET2
```

Statistics	Packets	pps	Bytes	bps
------------	---------	-----	-------	-----

Bundle:

Input :	89241	0	7140674	0
Output:	89244	0	8731668	0

Adaptive Statistics:

Adaptive Adjusts:	0
Adaptive Scans :	0
Adaptive Updates:	0

Link:

ge-0/0/1.0

Input :	47583	0	3807058	0
Output:	0	0	0	0

ge-0/0/2.0

Input :	41632	0	3331512	0
Output:	89243	0	8731574	0

Aggregate member links: 2

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
--------------------	-----------	---------	------------	------------

ge-0/0/1.0	0	0	0	0
ge-0/0/2.0	0	0	0	0

Protocol inet, MTU: 1500

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 177, Route table: 0

```

Flags: Sendbcast-pkt-to-re, 0x0
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.0/24, Local: 10.0.0.1, Broadcast: 10.0.0.255, Generation: 157
Protocol multiservice, MTU: Unlimited, Generation: 178, Route table: 0
Flags: Is-Primary, 0x0
Policer: Input: __default_arp_policer__

```

user@R0> **show interfaces ae1.0 extensive**

```

Logical interface ae1.0 (Index 362) (SNMP ifIndex 593) (Generation 175)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.16 ] Encapsulation: ENET2

```

Statistics	Packets	pps	Bytes	bps
------------	---------	-----	-------	-----

Bundle:

Input :	89631	0	7194074	312
Output:	89626	1	8793864	784

Adaptive Statistics:

Adaptive Adjusts:	0
Adaptive Scans :	0
Adaptive Updates:	0

Link:

ge-0/0/3.0

Input :	89631	0	7194074	312
Output:	89626	0	8793864	0

ge-0/0/4.0

Input :	0	0	0	0
Output:	0	0	0	0

Aggregate member links: 2

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
--------------------	-----------	---------	------------	------------

ge-0/0/3.0	0	0	0	0
ge-0/0/4.0	0	0	0	0

Protocol inet, MTU: 1500

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 180, Route table: 0

Flags: Sendbcast-pkt-to-re, 0x0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.0.1/24, Local: 10.0.1.1, Broadcast: 10.0.1.255, Generation: 159

Protocol multiservice, MTU: Unlimited, Generation: 181, Route table: 0

Flags: 0x0

Policer: Input: __default_arp_policer__

user@R0> **show interfaces ae2.0 extensive**

Logical interface ae2.0 (Index 364) (SNMP ifIndex 592) (Generation 177)

Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.26] Encapsulation: ENET2

Statistics	Packets	pps	Bytes	bps
------------	---------	-----	-------	-----

Bundle:

Input :	89612	0	7193002	0
Output:	89664	0	8797828	0

Adaptive Statistics:

Adaptive Adjusts:	0
Adaptive Scans :	0
Adaptive Updates:	0

Link:

ge-0/0/5.0

Input :	89612	0	7193002	0
Output:	89664	0	8797828	0

ge-0/0/6.0

Input :	0	0	0	0
Output:	0	0	0	0

Aggregate member links: 2

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
--------------------	-----------	---------	------------	------------

ge-0/0/5.0	0	0	0	0
ge-0/0/6.0	0	0	0	0

Protocol inet, MTU: 1500

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 183, Route table: 0

Flags: Sendbroadcast-pkt-to-re, 0x0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.2.2/24, Local: 10.2.2.1, Broadcast: 10.2.2.255, Generation: 161

Protocol multiservice, MTU: Unlimited, Generation: 184, Route table: 0

Flags: 0x0

Policer: Input: __default_arp_policer__

Meaning

OSPF distributes traffic equally when the three aggregated Ethernet bundles have the same bandwidth available.

Verifying Unequal Traffic Distribution Over Available Bandwidth

Purpose

To verify that OSPF distributes traffic unevenly when one of the aggregated link is down during per-packet load balancing depending on the available bandwidth.

Action

Disable one of the links on the ae0 bundle. From operational mode, enter the `show route 10.1.255.2 extensive` command.

```

user@R0> show route 10.1.255.2 extensive
inet.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
10.1.255.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.255.2/32 -> {list:10.0.0.2, 10.0.1.2, 10.2.2.2}
  *OSPF   Preference: 10
          Next hop type: Router, Next hop index: 0
          Address: 0x819ba14
          Next-hop reference count: 2, Next-hop session id: 0
          Kernel Table Id: 0
          Next hop: 10.0.0.2 via ae0.0 weight 0x1 balance 20%
          Session Id: 0
          Next hop: 10.0.1.2 via ae1.0 weight 0x1 balance 40%, selected
          Session Id: 0
          Next hop: 10.2.2.2 via ae2.0 weight 0x1 balance 40%
          Session Id: 0
          State: <Active Int>
          Age: 23           Metric: 1
          Validation State: unverified
          Area: 0.0.0.0
          Task: OSPF
          Announcement bits (1): 0-KRT
          AS path: I
          Thread: junos-main

```

Meaning

OSPF infers that the ae0 bundle has lesser bandwidth available. Therefore, modifies per-packet load balancing according to the available bandwidth. As per the output, only 20 percent of the bandwidth is available on ae0 because one of the aggregated Ethernet links is down. Thus OSPF distributes traffic unequally depending on the available bandwidth.

Appendix 1: Set Commands on All Devices

IN THIS SECTION

● [R0 | 229](#)

● [R1 | 230](#)

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

R0

```
set system host-name R0
set chassis maximum-ecmp 64
set chassis aggregated-devices ethernet device-count 3
set interfaces ge-0/0/1 description "LinkID: R0R1-1"
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 description "LinkID: R0R1-2"
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 description "LinkID: R0R1-3"
set interfaces ge-0/0/3 gigether-options 802.3ad ae1
set interfaces ge-0/0/4 description "LinkID: R0R1-4"
set interfaces ge-0/0/4 gigether-options 802.3ad ae1
set interfaces ge-0/0/5 description "LinkID: R0R1-5"
set interfaces ge-0/0/5 gigether-options 802.3ad ae2
set interfaces ge-0/0/6 description "LinkID: R0R1-6"
set interfaces ge-0/0/6 gigether-options 802.3ad ae2
set interfaces ae0 vlan-tagging
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 vlan-id 6
```

```

set interfaces ae0 unit 0 family inet address 10.0.0.1/24
set interfaces ae1 vlan-tagging
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 vlan-id 16
set interfaces ae1 unit 0 family inet address 10.0.1.1/24
set interfaces ae2 vlan-tagging
set interfaces ae2 aggregated-ether-options minimum-links 2
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 vlan-id 26
set interfaces ae2 unit 0 family inet address 10.2.2.1/24
set interfaces lo0 unit 0 family inet address 10.1.255.1/32
set policy-options policy-statement pplb then load-balance per-packet
set routing-options router-id 10.1.255.1
set routing-options forwarding-table export pplb
set protocols ospf spf-options multipath weighted one-hop
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

R1

```

set system host-name R1
set chassis maximum-ecmp 64
set chassis aggregated-devices ethernet device-count 3
set interfaces ge-0/0/1 description "LinkID: R0R1-1"
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 description "LinkID: R0R1-2"
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 description "LinkID: R0R1-3"
set interfaces ge-0/0/3 gigether-options 802.3ad ae1
set interfaces ge-0/0/4 description "LinkID: R0R1-4"
set interfaces ge-0/0/4 gigether-options 802.3ad ae1
set interfaces ge-0/0/5 description "LinkID: R0R1-5"
set interfaces ge-0/0/5 gigether-options 802.3ad ae2
set interfaces ge-0/0/6 description "LinkID: R0R1-6"
set interfaces ge-0/0/6 gigether-options 802.3ad ae2
set interfaces ae0 vlan-tagging
set interfaces ae0 aggregated-ether-options minimum-links 1

```

```
set interfaces ae0 aggregated-ether-options lACP active
set interfaces ae0 unit 0 vlan-id 6
set interfaces ae0 unit 0 family inet address 10.0.0.2/24
set interfaces ae1 vlan-tagging
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options lACP active
set interfaces ae1 unit 0 vlan-id 16
set interfaces ae1 unit 0 family inet address 10.0.1.2/24
set interfaces ae2 vlan-tagging
set interfaces ae2 aggregated-ether-options minimum-links 2
set interfaces ae2 aggregated-ether-options lACP active
set interfaces ae2 unit 0 vlan-id 26
set interfaces ae2 unit 0 family inet address 10.2.2.2/24
set interfaces lo0 unit 0 family inet address 10.1.255.2/32
set routing-options router-id 10.1.255.2
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

IN THIS SECTION

- [Requirements | 233](#)
- [Overview | 234](#)
- [Verification | 234](#)

This example shows how to dynamically adjust OSPF interface metrics based on bandwidth.

Configuration

CLI Quick Configuration

To quickly configure bandwidth threshold values and associated metric values for an OSPF interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface ae0.0 metric 5
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 1g metric 60
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 10g metric 50
```

Step-by-Step Procedure

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0 metric 5
```

3. Configure the bandwidth threshold values and associated metric values. With this configuration when aggregated Ethernet interface's bandwidth is 1g, OSPF considers metric 60 for this interface. When aggregated Ethernet interface's bandwidth is 10g , OSPF considers metric 50 for this interface.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 1g metric 60
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 10g metric 50
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ae0.0 {
    bandwidth-based-metrics {
      bandwidth 1g metric 60;
      bandwidth 10g metric 50;
    }
    metric 5;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election"](#) on page 57
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network"](#) on page 62.

Overview

IN THIS SECTION

- [Topology | 234](#)

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. When you configure bandwidth-based metric values, you typically configure multiple bandwidth and metric values.

In this example, you configure OSPF interface **ae0** for bandwidth-based metrics by including the **bandwidth-based-metrics** statement and the following settings:

- **bandwidth**—Specifies the bandwidth threshold in bits per second. The range is 9600 through 1,000,000,000,000,000.
- **metric**—Specifies the metric value to associate with a specific bandwidth value. The range is 1 through 65,535.

Topology

Verification

IN THIS SECTION

- [Verifying the Configured Metric | 235](#)

Confirm that the configuration is working properly.

Verifying the Configured Metric

Purpose

Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

Action

From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

Example: Controlling OSPF Route Preferences

IN THIS SECTION

- [Requirements | 236](#)
- [Overview | 237](#)
- [Verification | 238](#)

This example shows how to control OSPF route selection in the forwarding table. This example also shows how you might control route selection if you are migrating from OSPF to another IGP.

Configuration

CLI Quick Configuration

To quickly configure the OSPF route preference values, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]  
set protocols ospf preference 168 external-preference 169
```

Step-by-Step Procedure

To configure route selection:

1. Enter OSPF configuration mode and set the external and internal routing preferences.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf preference 168 external-preference 169
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
preference 168;
external-preference 169;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Requirements

This example assumes that OSPF is properly configured and running in your network, and you want to control route selection because you are planning to migrate from OSPF to a different IGP.

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the IGP that you want to migrate to.

Overview

IN THIS SECTION

- [Topology | 237](#)

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. You might want to modify this setting if you are planning to migrate from OSPF to a different IGP. Modifying the route preferences enables you to perform the migration in a controlled manner.

This example makes the following assumptions:

- OSPF is already running in your network.
- You want to migrate from OSPF to IS-IS.
- You configured IS-IS per your network requirements and confirmed it is working properly.

In this example, you increase the OSPF route preference values to make them less preferred than IS-IS routes by specifying 168 for internal OSPF routes and 169 for external OSPF routes. IS-IS internal routes have a preference of either 15 (for Level 1) or 18 (for Level 2), and external routes have a preference of 160 (for Level 1) or 165 (for Level 2). In general, it is preferred to leave the new protocol at its default settings to minimize complexities and simplify any future addition of routing devices to the network. To modify the OSPF route preference values, configure the following settings:

- `preference`—Specifies the route preference for internal OSPF routes. By default, internal OSPF routes have a value of 10. The range is from 0 through 4,294,967,295 ($2^{32} - 1$).
- `external-preference`—Specifies the route preference for external OSPF routes. By default, external OSPF routes have a value of 150. The range is from 0 through 4,294,967,295 ($2^{32} - 1$).

Topology

Verification

IN THIS SECTION

- [Verifying the Route | 238](#)

Confirm that the configuration is working properly.

Verifying the Route

Purpose

Verify that the IGP is using the appropriate route. After the new IGP becomes the preferred protocol (in this example, IS-IS), you should monitor the network for any issues. After you confirm that the new IGP is working properly, you can remove the OSPF configuration from the routing device by entering the `delete ospf` command at the `[edit protocols]` hierarchy level.

Action

From operational mode, enter the `show route` command.

Understanding OSPF Overload Function

If the time elapsed after the OSPF instance is enabled is less than the specified timeout, overload mode is set.

You can configure the local routing device so that it appears to be overloaded. An overloaded routing device determines it is unable to handle any more OSPF transit traffic, which results in sending OSPF transit traffic to other routing devices. OSPF traffic to directly attached interfaces continues to reach the routing device. You might configure overload mode for many reasons, including:

- If you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic. This could include a routing device that is connected to the network for analysis purposes, but is not considered part of the production network, such as network management routing devices.
- If you are performing maintenance on a routing device in a production network. You can move traffic off that routing device so network services are not interrupted during your maintenance window.

You configure or disable overload mode in OSPF with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the OSPF instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. When the timer expires, overload mode is cleared. In overload mode, the router link-state advertisement (LSA) is originated with all the transit router links (except stub) set to a metric of 0xFFFF. The stub router links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and to take paths around the routing device. However, the overloaded routing device's own links are still accessible.

The routing device can also dynamically enter the overload state, regardless of configuring the device to appear overloaded. For example, if the routing device exceeds the configured OSPF prefix limit, the routing device purges the external prefixes and enters into an overload state.

In cases of incorrect configurations, the huge number of routes might enter OSPF, which can hamper the network performance. To prevent this, `prefix-export-limit` should be configured which will purge externals and prevent the network from the bad impact.

By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area. You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem.

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. To prevent this, `prefix-export-limit` should be configured which will purge externals and prevent the network.

Starting from Junos OS Release 18.2 onward, the following functionalities are supported by Stub Router in your OSPF network, when the OSPF is overloaded:

- Allow Route leaking—external prefixes are redistributed during OSPF overload and the prefixes are originated with normal cost.
- Advertise stub network with max metric—stub networks are advertised with maximum metric during OSPF overload.
- Advertise intra-area prefix with max metric—intra-area prefixes are advertised with maximum metric during OSPF overload.
- Advertise external prefix with max possible metric—OSPF AS external prefixes are redistributed during OSPF overload and the prefixes are advertised with maximum cost.

You can now configure the following when OSPF is overloaded:

- `allow-route-leaking` at the `[edit protocols <ospf | ospf3> overload]` hierarchy level to advertise the external prefixes with normal cost.

- `stub-network` at the `[edit protocols ospf overload]` hierarchy level to advertise stub network with maximum metric.
- `intra-area-prefix` at the `[edit protocols ospf3 overload]` hierarchy level to advertise intra-area prefix with maximum metric.
- `as-external` at the `[edit protocols <ospf | ospf3> overload]` hierarchy level to advertise external prefix with maximum metric.

To limit the number of prefixes exported to OSPF:

```
[edit]
set protocols ospf prefix-export-limit number
```

The prefix export limit number can be a value from 0 through 4,294,967,295.

SEE ALSO

[*overload*](#)

[*allow-route-leaking*](#)

[*stub-network*](#)

[*intra-area-prefix*](#)

[*as-external*](#)

Example: Configuring OSPF to Make Routing Devices Appear Overloaded

IN THIS SECTION

● [Requirements | 241](#)

● [Overview | 241](#)

● [Configuration | 242](#)

● [Verification | 243](#)

This example shows how to configure a routing device running OSPF to appear to be overloaded.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 241](#)

You can configure a local routing device running OSPF to appear to be overloaded, which allows the local routing device to participate in OSPF routing, but not for transit traffic. When configured, the transit interface metrics are set to the maximum value of 65535.

This example includes the following settings:

- **overload**—Configures the local routing device so it appears to be overloaded. You might configure this if you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic, or you are performing maintenance on a routing device in a production network.
- **timeout *seconds***—(Optional) Specifies the number of seconds at which the overload is reset. If no timeout interval is specified, the routing device remains in the overload state until the overload statement is deleted or a timeout is set. In this example, you configure 60 seconds as the amount of time the routing device remains in the overload state. By default, the timeout interval is 0 seconds (this value is not configured). The range is from 60 through 1800 seconds.

Topology

Configuration

IN THIS SECTION

- [Procedure | 242](#)

Procedure

CLI Quick Configuration

To quickly configure a local routing device to appear as overloaded, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf overload timeout 60
```

Step-by-Step Procedure

To configure a local routing device to appear overloaded:

1. Enter OSPF configuration mode.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf
```

2. Configure the local routing device to be overloaded.

```
[edit protocols ospf]
user@host# set overload
```

3. (Optional) Configure the number of seconds at which overload is reset.

```
[edit protocols ospf]
user@host# set overload timeout 60
```

4. (Optional) Configure the limit on the number prefixes exported to OSPF, to minimise the load on the routing device and prevent the device from entering the overload mode.

```
[edit protocols ospf]
user@host# set prefix-export-limit 50
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration. The output includes the optional `timeout` and `prefix-export-limit` statements.

```
user@host# show protocols ospf

prefix-export-limit 50;
overload timeout 60;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying Traffic Has Moved Off Devices | 244](#)
- [Verifying Transit Interface Metrics | 244](#)
- [Verifying the Overload Configuration | 244](#)

- [Verifying the Viable Next Hop | 245](#)

Confirm that the configuration is working properly.

Verifying Traffic Has Moved Off Devices

Purpose

Verify that the traffic has moved off the upstream devices.

Action

From operational mode, enter the `show interfaces detail` command.

Verifying Transit Interface Metrics

Purpose

Verify that the transit interface metrics are set to the maximum value of 65535 on the downstream neighboring device.

Action

From operational mode, enter the `show ospf database router detail advertising-router address` command for OSPFv2, and enter the `show ospf3 database router detail advertising-router address` command for OSPFv3.

Verifying the Overload Configuration

Purpose

Verify that overload is configured by reviewing the Configured overload field. If the overload timer is also configured, this field also displays the time that remains before it is set to expire.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and the `show ospf3 overview` command for OSPFv3.

Verifying the Viable Next Hop

Purpose

Verify the viable next hop configuration on the upstream neighboring device. If the neighboring device is overloaded, it is not used for transit traffic and is not displayed in the output.

Action

From operational mode, enter the `show route address` command.

Understanding the SPF Algorithm Options for OSPF

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to reach each destination. The SPF algorithm describes how OSPF determines the route to reach each destination, and the SPF options control the timers that dictate when the SPF algorithm runs. Depending on your network environment and requirements, you might want to modify the SPF options. For example, consider a large-scale environment with a large number of devices flooding link-state advertisements (LSAs) through out the area. In this environment, it is possible to receive a large number of LSAs to process, which can consume memory resources. By configuring the SPF options, you continue to adapt to the changing network topology, but you can minimize the amount of memory resources being used by the devices to run the SPF algorithm.

You can configure the following SPF options:

- The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.
- The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.
- The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured number of times. If the network stabilizes during the holddown period and the SPF algorithm does not need to run again, the system reverts to the configured values for the **delay** and `rapid-runs` statements.

Example: Configuring SPF Algorithm Options for OSPF

IN THIS SECTION

- [Requirements | 246](#)
- [Overview | 246](#)
- [Configuration | 247](#)
- [Verification | 249](#)

This example shows how to configure the SPF algorithm options. The SPF options control the timers that dictate when the SPF algorithm runs.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 247](#)

OSPF uses the SPF algorithm to determine the route to reach each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topology databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. The SPF options control the timers used by the SPF algorithm.

Before you modify any of the default settings, you should have a good understanding of your network environment and requirements.

This example shows how to configure the options for running the SPF algorithm. You include the `spf-options` statement and the following options:

- **delay**—Configures the amount of time (in milliseconds) between the detection of a topology and when the SPF actually runs. When you modify the delay timer, consider your requirements for network reconvergence. For example, you want to specify a timer value that can help you identify abnormalities in the network, but allow a stable network to reconverge quickly. By default, the SPF algorithm runs 200 milliseconds after the detection of a topology. The range is from 50 through 8000 milliseconds.
- **rapid-runs**—Configures the maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins. By default, the number of SPF calculations that can occur in succession is 3. The range is from 1 through 10. Each SPF algorithm is run after the configured SPF delay. When the maximum number of SPF calculations occurs, the hold-down timer begins. Any subsequent SPF calculation is not run until the hold-down timer expires.
- **holddown**—Configures the time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times. By default, the hold down time is 5000 milliseconds. The range is from 2000 through 20,000 milliseconds. If the network stabilizes during the holddown period and the SPF algorithm does not need to run again, the system reverts to the configured values for the **delay** and **rapid-runs** statements.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 248](#)
- [Procedure | 248](#)

CLI Quick Configuration

To quickly configure the SPF options, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf spf-options delay 210
set protocols ospf spf-options rapid-runs 4
set protocols ospf spf-options holddown 5050
```

Procedure

Step-by-Step Procedure

To configure the SPF options:

1. Enter OSPF configuration mode.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf
```

2. Configure the SPF delay time.

```
[edit protocols ospf]
user@host# set spf-options delay 210
```

3. Configure the maximum number of times that the SPF algorithm can run in succession.

```
[edit protocols ospf]
user@host# set spf-options rapid-runs 4
```

4. Configure the SPF hold-down timer.

```
[edit protocols ospf]
user@host# set spf-options holddown 5050
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
spf-options {
  delay 210;
  holddown 5050;
  rapid-runs 4;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying SPF Options | 250](#)

Confirm that the configuration is working properly.

Verifying SPF Options

Purpose

Verify that SPF is operating per your network requirements. Review the SPF delay field, the SPF holddown field, and the SPF rapid runs fields.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

The OSPF standard requires that every link-state advertisement (LSA) be refreshed every 30 minutes. The Juniper Networks implementation refreshes LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. This requirement can result in traffic overhead that makes it difficult to scale OSPF networks. You can override the default behavior by specifying that the DoNotAge bit be set in self-originated LSAs when they are initially sent by the router or switch. Any LSA with the DoNotAge bit set is reflooded only when a change occurs in the LSA. This feature thus reduces protocol traffic overhead while permitting any changed LSAs to be flooded immediately. Routers or switches enabled for flood reduction continue to send hello packets to their neighbors and to age self-originated LSAs in their databases.

The Juniper implementation of OSPF refresh and flooding reduction is based on RFC 4136, *OSPF Refresh and Flooding Reduction in Stable Topologies*. However, the Juniper implementation does not include the forced-flooding interval defined in the RFC. Not implementing the forced-flooding interval ensures that LSAs with the DoNotAge bit set are reflooded only when a change occurs.

This feature is supported for the following:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 realms
- OSPFv2 and OSPFv3 virtual links
- OSPFv2 sham links
- OSPFv2 peer interfaces
- All routing instances supported by OSPF

- Logical systems

To configure flooding reduction for an OSPF interface, include the flood-reduction statement at the [edit protocols (ospf | ospf3) area *area-id* interface *interface-id*] hierarchy level.



NOTE: If you configure flooding reduction for an interface configured as a demand circuit, the LSAs are not initially flooded, but sent only when their content has changed. Hello packets and LSAs are sent and received on a demand-circuit interface only when a change occurs in the network topology.

In the following example, the OSPF interface so-0/0/1.0 is configured for flooding reduction. As a result, all the LSAs generated by the routes that traverse the specified interface have the DoNotAge bit set when they are initially flooded, and LSAs are refreshed only when a change occurs.

```
[edit]
protocols ospf {
  area 0.0.0.0 {
    interface so-0/0/1.0 {
      flood-reduction;
    }
    interface lo0.0;
    interface so-0/0/0.0;
  }
}
```



NOTE: Beginning with Junos OS Release 12.2, you can configure a global default link-state advertisement (LSA) flooding interval in OSPF for self-generated LSAs by including the `lsa-refresh-interval minutes` statement at the [edit protocols (ospf | ospf3)] hierarchy level. The Juniper Networks implementation refreshes LSAs every 50 minutes. The range is 25 through 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes.

If you have both the global LSA refresh interval configured for OSPF and OSPF flooding reduction configured for a specific interface in an OSPF area, the OSPF flood reduction configuration takes precedence for that specific interface.

Understanding Synchronization Between LDP and IGPs

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the interior gateway protocol (IGP). If synchronization between LDP and the IGP is not maintained, the label-switch path (LSP) goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

SEE ALSO

[Example: Configuring Synchronization Between LDP and OSPF | 252](#)

[Junos OS MPLS Applications User Guide](#)

Example: Configuring Synchronization Between LDP and OSPF

IN THIS SECTION

- [Requirements | 252](#)
- [Overview | 253](#)
- [Configuration | 254](#)
- [Verification | 257](#)

This example shows how to configure synchronization between LDP and OSPFv2.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62.](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)

Overview

IN THIS SECTION

- [Topology | 253](#)

In this example, configure synchronization between LDP and OSPFv2 by performing the following tasks:

- Enable LDP on interface **so-1/0/3**, which is a member of OSPF area 0.0.0.0, by including the `ldp` statement at the `[edit protocols]` hierarchy level. You can configure one or more interfaces. By default, LDP is disabled on the routing device.
- Enable LDP synchronization by including the `ldp-synchronization` statement at the `[edit protocols ospf area area-id interface interface-name]` hierarchy level. This statement enables LDP synchronization by advertising the maximum cost metric until LDP is operational on the link.
- Configure the amount of time (in seconds) the routing device advertises the maximum cost metric for a link that is not fully operational by including the `hold-time` statement at the `[edit protocols ospf area area-id interface interface-name ldp-synchronization]` hierarchy level. If you do not configure the `hold-time` statement, the `hold-time` value defaults to infinity. The range is from 1 through 65,535 seconds. In this example, configure 10 seconds for the `hold-time` interval.

This example also shows how to disable synchronization between LDP and OSPFv2 by including the `disable` statement at the `[edit protocols ospf area area-id interface interface-name ldp-synchronization]` hierarchy level.

Topology

Configuration

IN THIS SECTION

- [Enabling Synchronization Between LDP and OSPFv2 | 254](#)
- [Disabling Synchronization Between LDP and OSPFv2 | 256](#)

Enabling Synchronization Between LDP and OSPFv2

CLI Quick Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in [CLI User Guide](#).

To quickly enable synchronization between LDP and OSPFv2, copy the following commands, remove any line breaks, and then paste them into the CLI.

```
[edit]
set protocols ldp interface so-1/0/3
set protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization hold-time 10
```

Step-by-Step Procedure

To enable synchronization between LDP and OSPFv2:

1. Enable LDP on the interface.

```
[edit]
user@host# set protocols ldp interface so-1/0/3
```

2. Configure LDP synchronization and optionally configure a time period of 10 seconds to advertise the maximum cost metric for a link that is not fully operational.

```
[edit ]
user@host# edit protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization
```

3. Configure a time period of 10 seconds to advertise the maximum cost metric for a link that is not fully operational.

```
[edit protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization ]
user@host# set hold-time 10
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ldp` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ldp
interface so-1/0/3.0;
```

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface so-1/0/3.0 {
    ldp-synchronization {
      hold-time 10;
    }
  }
}
```

Disabling Synchronization Between LDP and OSPFv2

CLI Quick Configuration

To quickly disable synchronization between LDP and OSPFv2, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization disable
```

Step-by-Step Procedure

To disable synchronization between LDP and OSPF:

1. Disable synchronization by including the disable statement.

```
[edit ]
user@host# set protocols ospf area 0.0.0.0 interface so-1/0/3 ldp-synchronization disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface so-1/0/3.0 {
    ldp-synchronization {
      disable;
    }
  }
}
```

Verification

IN THIS SECTION

- [Verifying the LDP Synchronization State of the Interface | 257](#)

Confirm that the configuration is working properly.

Verifying the LDP Synchronization State of the Interface

Purpose

Verify the current state of LDP synchronization on the interface. The LDP sync state displays information related to the current state, and the config holdtime field displays the configured hold-time interval.

Action

From operational mode, enter the `show ospf interface extensive` command.

OSPFv2 Compatibility with RFC 1583 Overview

By default, the Junos OS implementation of OSPFv2 is compatible with RFC 1583, *OSPF Version 2*. This means that Junos OS maintains a single best route to an autonomous system (AS) boundary router in the OSPF routing table, rather than multiple intra-AS paths, if they are available. You can now disable compatibility with RFC 1583. It is preferable to do so when the same external destination is advertised by AS boundary routers that belong to different OSPF areas. When you disable compatibility with RFC 1583, the OSPF routing table maintains the multiple intra-AS paths that are available, which the router uses to calculate AS external routes as defined in RFC 2328, *OSPF Version 2*. Being able to use multiple available paths to calculate an AS external route can prevent routing loops.

SEE ALSO

| [Example: Disabling OSPFv2 Compatibility with RFC 1583 | 258](#)

Example: Disabling OSPFv2 Compatibility with RFC 1583

IN THIS SECTION

- [Requirements | 258](#)
- [Overview | 258](#)
- [Configuration | 259](#)
- [Verification | 260](#)

This example shows how to disable OSPFv2 compatibility with RFC 1583 on the routing device.

Requirements

No special configuration beyond device initialization is required before disabling OSPFv2 compatibility with RFC 1583.

Overview

IN THIS SECTION

- [Topology | 258](#)

By default, the Junos OS implementation of OSPF is compatible with RFC 1583. This means that Junos OS maintains a single best route to an autonomous system (AS) boundary router in the OSPF routing table, rather than multiple intra-AS paths, if they are available. You can disable compatibility with RFC 1583. It is preferable to do so when the same external destination is advertised by AS boundary routers that belong to different OSPF areas. When you disable compatibility with RFC 1583, the OSPF routing table maintains the multiple intra-AS paths that are available, which the router uses to calculate AS external routes as defined in RFC 2328. Being able to use multiple available paths to calculate an AS external route can prevent routing loops. To minimize the potential for routing loops, configure the same RFC compatibility on all OSPF devices in an OSPF domain.

Topology

Configuration

IN THIS SECTION

- Procedure | 259
- Results | 260

Procedure

CLI Quick Configuration

To quickly disable OSPFv2 compatibility with RFC 1583, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode. You configure this setting on all devices that are part of the OSPF domain.

```
[edit]  
set protocols ospf no-rfc-1583
```

Step-by-Step Procedure

To disable OSPFv2 compatibility with RFC 1583:

1. Disable RFC 1583.

```
[edit]  
user@host# set protocols ospf no-rfc-1583
```

2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```



NOTE: Repeat this configuration on each routing device that participates in an OSPF routing domain.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-rfc-1583;
```

Verification

IN THIS SECTION

- [Verifying the OSPF Routes | 260](#)

Confirm that the configuration is working properly.

Verifying the OSPF Routes

Purpose

Verify that the OSPF routing table maintains the intra-AS paths with the largest metric, which the router uses to calculate AS external routes.

Action

From operational mode, enter the `show ospf route detail` command.

RELATED DOCUMENTATION

[OSPF Overview | 2](#)

[Understanding OSPF Configurations | 14](#)

6

CHAPTER

Configure OSPF Authentication

IN THIS CHAPTER

- [Configuring OSPF Authentication | 262](#)
-

Configuring OSPF Authentication

IN THIS SECTION

- [Understanding IPsec Authentication for OSPF Packets on EX Series Switches | 262](#)
- [Understanding OSPFv2 Authentication | 265](#)
- [Understanding OSPFv3 Authentication | 268](#)
- [Example: Configuring Simple Authentication for OSPFv2 Exchanges | 270](#)
- [Example: Configuring MD5 Authentication for OSPFv2 Exchanges | 273](#)
- [Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface | 277](#)
- [Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\) | 282](#)
- [Example: Configuring IPsec Authentication for an OSPF Interface | 284](#)

Understanding IPsec Authentication for OSPF Packets on EX Series Switches

IN THIS SECTION

- [Authentication Algorithms | 263](#)
- [Encryption Algorithms | 263](#)
- [IPsec Protocols | 264](#)
- [Security Associations | 264](#)
- [IPsec Modes | 264](#)

IP Security (IPsec) provides a secure way to authenticate senders and encrypt IP version 4 (IPv4) traffic between network devices. IPsec offers network administrators for Juniper Networks EX Series Ethernet Switches and their users the benefits of data confidentiality, data integrity, sender authentication, and anti-replay services.

IPsec is a framework for ensuring secure private communication over IP networks and is based on standards developed by the International Engineering Task Force (IETF). IPsec provides security services at the network layer of the Open Systems Interconnection (OSI) model by enabling a system to select required security protocols, determine the algorithms to use for the security services, and implement any cryptographic keys required to provide the requested services. You can use IPsec to protect one or more paths between a pair of hosts, between a pair of security gateways (such as switches), or between a security gateway and a host.

OSPF version 3 (OSPFv3), unlike OSPF version 2 (OSPFv2), does not have a built-in authentication method and relies on IPsec to provide this functionality. You can secure specific OSPFv3 interfaces and protect OSPFv3 virtual links.

Authentication Algorithms

Authentication is the process of verifying the identity of the sender. Authentication algorithms use a shared key to verify the authenticity of the IPsec devices. The Juniper Networks Junos operating system (Junos OS) uses the following authentication algorithms:

- Message Digest 5 (MD5) uses a one-way hash function to convert a message of arbitrary length to a fixed-length message digest of 128 bits. Because of the conversion process, it is mathematically infeasible to calculate the original message by computing it backwards from the resulting message digest. Likewise, a change to a single character in the message will cause it to generate a very different message digest number.

To verify that the message has not been tampered with, Junos OS compares the calculated message digest against a message digest that is decrypted with a shared key. Junos OS uses the MD5 hashed message authentication code (HMAC) variant that provides an additional level of hashing. MD5 can be used with an authentication header (AH) and Encapsulating Security Payload (ESP).

- Secure Hash Algorithm 1 (SHA-1) uses a stronger algorithm than MD5. SHA-1 takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest ensures that the data has not been changed and that it originates from the correct source. Junos OS uses the SHA-1 HMAC variant that provides an additional level of hashing. SHA-1 can be used with AH, ESP, and Internet Key Exchange (IKE).

Encryption Algorithms

Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. As with authentication algorithms, a shared key is used with encryption algorithms to verify the authenticity of IPsec devices. Junos OS uses the following encryption algorithms:

- Data Encryption Standard cipher-block chaining (DES-CBC) is a symmetric secret-key block algorithm. DES uses a key size of 64 bits, where 8 bits are used for error detection and the remaining 56 bits provide encryption. DES performs a series of simple logical operations on the shared key,

including permutations and substitutions. CBC takes the first block of 64 bits of output from DES, combines this block with the second block, feeds this back into the DES algorithm, and repeats this process for all subsequent blocks.

- Triple DES-CBC (3DES-CBC) is an encryption algorithm that is similar to DES-CBC but provides a much stronger encryption result because it uses three keys for 168-bit (3 x 56-bit) encryption. 3DES works by using the first key to encrypt the blocks, the second key to decrypt the blocks, and the third key to reencrypt the blocks.

IPsec Protocols

IPsec protocols determine the type of authentication and encryption applied to packets that are secured by the switch. Junos OS supports the following IPsec protocols:

- AH—Defined in *RFC 2402*, AH provides connectionless integrity and data origin authentication for IPv4. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. In an IP header, AH can be identified with a value of 51 in the Protocol field of an IPv4 packet.
- ESP—Defined in *RFC 2406*, ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service. In an IP header, ESP can be identified with a value of 50 in the Protocol field of an IPv4 packet.

Security Associations

An IPsec consideration is the type of security association (SA) that you wish to implement. An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA can be either unidirectional or bidirectional, depending on the choices made by the network administrator. An SA is uniquely identified by a Security Parameter Index (SPI), an IPv4 or IPv6 destination address, and a security protocol (AH or ESP) identifier.

IPsec Modes

Junos OS supports the following IPsec modes:

- Tunnel mode is supported for both AH and ESP in Junos OS. In tunnel mode, the SA and associated protocols are applied to tunneled IPv4 or IPv6 packets. For a tunnel mode SA, an outer IP header specifies the IPsec processing destination and an inner IP header specifies the ultimate destination for the packet. The security protocol header appears after the outer IP header and before the inner

IP header. In addition, there are slight differences for tunnel mode when you implement it with AH and ESP:

- For AH, portions of the outer IP header are protected, as well as the entire tunneled IP packet.
- For ESP, only the tunneled packet is protected, not the outer header.

When one side of an SA is a security gateway (such as a switch), the SA must use tunnel mode. However, when traffic (for example, SNMP commands or BGP sessions) is destined for a switch, the system acts as a host. Transport mode is allowed in this case because the system does not act as a security gateway and does not send or receive transit traffic.



NOTE: Tunnel mode is not supported for OSPF v3 control packet authentication.

- Transport mode provides an SA between two hosts. In transport mode, the protocols provide protection primarily for upper-layer protocols. A transport mode security protocol header appears immediately after the IP header and any options and before any higher-layer protocols (for example, TCP or UDP). There are slight differences for transport mode when you implement it with AH and ESP:
 - For AH, selected portions of the IP header are protected, as well as selected portions of the extension headers and selected options within the IPv4 header.
 - For ESP, only the higher-layer protocols are protected, not the IP header or any extension headers preceding the ESP header.

Understanding OSPFv2 Authentication

All OSPFv2 protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in the autonomous system's routing. By default, OSPFv2 authentication is disabled.



NOTE: OSPFv3 does not have a built-in authentication method and relies on IP Security (IPsec) to provide this functionality.

You can enable the following authentication types:

- Simple authentication—Authenticates by using a plain-text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet.

- MD5 authentication—Authenticates by using an encoded MD5 checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet.

You define an MD5 key for each interface. If MD5 is enabled on an interface, that interface accepts routing updates only if MD5 authentication succeeds. Otherwise, updates are rejected. The routing device only accepts OSPFv2 packets sent using the same key identifier (ID) that is defined for that interface.

- IPsec authentication (beginning with Junos OS Release 8.3)—Authenticates OSPFv2 interfaces, the remote endpoint of a sham link, and the OSPFv2 virtual link by using manual security associations (SAs) to ensure that a packet's contents are secure between the routing devices. You configure the actual IPsec authentication separately.



NOTE: You can configure IPsec authentication together with either MD5 or simple authentication.

The following restrictions apply to IPsec authentication for OSPFv2:

- Dynamic Internet Key Exchange (IKE) SAs are not supported.
- Only IPsec transport mode is supported. Tunnel mode is not supported.
- Because only bidirectional manual SAs are supported, all OSPFv2 peers must be configured with the same IPsec SA. You configure a manual bidirectional SA at the [edit security ipsec] hierarchy level.
- You must configure the same IPsec SA for all virtual links with the same remote endpoint address, for all neighbors on OSPF nonbroadcast multiaccess (NBMA) or point-to-multipoint links, and for every subnet that is part of a broadcast link.
- OSPFv2 peer interfaces are not supported.

Because OSPF performs authentication at the area level, all routing devices within the area must have the same authentication and corresponding password (key) configured. For MD5 authentication to work, both the receiving and transmitting routing devices must have the same MD5 key. In addition, a simple password and MD5 key are mutually exclusive. You can configure only one simple password, but multiple MD5 keys.

As part of your security measures, you can change MD5 keys. You can do this by configuring multiple MD5 keys, each with a unique key ID, and setting the date and time to switch to the new key. Each unique MD5 key has a unique ID. The ID is used by the receiver of the OSPF packet to determine which key to use for authentication. The key ID, which is required for MD5 authentication, specifies the identifier associated with the MD5 key.

Starting in Junos OS Release 22.4R1, we support advertising OSPF MD5 authentication with multiple active keys to send packets with a maximum limit of two keys per interface. Having multiple keys active at any one time at the interface enables the smooth transition from one key to another for OSPF. You can delete old keys without any impact on the OSPF session.

Starting in Junos OS Release 23.3R1 and Junos OS Evolved Release 23.3R1, you can enable OSPFv2 HMAC-SHA1 authentication with keychain to authenticate packets reaching or originating from an OSPF interface. This ensures smooth transition from one key to another for OSPFv2 with enhanced security. You can enable OSPFv2 to send packets authenticated with only the latest MD5 key once all the neighbours switch to the latest configured key. Prior to this release, we support advertising authenticated OSPF packets always with multiple active MD5 keys with a maximum limit of two keys per interface.

HMAC-SHA1 authentication for OSPFv2 does not support:

- Keychain with no active keys.
- Migration from other existing authentication types to keychain with hitless session.
- Migration from no authentication to keychain with hitless session.
- Keyed MD5 as a part of Keychain configuration.



NOTE:

- When multi-active MD5 optimization with `delete-if-not-inuse` configuration statement is enabled and once the authentication negotiation happens with its neighbors, from then on the device uses only active key for transmission for that negotiated neighbor. That is, we do not support rolling back to the old key.

For example: R0 and R1 are configured with key-id 1 as `delete-if-not-inuse` and key-id 2. Later, if R1 is configured to remove key-id 2, then R0 does not rollback to using both keys (key-id 1 and key-id 2) for transmission.

- Keychain activeness is based on the absolute time (wall clock) and the wall clock might go backward after commit. This type of error does not reflect at the time of commit. So it is important to have the system time synchronized on all the devices when keychain is active on a OSPF session.

Starting in Junos OS Evolved Release 24.2R1, you can enable OSPFv2 keychain module with HMAC-SHA2 (OSPFv2 HMAC-SHA2) authentication to authenticate packets reaching or originating from an OSPF interface. HMAC SHA2 algorithms include HMAC-SHA2-256, HMAC-SHA2-384, and HMAC-SHA2-512 as defined in RFC 5709. We support these algorithms along with HMAC-SHA2-224. This feature ensures smooth transition from one key to another for OSPFv2 with enhanced security. We also support HMAC-SHA1 and HMAC-SHA2 authentication for virtual and sham links.

HMAC-SHA2 authentication for OSPFv2 does not support:

- Keychain with no active keys.
- Migration from other existing authentication types to keychain with hitless session.
- Migration from no authentication to keychain with hitless session.



NOTE:

- SHA1 (without HMAC) algorithm under keychain configuration is not supported.
- It is important to have the system time synchronized on all the devices when keychain is active on a OSPF session.

SEE ALSO

[Overview of IPsec](#)

authentication (Protocols OSPF)

Understanding OSPFv3 Authentication

OSPFv3 does not have a built-in authentication method and relies on the IP Security (IPsec) suite to provide this functionality. IPsec provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. You can use IPsec to secure specific OSPFv3 interfaces and protect OSPFv3 virtual links.



NOTE:

You configure the actual IPsec authentication separately from your OSPFv3 configuration and then apply IPsec to the OSPFv3 interfaces or OSPFv3 virtual links.

OSPFv3 uses the IP authentication header (AH) and the IP Encapsulating Security Payload (ESP) portions of the IPsec Protocol to authenticate routing information between peers. AH can provide connectionless integrity and data origin authentication. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service.

IPsec is based on security associations (SAs). An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. This simplex connection provides security services to the packets carried by the SA. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA is used to encrypt and authenticate a particular flow in one direction. Therefore, in normal bidirectional traffic, the flows are secured by a pair of SAs. An SA to be used with OSPFv3 must be configured manually and use transport mode. Static values must be configured on both ends of the SA.

Manual SAs require no negotiation between the peers. All values, including the keys, are static and specified in the configuration. Manual SAs statically define the security parameter index (SPI) values, algorithms, and keys to be used and require matching configurations on both end points (OSPFv3 peers). As a result, each peer must have the same configured options for communication to take place.

The actual choice of encryption and authentication algorithms is left to your IPsec administrator; however, we have the following recommendations:

- Use ESP with NULL encryption to provide authentication to the OSPFv3 protocol headers only. With NULL encryption, you are choosing not to provide encryption on OSPFv3 headers. This can be useful for troubleshooting and debugging purposes. For more information about NULL encryption, see RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*.
- Use ESP with non-NULL encryption for full confidentiality. With non-NULL encryption, you are choosing to provide encryption. For more information about NULL encryption, see RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*.
- Use AH to provide authentication to the OSPFv3 protocol headers, portions of the IPv6 header, and portions of the extension headers.

The following restrictions apply to IPsec authentication for OSPFv3:

- Dynamic Internet Key Exchange (IKE) security associations (SAs) are not supported.
- Only IPsec transport mode is supported. In transport mode, only the payload (the data you transfer) of the IP packet is encrypted and/or authenticated. Tunnel mode is not supported.
- Because only bidirectional manual SAs are supported, all OSPFv3 peers must be configured with the same IPsec SA. You configure a manual bidirectional SA at the [edit security ipsec] hierarchy level.
- You must configure the same IPsec SA for all virtual links with the same remote endpoint address.

SEE ALSO

| [Overview of IPsec](#)

Example: Configuring Simple Authentication for OSPFv2 Exchanges

IN THIS SECTION

- Requirements | 270
- Overview | 270
- Configuration | 271
- Verification | 272

This example shows how to enable simple authentication for OSPFv2 exchanges.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#) or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

Simple authentication uses a plain-text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet. Plain-text passwords are not encrypted and might be subject to packet interception. This method is the least secure and should only be used if network security is not your goal.

You can configure only one simple authentication key (password) on the routing device. The simple key can be from 1 through 8 characters and can include ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

In this example, you specify OSPFv2 interface **so-0/1/0** in area 0.0.0.0, set the authentication type to simple-password, and define the key as PssWd4.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 271](#)
- [Procedure | 271](#)
- [Results | 272](#)

CLI Quick Configuration

To quickly configure simple authentication, copy the following command, removing any line breaks, and then paste the command into the CLI. You must configure all routing devices within the area with the same authentication and corresponding password.

```
[edit]  
set protocols ospf area 0.0.0.0 interface so-0/1/0 authentication simple-password PssWd4
```

Procedure

Step-by-Step Procedure

To enable simple authentication for OSPFv2 exchanges:

1. Create an OSPF area.

```
[edit]  
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]  
user@host# edit interface so-0/1/0
```

3. Set the authentication type and the password.

```
[edit protocols ospf area 0.0.0.0 interface so-0/1/0.0]
user@host# set authentication simple-password PssWd4
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 interface so-0/1/0.0]
user@host# commit
```



NOTE: Repeat this entire configuration on all peer OSPFv2 routing devices in the area.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.



NOTE: After you configure the password, you do not see the password itself. The output displays the encrypted form of the password you configured.

```
user@host# show protocols ospf
  area 0.0.0.0 {
    interface so-0/1/0.0 {
      authentication {
        simple-password "$9$-3dY4ZUHm5FevX-db2g"; ## SECRET-DATA
      }
    }
  }
```

Verification

IN THIS SECTION

- [Verifying the Configured Authentication Method | 273](#)

Confirm that the configuration is working properly.

Verifying the Configured Authentication Method

Purpose

Verify that the authentication method for sending and receiving OSPF protocol packets is configured. The Authentication Type field displays Password when configured for simple authentication.

Action

From operational mode, enter the `show ospf interface` and the `show ospf overview` commands.

Example: Configuring MD5 Authentication for OSPFv2 Exchanges

IN THIS SECTION

- [Requirements | 273](#)
- [Overview | 274](#)
- [Configuration | 274](#)
- [Verification | 276](#)

This example shows how to enable MD5 authentication for OSPFv2 exchanges.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#) or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57

- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62.](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)

Overview

IN THIS SECTION

- [Topology | 274](#)

MD5 authentication uses an encoded MD5 checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet.

You define an MD5 key for each interface. If MD5 is enabled on an interface, that interface accepts routing updates only if MD5 authentication succeeds. Otherwise, updates are rejected. The routing device only accepts OSPFv2 packets sent using the same key identifier (ID) that is defined for that interface.

In this example, you create the backbone area (area 0.0.0.0), specify OSPFv2 interface **so-0/2/0**, set the authentication type to md5, and then define the authentication key ID as 5 and the password as PssWd8.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 275](#)
- [Procedure | 275](#)
- [Results | 276](#)

CLI Quick Configuration

To quickly configure MD5 authentication, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface so-0/2/0 authentication md5 5 key PssWd8
```

Procedure

Step-by-Step Procedure

To enable MD5 authentication for OSPFv2 exchanges:

1. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# edit interface so-0/2/0
```

3. Configure MD5 authentication and set a key ID and an authentication password.

```
[edit protocols ospf area 0.0.0.0 interface s0-0/2/0.0]
user@host# set authentication md5 5 key PssWd8
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 interface s0-0/2/0.0]
user@host# commit
```



NOTE: Repeat this entire configuration on all peer OSPFv2 routing devices.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.



NOTE: After you configure the password, you do not see the password itself. The output displays the encrypted form of the password you configured.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface so-0/2/0.0 {
    authentication {
      md5 5 key "$9$pXXhuIhreWx-wQF9puBEh"; ## SECRET-DATA
    }
  }
}
```

Verification

IN THIS SECTION

- [Verifying the Configured Authentication Method | 276](#)

Confirm that the configuration is working properly.

Verifying the Configured Authentication Method

Purpose

Verify that the authentication method for sending and receiving OSPF protocol packets is configured. When configured for MD5 authentication, the Authentication Type field displays MD5, the Active key ID field displays the unique number you entered that identifies the MD5 key, and the Start time field displays the date as Start time 1970 Jan 01 00:00:00 PST. Do not be alarmed by this start time. This is the default start time that the routing device displays if the MD5 key is effective immediately.

Action

From operational mode, enter the `show ospf interface` and the `show ospf overview` commands.

Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface

IN THIS SECTION

- [Requirements | 277](#)
- [Overview | 278](#)
- [Configuration | 279](#)
- [Verification | 281](#)

This example shows how to configure a transition of MD5 keys on an OSPFv2 interface.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#) or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 278](#)

MD5 authentication uses an encoded MD5 checksum that is included in the transmitted packet. For MD5 authentication to work, both the receiving and transmitting routing devices must have the same MD5 key.

You define an MD5 key for each interface. If MD5 is enabled on an interface, that interface accepts routing updates only if MD5 authentication succeeds. Otherwise, updates are rejected. The routing device only accepts OSPFv2 packets sent using the same key identifier (ID) that is defined for that interface.

For increased security, you can configure multiple MD5 keys, each with a unique key ID, and set the date and time to switch to a new key. The receiver of the OSPF packet uses the ID to determine which key to use for authentication.

In this example, you configure new keys to take effect at 12:01 AM on the first day of the next three months on OSPFv2 interface **fe-0/0/1** in the backbone area (area 0.0.0.0), and you configure the following MD5 authentication settings:

- **md5**—Specifies the MD5 authentication key ID. The key ID can be set to any value between 0 and 255, with a default value of 0. The routing device only accepts OSPFv2 packets sent using the same key ID that is defined for that interface.
- **key**—Specifies the MD5 key. Each key can be a value from 1 through 16 characters long. Characters can include ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
- **start-time**—Specifies the time to start using the MD5 key. This option enables you to configure a smooth transition mechanism for multiple keys. The start time is relevant for transmission but not for receiving OSPF packets.



NOTE: You must set the same passwords and transition dates and times on all devices in the area so that OSPFv2 adjacencies remain active.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 279](#)
- [Procedure | 279](#)
- [Results | 281](#)

CLI Quick Configuration

To quickly configure multiple MD5 keys on an OSPFv2 interface, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/1/0 authentication md5 1 key $2010HaL
set protocols ospf area 0.0.0.0 interface fe-0/1/0 authentication md5 2 key NeWpsswdFEB start-
time 2011-02-01.00:01
set protocols ospf area 0.0.0.0 interface fe-0/1/0 authentication md5 3 key NeWpsswdMAR start-
time 2011-03-01.00:01
set protocols ospf area 0.0.0.0 interface fe-0/1/0 authentication md5 4 key NeWpsswdAPR start-
time 2011-04-01.00:01
```

Procedure

Step-by-Step Procedure

To configure multiple MD5 keys on an OSPFv2 interface:

1. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# edit interface fe-0/1/0
```

3. Configure MD5 authentication and set an authentication password and key ID.

```
[edit protocols ospf area 0.0.0.0 interface fe-0/1/0.0]
user@host# set authentication md5 1 key $2010HaL
```

4. Configure a new key to take effect at 12:01 AM on the first day of February, March, and April.

You configure a new authentication password and key ID for each month.

- a. For the month of February, enter the following:

```
[edit protocols ospf area 0.0.0.0 interface fe-0/1/0.0]
user@host# set authentication md5 2 key NeWpsswdFEB start-time 2011-02-01.00:01
```

- b. For the month of March, enter the following:

```
[edit protocols ospf area 0.0.0.0 interface fe-0/1/0.0]
user@host# set authentication md5 3 key NeWpsswdMAR start-time 2011-03-01.00:01
```

- c. For the month of April, enter the following:

```
[edit protocols ospf area 0.0.0.0 interface fe-0/1/0.0]
user@host# set authentication md5 4 key NeWpsswdAPR start-time 2011-04-01.00:01
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 interface fe-0/1/0.0]
user@host# commit
```



NOTE: Repeat this entire configuration on all peer OSPFv2 routing devices.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.



NOTE: After you configure the password, you do not see the password itself. The output displays the encrypted form of the password you configured.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/1/0.0 {
    authentication {
      md5 1 key "$9$wzs24JGDjk.2gfTQ3CAp0B1hy"; ## SECRET-DATA
      md5 2 key "$9$Q9gz39t1IcML7EcwgJZq.RhSy1MN-b4oZDi" start-time "2011-2-1.00:01:00
-0800"; ## SECRET-DATA
      md5 3 key "$9$zjo2nCpIRSWXNhSs4ZG.mEcyreW2gaZGjCt" start-time "2011-3-1.00:01:00
-0800"; ## SECRET-DATA
      md5 4 key "$9$fQn900ReML1Rds4oiHBIEhSevMLXNVqm" start-time "2011-4-1.00:01:00
-0700"; ## SECRET-DATA
    }
  }
}
```

Verification

IN THIS SECTION

- [Verifying the Configured Authentication Method | 281](#)

Confirm that the configuration is working properly.

Verifying the Configured Authentication Method

Purpose

Verify that the authentication method for sending and receiving OSPF protocol packets is configured. When configured for MD5 authentication with a transition of keys, the Auth type field displays MD5,

the Active key ID field displays the unique number you entered that identifies the MD5 key, and the Start time field displays the time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface you configured.

Action

From operational mode, enter the `show ospf interface` and the `show ospf overview` commands.

Using IPsec to Secure OSPFv3 Networks (CLI Procedure)

IN THIS SECTION

- [Configuring Security Associations | 282](#)
- [Securing OPSFv3 Networks | 283](#)

OSPF version 3 (OSPFv3) does not have a built-in authentication method and relies on IP Security (IPsec) to provide this functionality. You can use IPsec to secure OSPFv3 interfaces on EX Series switches.

This topic includes:

Configuring Security Associations

When you configure a security association (SA), include your choices for authentication, encryption, direction, mode, protocol, and security parameter index (SPI).

To configure a security association:

1. Specify a name for the security association:

```
[edit security ipsec]
user@switch# set security-association sa-name
```

2. Specify the mode of the security association:

```
[edit security ipsec security-association sa-name]
user@switch# set mode transport
```

3. Specify the type of security association:

```
[edit security ipsec security-association sa-name]
user@switch# set type manual
```

4. Specify the direction of the security association:

```
[edit security ipsec security-association sa-name]
user@switch# set direction bidirectional
```

5. Specify the value of the security parameter index:

```
[edit security ipsec security-association sa-name]
user@switch# set spi spi-value
```

6. Specify the type of authentication to be used:

```
[edit security ipsec security-association sa-name]
user@switch# set authentication algorithm type
```

7. Specify the encryption algorithm and key:

```
[edit security ipsec security-association sa-name]
user@switch# set encryption algorithm algorithm key type
```

Securing OSPFv3 Networks

You can secure the OSPFv3 network by applying the SA to the OSPFv3 configuration.

To secure the OSPFv3 network:

```
[edit protocols ospf3 area area-number interface interface-name]
user@switch# set ipsec-sa sa-name
```

Example: Configuring IPsec Authentication for an OSPF Interface

IN THIS SECTION

- Requirements | 284
- Overview | 284
- Configuration | 287
- Verification | 291

This example shows how to enable IP Security (IPsec) authentication for an OSPF interface.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#) or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- Topology | 287

You can use IPsec authentication for both OSPFv2 and OSPFv3. You configure the actual IPsec authentication separately and apply it to the applicable OSPF configuration.

OSPFv2

Beginning with Junos OS Release 8.3, you can use IPsec authentication to authenticate OSPFv2 interfaces, the remote endpoint of a sham link, and the OSPFv2 virtual link by using manual security associations (SAs) to ensure that a packet's contents are secure between the routing devices.



NOTE: You can configure IPsec authentication together with either MD5 or simple authentication.

To enable IPsec authentication, do one of the following:

- For an OSPFv2 interface, include the `ipsec-sa name` statement for a specific interface:

```
interface interface-name ipsec-sa name;
```

- For a remote sham link, include the `ipsec-sa name` statement for the remote end point of the sham link:

```
sham-link-remote address ipsec-sa name;
```



NOTE: If a Layer 3 VPN configuration has multiple sham links with the same remote endpoint IP address, you must configure the same IPsec security association for all the remote endpoints. You configure a Layer 3 VPN at the `[edit routing-instances routing-instance-name instance-type]` hierarchy level. For more information about Layer 3 VPNs, see the [Junos OS VPNs Library for Routing Devices](#).

- For a virtual link, include the `ipsec-sa name` statement for a specific virtual link:

```
virtual-link neighbor-id router-id transit-area area-id ipsec-sa name;
```

OSPFv3

OSPFv3 does not have a built-in authentication method and relies on IPsec to provide this functionality. You use IPsec authentication to secure OSPFv3 interfaces and protect OSPFv3 virtual links by using manual SAs to ensure that a packet's contents are secure between the routing devices.

To apply authentication, do one of the following:

- For an OSPFv3 interface, include the `ipsec-sa name` statement for a specific interface:

```
interface interface-name ipsec-sa name;
```

- For a virtual link, include the `ipsec-sa name` statement for a specific virtual link:

```
virtual-link neighbor-id router-id transit-area area-id ipsec-sa name;
```

Tasks to Complete for Both OSPFv2 and OSPFv3

In this example, you perform the following tasks:

1. Configure IPsec authentication. To do this, define a manual SA named **sa1** and specify the processing direction, the protocol used to protect IP traffic, the security parameter index (SPI), and the authentication algorithm and key.

- a. Configure the following option at the [edit security ipsec security-association *sa-name* mode] hierarchy level:

transport—Specifies transport mode. This mode protects traffic when the communication endpoint and the cryptographic endpoint are the same. The data portion of the IP packet is encrypted, but the IP header is not.

- b. Configure the following option at the [edit security ipsec security-association *sa-name* manual direction] hierarchy level:

bidirectional—Defines the direction of IPsec processing. By specifying bidirectional, the same algorithms, keys, and security parameter index (SPI) values you configure are used in both directions.

- c. Configure the following options at the [edit security ipsec security-association *sa-name* manual direction bidirectional] hierarchy level:

protocol—Defines the IPsec protocol used by the manual SA to protect IP traffic. You can specify either the authentication header (AH) or the Encapsulating Security Payload (ESP). If you specify AH, which you do in this example, you cannot configure encryption.

spi—Configures the SPI for the manual SA. An SPI is an arbitrary value that uniquely identifies which SA to use at the receiving host. The sending host uses the SPI to identify and select which SA to use to secure every packet. The receiving host uses the SPI to identify and select the encryption algorithm and key used to decrypt packets. In this example, you specify 256.

authentication—Configures the authentication algorithm and key. The **algorithm** option specifies the hash algorithm that authenticates packet data. In this example, you specify **hmac-md5-96**,

which produces a 128-bit digest. The **key** option indicates the type of authentication key. In this example, you specify **ascii-text-key**, which is 16 ASCII characters for the **hmac-md5-96** algorithm.

2. Enable IPsec authentication on OSPF interface **so-0/2/0.0** in the backbone area (area 0.0.0.0) by including the name of the manual SA **sa1** that you configured at the [edit security ipsec] hierarchy level.

Topology

Configuration

IN THIS SECTION

- [Configuring Security Associations | 287](#)
- [Enabling IPsec Authentication for an OSPF Interface | 289](#)

Configuring Security Associations

CLI Quick Configuration

To quickly configure a manual SA to be used for IPsec authentication on an OSPF interface, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set security ipsec security-association sa1
set security ipsec security-association sa1 mode transport
set security ipsec security-association sa1 manual direction bidirectional
set security ipsec security-association sa1 manual direction bidirectional protocol ah
set security ipsec security-association sa1 manual direction bidirectional spi 256
set security ipsec security-association sa1 manual direction bidirectional authentication
algorithm hmac-md5-96 key ascii-text 123456789012abcd
```

Step-by-Step Procedure

To configure a manual SA to be used on an OSPF interface:

1. Specify a name for the SA.

```
[edit]
user@host# edit security ipsec security-association sa1
```

2. Specify the mode of the SA.

```
[edit security ipsec security-association sa1 ]
user@host# set mode transport
```

3. Configure the direction of the manual SA.

```
[edit security ipsec security-association sa1 ]
user@host# set manual direction bidirectional
```

4. Configure the IPsec protocol to use.

```
[edit security ipsec security-association sa1 ]
user@host# set manual direction bidirectional protocol ah
```

5. Configure the value of the SPI.

```
[edit security ipsec security-association sa1 ]
user@host# set manual direction bidirectional spi 256
```

6. Configure the authentication algorithm and key.

```
[edit security ipsec security-association sa1 ]
user@host# set manual direction bidirectional authentication algorithm hmac-md5-96 key ascii-
text 123456789012abcd
```

7. If you are done configuring the device, commit the configuration.

```
[edit security ipsec security-association sa1 ]
user@host# commit
```



NOTE: Repeat this entire configuration on all peer OSPF routing devices.

Results

Confirm your configuration by entering the `show security ipsec` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.



NOTE: After you configure the password, you do not see the password itself. The output displays the encrypted form of the password you configured.

```
user@host# show security ipsec
security-association sa1 {
  mode transport;
  manual {
    direction bidirectional {
      protocol ah;
      spi 256;
      authentication {
        algorithm hmac-md5-96;
        key ascii-text "$9$AP5Hp1RcylMLxSygoZUHK1REhKMVwY2oJx7jHq.zF69A00R"; ## SECRET-
DATA
      }
    }
  }
}
```

Enabling IPsec Authentication for an OSPF Interface

CLI Quick Configuration

To quickly apply a manual SA used for IPsec authentication to an OSPF interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface so-0/2/0 ipsec-sa sa1
```

Step-by-Step Procedure

To enable IPsec authentication for an OSPF interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# edit interface so-0/2/0
```

3. Apply the IPsec manual SA.

```
[edit protocols ospf area 0.0.0.0 interface so-0/2/0.0]
user@host# set ipsec-sa sa1
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 interface so-0/2/0.0]
user@host# commit
```



NOTE: Repeat this entire configuration on all peer OSPF routing devices.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface so-0/2/0.0 {
    ipsec-sa sa1;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the IPsec Security Association Settings | 291](#)
- [Verifying the IPsec Security Association on the OSPF Interface | 292](#)

Confirm that the configuration is working properly.

Verifying the IPsec Security Association Settings

Purpose

Verify the configured IPsec security association settings. Verify the following information:

- The Security association field displays the name of the configured security association.
- The SPI field displays the value you configured.
- The Mode field displays transport mode.
- The Type field displays manual as the type of security association.

Action

From operational mode, enter the `show ipsec security-associations` command.

Verifying the IPsec Security Association on the OSPF Interface

Purpose

Verify that the IPsec security association that you configured has been applied to the OSPF interface. Confirm that the IPsec SA name field displays the name of the configured IPsec security association.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

SEE ALSO

[Junos OS Administration Library for Routing Devices](#)

[Junos OS Services Interfaces Library for Routing Devices](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
22.4R1	

7

CHAPTER

Configure OSPF Routing Instances

IN THIS CHAPTER

- [Configuring OSPF Routing Instances | 294](#)
 - [Multiple Independent IGP Instances of OSPFv2 | 306](#)
 - [Example: Configure Multiple Independent Instances of OSPFv2 with Segment Routing | 309](#)
-

Configuring OSPF Routing Instances

IN THIS SECTION

- [Understanding OSPF Routing Instances | 294](#)
- [Installing Routes from OSPF Routing Instances into the OSPF Routing Table Group | 296](#)
- [Example: Configuring Multiple Routing Instances of OSPF | 297](#)

Understanding OSPF Routing Instances

IN THIS SECTION

- [Minimum Routing-Instance Configuration for OSPFv2 | 295](#)
- [Minimum Routing-Instance Configuration for OSPFv3 | 295](#)
- [Multiple Routing Instances of OSPF | 296](#)

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the OSPF routing protocol parameters control the information in the routing tables. You can further install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.



NOTE: The default routing instance, primary, refers to the main **inet.0** routing table. The primary routing instance is reserved and cannot be specified as a routing instance.

You can configure the following types of routing instances:

- OSPFv2—Forwarding, Layer 2 virtual private network (VPN), nonforwarding, VPN routing and forwarding (VRF), virtual router, and virtual private LAN service (VPLS).
- OSPFv3—Nonforwarding, VRF, and virtual router.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, the corresponding IP unicast table is **my-instance.inet.0**. All routes for **my-instance** are installed into **my-instance.inet.0**.

You can also configure multiple routing instances of OSPF.

Minimum Routing-Instance Configuration for OSPFv2

To configure a routing instance for OSPFv2, you must include at least the following statements in the configuration:

```
[edit]
routing-instances {
  routing-instance-name {
    interface interface-name;
    instance-type (forwarding | l2vpn | no-forwarding | virtual-router | vpls | vrf);
    route-distinguisher (as-number:number | ip-address:number);
    vrf-import [ policy-names ];
    vrf-export [ policy-names ];
    protocols {
      ospf {
        ... ospf-configuration ...
      }
    }
  }
}
```



NOTE: You can configure a *logical interface* under only one routing instance.

Minimum Routing-Instance Configuration for OSPFv3

To configure a routing instance for OSPFv3, you must include at least the following statements in the configuration:

```
[edit]
routing-instances {
  routing-instance-name {
    interface interface-name;
    instance-type (no-forwarding | virtual-router | vrf);
```

```

vrf-import [ policy-names ];
vrf-export [ policy-names ];
protocols {
    ospf3 {
        ... ospf3-configuration ...
    }
}
}
}

```



NOTE: You can configure a logical interface under only one routing instance.

Multiple Routing Instances of OSPF

Multiple instances of OSPF are used for Layer 3 VPN implementations. The multiple instances of OSPF keep routing information for different VPNs separate. The VRF instance advertises routes from the customer edge (CE) router to the provider edge (PE) router and advertises routes from the PE router to the CE router. Each VPN receives only routing information belonging to that VPN.

You can create multiple instances of OSPF by including statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* (ospf | ospf3)]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* (ospf | ospf3)]

Installing Routes from OSPF Routing Instances into the OSPF Routing Table Group

To install routes learned from OSPF routing instances into routing tables in the OSPF routing table group, include the `rib-group` statement:

```
rib-group group-name;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Example: Configuring Multiple Routing Instances of OSPF

IN THIS SECTION

- [Requirements | 297](#)
- [Overview | 297](#)
- [Configuration | 300](#)
- [Verification | 306](#)

This example shows how to configure multiple routing instances of OSPF.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57

Overview

IN THIS SECTION

- [Topology | 299](#)

When you configure multiple routing instances of OSPF, we recommend that you perform the following tasks:

1. Configure the OSPFv2 or OSPFv3 default instance at the [edit protocols (ospf | ospf3)] and [edit logical-systems *logical-system-name* protocols (ospf | ospf3)] hierarchy levels with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

2. Configure an OSPFv2 or OSPFv3 routing instance for each additional OSPFv2 or OSPFv3 routing entity, configuring the following:
 - Interfaces
 - Routing options
 - OSPF protocol statements belonging to that entity
 - Routing table group
3. Configure a routing table group to install routes from the default route table, **inet.0**, into a routing instance's route table.
4. Configure a routing table group to install routes from a routing instance into the default route table, **inet.0**.




NOTE: Nonforwarding routing instances do not have forwarding tables that correspond to their routing tables.

5. Create an export policy to export routes with a specific tag, and use that tag to export routes back into the instances. For more information, see the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#).

Figure 21 on page 299 shows how you can use multiple routing instances of OSPFv2 or OSPFv3 to segregate prefixes within a large network. The network consists of three administrative entities: **voice-policy**, **other-policy**, and the default routing instance. Each entity is composed of several geographically separate sites that are connected by the backbone and managed by the backbone entity.

Topology

Figure 21: Configuration for Multiple Routing Instances



Sites A and D belong to the **voice-policy** routing instance. Sites B and C belong to the **other-policy** instance. Device 1 and Device 3 at the edge of the backbone connect the routing instances. Each runs a separate OSPF or OSPFv3 instance (one per entity).

Device 1 runs three OSPFv2 or OSPFv3 instances: one each for Site A (**voice-policy**), Site C (**other-policy**), and the backbone, otherwise known as the default instance. Device 3 also runs three OSPFv2 or OSPFv3 instances: one each for Site B (**other-policy**), Site D (**voice-policy**), and the backbone (default instance).

When Device 1 runs the OSPFv2 or OSPFv3 instances, the following occur:

- Routes from the default instance routing table are placed in the voice-policy and other-policy instance routing tables.
- Routes from the voice-policy routing instance are placed in the default instance routing table.
- Routes from the other-policy routing instance are placed in the default instance routing table.
- Routes from the voice-policy routing instance do not enter the other-policy instance routing table.
- Routes from the other-policy routing instance do not enter the voice-policy instance routing table.

Configuration

IN THIS SECTION

- [Procedure | 300](#)

Procedure

CLI Quick Configuration

To quickly configure multiple routing instances of OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Configuration on Device 1:

```
[edit]
set routing-instances voice-policy interface so-2/2/2
set routing-instances voice-policy protocols ospf rib-group voice-to-inet area 0.0.0.0 interface
so-2/2/2
set routing-instances other-policy interface so-4/2/2
set routing-instances other-policy protocols ospf rib-group other-to-inet area 0.0.0.0 interface
so-4/2/2
set routing-options rib-groups inet-to-voice-and-other import-rib [ inet.0 voice-policy.inet.0
other-policy.inet.0 ]
set routing-options rib-groups voice-to-inet import-rib [ voice-policy.inet.0 inet.0 ]
set routing-options rib-groups other-to-inet import-rib [ other-policy.inet.0 inet.0 ]
set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-2/2/2
set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-4/2/2
```

Configuration on Device 3:

```
[edit]
set routing-instances voice-policy interface so-3/2/2
set routing-instances voice-policy protocols ospf rib-group voice-to-inet area 0.0.0.0 interface
so-3/2/2
set routing-instances other-policy interface so-5/2/2
```

```

set routing-instances other-policy protocols ospf rib-group other-to-inet area 0.0.0.0 interface
so-5/2/2
set routing-options rib-groups inet-to-voice-and-other import-rib [ inet.0 voice-policy.inet.0
other-policy.inet.0 ]
set routing-options rib-groups voice-to-inet import-rib [ voice-policy.inet.0 inet.0 ]
set routing-options rib-groups other-to-inet import-rib [ other-policy.inet.0 inet.0 ]
set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-3/2/2
set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-5/2/2

```

Step-by-Step Procedure

To configure multiple routing instances of OSPF:

1. Configure the routing instances for **voice-policy** and **other-policy**.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit routing-instances protocols] hierarchy level.

```

[edit]
user@D1# set routing-instances voice-policy interface so-2/2/2
user@D1# set routing-instances voice-policy protocols ospf rib-group voice-to-inet area
0.0.0.0 interface so-2/2/2
user@D1# set routing-instances other-policy interface so-4/2/2
user@D1# set routing-instances other-policy protocols ospf rib-group other-to-inet area
0.0.0.0 interface so-4/2/2

```

```

[edit]
user@D3# set routing-instances voice-policy interface so-3/2/2
user@D3# set routing-instances voice-policy protocols ospf rib-group voice-to-inet area
0.0.0.0 interface so-3/2/2
user@D3#set routing-instances other-policy interface so-5/2/2
user@D3# set routing-instances other-policy protocols ospf rib-group other-to-inet area
0.0.0.0 interface so-5/2/2

```

2. Configure the routing table group **inet-to-voice-and-other** to take routes from **inet.0** (default routing table) and place them in the **voice-policy.inet.0** and **other-policy.inet.0** routing tables.

```
[edit]
user@D1# set routing-options rib-groups inet-to-voice-and-other import-rib [ inet.0 voice-
policy.inet.0 other-policy.inet.0 ]
```

```
[edit]
user@D3# set routing-options rib-groups inet-to-voice-and-other import-rib [ inet.0 voice-
policy.inet.0 other-policy.inet.0 ]
```

3. Configure the routing table group **voice-to-inet** to take routes from **voice-policy.inet.0** and place them in the **inet.0** default routing table.

```
[edit]
user@D1# set routing-options rib-groups voice-to-inet import-rib [ voice-policy.inet.0
inet.0 ]
```

```
[edit]
user@D3# set routing-options rib-groups voice-to-inet import-rib [ voice-policy.inet.0
inet.0 ]
```

4. Configure the routing table group **other-to-inet** to take routes from **other-policy.inet.0** and place them in the **inet.0** default routing table.

```
[edit]
user@D1# set routing-options rib-groups other-to-inet import-rib [ other-policy.inet.0
inet.0 ]
```

```
[edit]
user@D3# set routing-options rib-groups other-to-inet import-rib [ other-policy.inet.0
inet.0 ]
```

5. Configure the default OSPF instance.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit routing-instances protocols]` hierarchy level.

```
[edit]
user@D1# set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-2/2/2
user@D1# set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-4/2/2
```

```
[edit]
user@D3# set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-3/2/2
user@D3# set protocols ospf rib-group inet-to-voice-and-other area 0.0.0.0 interface so-5/2/2
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-instances`, `show routing-options`, and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on Device 1:

```
user@D1# show routing-instances
voice-policy {
  interface so-2/2/2.0;
  protocols {
    ospf {
      rib-group voice-to-inet;
      area 0.0.0.0 {
        interface so-2/2/2.0;
      }
    }
  }
}
other-policy {
```

```

interface so-4/2/2.0;
protocols {
    ospf {
        rib-group other-to-inet;
        area 0.0.0.0 {
            interface so-4/2/2.0;
        }
    }
}
}

```

```

user@D1# show routing-options
rib-groups {
    inet-to-voice-and-other {
        import-rib [ inet.0 voice-policy.inet.0 other-policy.inet.0 ];
    }
    voice-to-inet {
        import-rib [ voice-policy.inet.0 inet.0 ];
    }
    other-to-inet {
        import-rib [ other-policy.inet.0 inet.0 ];
    }
}

```

```

user@D1# show protocols ospf
rib-group inet-to-voice-and-other;
area 0.0.0.0 {
    interface so-2/2/2.0;
    interface so-4/2/2.0;
}

```

Configuration on Device 3:

```

user@D3# show routing-instances
voice-policy {
    interface so-3/2/2.0;
    protocols {
        ospf {
            rib-group voice-to-inet;

```

```

        area 0.0.0.0 {
            interface so-3/2/2.0;
        }
    }
}
other-policy {
    interface so-5/2/2.0;
    protocols {
        ospf {
            rib-group other-to-inet;
            area 0.0.0.0 {
                interface so-5/2/2.0;
            }
        }
    }
}
}

```

```

user@D3# show routing-options
rib-groups {
    inet-to-voice-and-other {
        import-rib [ inet.0 voice-policy.inet.0 other-policy.inet.0 ];
    }
    voice-to-inet {
        import-rib [ voice-policy.inet.0 inet.0 ];
    }
    other-to-inet {
        import-rib [ other-policy.inet.0 inet.0 ];
    }
}

```

```

user@D3# show protocols ospf
rib-group inet-to-voice-and-other;
area 0.0.0.0 {
    interface so-3/2/2.0;
    interface so-5/2/2.0;
}

```

To confirm your OSPFv3 configuration, enter the `show routing-instances`, `show routing-options`, and `show protocols ospf3` commands.

Verification

IN THIS SECTION

- [Verifying the Routing Instances | 306](#)

Confirm that the configuration is working properly.

Verifying the Routing Instances

Purpose

Verify the configured routing instance settings.

Action

From operational mode, enter the `show route instance detail` command.

SEE ALSO

| [rib-group \(Protocols OSPF\)](#)

RELATED DOCUMENTATION

| [Routing Instances Overview](#)

Multiple Independent IGP Instances of OSPFv2

IN THIS SECTION

- [Benefits of Multi-Instance OSPFv2 | 307](#)
- [Multi-Instance OSPF Overview | 307](#)

Benefits of Multi-Instance OSPFv2

- You can use multiple IGP instances of OSPFv2 and redistribute routes among independent OSPFv2 domains on a single router.
- You can construct flexible OSPFv2 hierarchies across independent IGP domains.
- Allows decoupling of multiple OSPFv2 flooding domains and therefore achieve a more scalable OSPFv2 deployment.

Figure 22: Multi-Instance OSPFv2 Deployment Topology




Figure 22 on page 307 illustrates several benefits of configuring multiple IGP instances of OSPFv2 on the router. For example, Router F participates in two independent OSPF instances. Router F treats OSPF Aggregation Network-1 and OSPF Core Network as two independent IGP domains, while at the same time redistributing routes between those domains. Network operators can use this flexibility to construct a hierarchy of OSPF domains.

Figure 22 on page 307 also illustrates the use of multiple IGP instances of OSPF to separate metro networks into independent OSPF flooding domains. In the example, routers D and E participate in the OSPF metro-a, OSPF metro-b, and OSPF metro-c networks, as well as in OSPF Aggregation Network-1. Routers D and E do not flood the different OSPF domains with OSPF advertisements. Instead they redistribute specific routes among the different OSPF domains, which allows for more scalable metro deployments.

Multi-Instance OSPF Overview

You can configure and run multiple independent IGP instances of OSPFv2 simultaneously on a router. These instances are associated with the default routing instance, and they install routes in the default routing table. Each OSPF instance can also export the routes installed in the routing table by other

OSPF instances using the standard Junos OS routing policy configuration. By default, the routes installed by the different OSPF instances have the same route preference.



NOTE: Junos OS does not support configuring the same logical interface in multiple IGP instances of OSPF.

In most deployment scenarios, only one OSPF instance on a router installs a route for a given prefix. Therefore, you don't need to configure different route preferences for multiple OSPF instances. However, for certain deployment scenarios where multiple OSPF instances install the routes for the same prefix in the routing table, you can set a different route preference for the routes installed by other OSPF instances. This allows the routing table to choose the routes with the best route preference and installs those routes in the forwarding table.

You can use the multiple OSPF instance feature for both hierarchical and parallel deployments. In the case of hierarchical deployments, there are well-defined borders between the groups of routers participating in different IGP instances. In parallel deployments, different IGP instances (typically not more than two or three) span entire groups of routers. You can also have mixed deployments, with some domains in a hierarchical deployment running IGP instances in parallel.

You can configure multiple independent IGP instances of OSPFv2 by including the `ospf-instance` configuration statement at the `[edit protocols]` hierarchy level. The configuration statements that you use at the `[edit protocols ospf-instance igp-instance-name]` hierarchy level are the same as those available at the `[edit protocols ospf]` hierarchy level.



NOTE: The `ospf-instance` configuration statement is not supported at the `[edit routing-instances routing-instance-name protocols]` hierarchy level.

You can configure and run multiple independent interior gateway protocol (IGP) instances of OSPFv2 with segment routing (SR) on a router. You can create two or more OSPF instances and apply SR-MPLS on each instance. Multiple instances of OSPF can advertise different prefix-segment identifiers (prefix-SIDs). Other instances can use these SIDs for making routing decisions.

Multi-instance OSPF combined with SR enhances network flexibility, scalability, and control over traffic engineering, especially in large and complex networks.

Example: Configure Multiple Independent Instances of OSPFv2 with Segment Routing

SUMMARY

Use this example to configure multiple IGP instances of OSPFv2 with segment routing.

IN THIS SECTION

- [Example Prerequisites | 309](#)
- [Before You Begin | 310](#)
- [Functional Overview | 310](#)
- [Topology Overview | 311](#)
- [Topology Illustration | 311](#)
- [R2 Configuration Steps | 311](#)
- [Verification | 315](#)
- [Appendix 1: Set Commands on All Devices | 321](#)



NOTE: Our content testing team has validated and updated this example.



TIP:
Table 3: Readability Score and Time Estimates

Reading Time	30 minutes
Configuration Time	20 minutes

Example Prerequisites

Hardware requirements	Three MX Series routers.
Software requirements	Junos OS Release 24.4R1 or later running on all devices.

Before You Begin

Benefits	Configuring multiple independent instances of OSPFv2 with segment routing enhances network flexibility, scalability, and control over traffic engineering, especially in large and complex networks.
Know more	Multiple Independent IGP Instances of OSPFv2

Functional Overview

Technologies used	<ul style="list-style-type: none"> • Routing Protocols:OSPF • Segment Routing with Multiprotocol Label Switching (SR-MPLS) • VLAN Tagging
Primary verification tasks	<ul style="list-style-type: none"> • Verify that multiple independent OSPF instances are running. • Verify the OSPF segment routing database for different prefix-SIDs advertised by the multiple instances of OSPF


Topology Overview

This configuration example depicts three devices R1, R2, and R3. There are two sub-interfaces configured between device R1 and device R2 and between device R2 and device R3. Each device runs multiple OSPF instances with segment routing enabled. We configure SR-MPLS to provide path control through the network. There are OSPF instances named `metro-edge` running on each of the two subinterfaces of the devices.

Hostname	Role	Function
R1, R2, and R3	The devices have multi-instance OSPF configured on the subinterfaces, with segment routing enabled.	The devices participate in OSPF multi-instances, advertise routes, and forward traffic using prefix-SIDs to other devices.

Topology Illustration

Figure 23:



Router	Loopback Address
R1	10.1.255.1/32
R2	10.1.255.2/32
R3	10.1.255.3/32

IGP Instance: Metro Edge

jin-001203

R2 Configuration Steps

For complete sample configurations on R2, see: ["Appendix 1: Set Commands on All Devices" on page 321](#)

This section highlights the main configuration tasks needed to configure the R2 device for this example.

1.
 - a. Configure the basic device settings such as hostname, enhanced-ip mode, IPv4 addresses on the logical units of the device interfaces.
 - b. Configure the loopback interface with an IP address and enable MPLS.
 - c. Configure the router ID and autonomous system (AS) number to propagate routing information within a set of routing devices that belong to the same AS.
 - d. Enable VLAN tagging and configure the logical units of both the interfaces with different VLAN IDs.
 - e. Enable MPLS on each logical unit. Configure the maximum number of MPLS labels that can be applied to outgoing packets on logical units of each interface.
 - f. Define a policy to load balance packets and apply the per-packet policy to enable load balancing of traffic.
 - g. Configure a policy statement that matches routes based on the exact prefix and assign a segment identifier to the matched route.
 - h. Configure MPLS traffic engineering, segment routing global block (SRGB) label range at the edit protocol `mpls` hierarchy level to ensure the labels are more predictable across segment routing domain, MPLS label range to assign labels from the configured `srgb` labels for the links.

```
[edit]
set system host-name R2
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 1 family inet address 10.0.0.2/24
set interfaces ge-0/0/0 unit 2 family inet address 10.0.0.12/24
```

```
set interfaces ge-0/0/1 unit 1 family inet address 10.2.2.1/24
set interfaces ge-0/0/1 unit 2 family inet address 10.2.2.11/24
```

```
[edit]
set interfaces lo0 unit 0 family inet address 10.1.255.2/32
set interfaces lo0 unit 0 family mpls
```

```
[edit]
set routing-options router-id 10.1.255.2
set routing-options autonomous-system 100
```

```
[edit]
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 2 vlan-id 2
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 1 vlan-id 1
set interfaces ge-0/0/1 unit 2 vlan-id 2
```

```
[edit]
set interfaces ge-0/0/0 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/0 unit 2 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 2 family mpls maximum-labels 5
```

```
[edit]
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
set routing-options forwarding-table export pplb
```

```
[edit]
set policy-options policy-statement prefix-sid term 1 from route-filter 10.1.255.2/32 exact
```

```
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1001
set policy-options policy-statement prefix-sid term 1 then accept
```

```
[edit]
set protocols mpls traffic-engineering
set protocols mpls label-range srgb-label-range 800000 879999
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```

2. Configure the `ospf-instance metro-edge` on the subinterfaces (connecting from R2 to R1 and from R2 to R3).

```
[edit]
set protocols ospf-instance metro-edge area 0.0.0.0 interface all
```

3. Enable the OSPF `metro-edge` instance to use segment routing with prefix-sids.

```
[edit]
set protocols ospf-instance metro-edge source-packet-routing prefix-segment prefix-sid
```

4. Configure the IPv4 index value of the node segment.

```
[edit]
set protocols ospf-instance metro-edge source-packet-routing node-segment ipv4-index 1
```

5. Configure the loopback address of the OSPF `metro-edge` instance as passive and disable the management interface (`fxp0.0`).

```
[edit]
set protocols ospf-instance metro-edge area 0.0.0.0 interface lo0.0 passive
set protocols ospf-instance metro-edge area 0.0.0.0 interface fxp0.0 disable
```

Verification

IN THIS SECTION

- [Verify the Routing Table | 315](#)
- [Verify OSPF Advertisements | 317](#)
- [Verify the Routes in the OSPF Routing Table | 318](#)
- [Verify the OSPF segment routing database | 319](#)
- [Verify the OSPF Interfaces | 319](#)
- [Verify the OSPF Neighbor | 320](#)

Command	Verification Task
show route protocol ospf table inet.0 extensive	<ul style="list-style-type: none"> ● Verify the route entries in the routing table. ● Verify the loopback address of R1 and R3 is mapped to the igp-instance as configured in R2.
show ospf spring sid-database igp-instance <i>igp-instance</i>	Verify the OSPF segment routing database for the OSPF instance.
show ospf neighbor igp-instance <i>igp-instance</i>	Verify neighbors for the specific OSPF instance.
show ospf database igp-instance <i>igp-instance</i>	Verify the OSPF advertisement entries in the OSPF link-state database (LSDB) associated with the IGP instance.
show ospf interface igp-instance <i>igp-instance</i>	Verify the interfaces mapped to the IGP instance.
show ospf route igp-instance <i>igp-instance</i>	Verify the routes and OSPF instance mapping information of R1 and R3.

Verify the Routing Table

Purpose

Verify the route entries in the routing table

Action

From operational mode, run the `show route table inet.0 route-destination address extensive` command.

```

user@R2>show route protocol ospf table inet.0 10.1.255.1 extensive
inet.0: 19 destinations, 21 routes (19 active, 0 holddown, 0 hidden)
10.1.255.1/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.255.1/32 -> {list:10.0.0.1, 10.0.0.11}
  *OSPF Preference: 10/10
    Next hop type: Router, Next hop index: 0
    Address: 0x8b32234
    Next-hop reference count: 2, Next-hop session id: 0
    Kernel Table Id: 0
    Next hop: 10.0.0.1 via ge-0/0/0.1, selected
    Session Id: 0
    Next hop: 10.0.0.11 via ge-0/0/0.2
    Session Id: 0
    State: <Active Int>
    Local AS: 100
    Age: 1w4d 16:01:19 Metric: 1
    Validation State: unverified
    Area: 0.0.0.0
    Task: OSPF-metro-edge
    Announcement bits (1): 0-KRT
    AS path: I
    Thread: junos-main

```

```

user@R2>show route protocol ospf table inet.0 10.1.255.3 extensive
inet.0: 19 destinations, 21 routes (19 active, 0 holddown, 0 hidden)
10.1.255.3/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.255.3/32 -> {list:10.2.2.2, 10.2.2.12}
  *OSPF Preference: 10/10
    Next hop type: Router, Next hop index: 0
    Address: 0x8b316f4
    Next-hop reference count: 2, Next-hop session id: 0
    Kernel Table Id: 0
    Next hop: 10.2.2.2 via ge-0/0/1.1, selected
    Session Id: 0
    Next hop: 10.2.2.12 via ge-0/0/1.2
    Session Id: 0

```

```

State: <Active Int>
Local AS: 100
Age: 1w4d 16:13:55      Metric: 1
Validation State: unverified
Area: 0.0.0.0
Task: OSPF-metro-edge
Announcement bits (1): 0-KRT
AS path: I
Thread: junos-main

```

Meaning

The output illustrates that the loopback address of R1 (10.1.255.1) and the loopback address of R3 (10.1.255.2) is mapped to the OSPF igp-instance **metro-edge** as configured in R2.

Verify OSPF Advertisements

Purpose

Verify the OSPF advertisement entries in the OSPF link-state database (LSDB) associated with the IGP instance.

Action

From the operational mode, run the `show ospf database igp-instance igp-instance` command.

```

user@R2>show ospf database igp-instance metro-edge
  OSPF database, Area 0.0.0.0

```

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.1.255.1	10.1.255.1	0x80000013	1110	0x22	0xe6e9	72
Router	*10.1.255.2	10.1.255.2	0x80000015	1084	0x22	0x7be2	96
Router	10.1.255.3	10.1.255.3	0x80000013	1585	0x22	0x491	72
Network	*10.0.0.2	10.1.255.2	0x80000010	2959	0x22	0x6791	32
Network	*10.0.0.12	10.1.255.2	0x80000010	2209	0x22	0x3eb	32
Network	10.2.2.2	10.1.255.3	0x80000010	2085	0x22	0x4ba6	32
Network	10.2.2.12	10.1.255.3	0x80000010	1085	0x22	0xe601	32
OpaqArea	7.0.0.1	10.1.255.1	0x80000012	193	0x22	0x8c0	44
OpaqArea*	7.0.0.1	10.1.255.2	0x80000012	511	0x22	0x2a9b	44
OpaqArea	7.0.0.1	10.1.255.3	0x80000012	585	0x22	0x4c76	44
OpaqArea	8.0.0.1	10.1.255.1	0x80000010	2610	0x22	0x4683	48
OpaqArea*	8.0.0.1	10.1.255.2	0x80000010	2584	0x22	0xac01	52

OpaqArea 8.0.0.1	10.1.255.3	0x80000010	2584	0x22	0x7d06	52
OpaqArea 8.0.0.2	10.1.255.1	0x80000010	1860	0x22	0x4f55	48
OpaqArea*8.0.0.2	10.1.255.2	0x80000011	334	0x22	0xf393	52
OpaqArea 8.0.0.2	10.1.255.3	0x80000011	84	0x22	0xc498	52
OpaqArea*8.0.0.3	10.1.255.2	0x80000010	1834	0x22	0x445a	48
OpaqArea*8.0.0.4	10.1.255.2	0x80000010	1459	0x22	0x4d2c	48

Meaning

Verify the Routes in the OSPF Routing Table

Purpose

Verify the routes in the OSPF routing table

Action

From the operational mode, run the `show ospf route` command.

```

user@R2>show ospf route igp-instance metro-edge
Topology default Route Table:

Prefix          Path Route   NH      Metric NextHop      Nexthop
                Type Type     Type                    Interface    Address/LSP
10.1.255.1      Intra Router IP        1 ge-0/0/0.1   10.0.0.1
                ge-0/0/0.2   10.0.0.11
10.1.255.3      Intra Router IP        1 ge-0/0/1.1   10.2.2.2
                ge-0/0/1.2   10.2.2.12
10.0.0.0/24     Intra Network IP        1 ge-0/0/0.1
                ge-0/0/0.2
10.1.255.1/32  Intra Network IP        1 ge-0/0/0.1   10.0.0.1
                ge-0/0/0.2   10.0.0.11
10.1.255.2/32   Intra Network IP        0 lo0.0
10.1.255.3/32  Intra Network IP        1 ge-0/0/1.1   10.2.2.2
                ge-0/0/1.2   10.2.2.12
10.2.2.0/24     Intra Network IP        1 ge-0/0/1.1
                ge-0/0/1.2
299840          Intra Network Mpls     0 ge-0/0/0.2   10.0.0.11
299840 (S=0)    Intra Network Mpls     0 ge-0/0/0.2   10.0.0.11

```

299856	Intra Network	Mpls	0 ge-0/0/0.1	10.0.0.1
299856 (S=0)	Intra Network	Mpls	0 ge-0/0/0.1	10.0.0.1
299904	Intra Network	Mpls	0 ge-0/0/1.2	10.2.2.12
299904 (S=0)	Intra Network	Mpls	0 ge-0/0/1.2	10.2.2.12
299920	Intra Network	Mpls	0 ge-0/0/1.1	10.2.2.2
299920 (S=0)	Intra Network	Mpls	0 ge-0/0/1.1	10.2.2.2

Meaning

The output on R2 shows the loopback addresses and OSPF instance mapping information of R1 and R3.

Verify the OSPF segment routing database

Purpose

Verify the OSPF segment routing database for the OSPF instance metro-edge.

Action

From the operational mode, run the `show ospf spring sid-database igp-instance igp-instance` command.

```
user@R2>show ospf spring sid-database igp-instance metro-edge
  OSPF database, Area 0.0.0.0
SID    Prefix          Advertised-by  Route-type
1000   10.1.255.1/32   10.1.255.1    Intra-Area
1001   10.1.255.2/32   10.1.255.2    Intra-Area
1002   10.1.255.3/32   10.1.255.3    Intra-Area
```

Meaning

The output illustrates the multiple instances of OSPF (metro-edge) advertise prefix-SIDs.

Verify the OSPF Interfaces

Purpose

Verify the status information about OSPF-instance enabled interfaces.

Action

From the operational mode, run the `show ospf interface igp-instance igp-instance` command.

```
user@R2>show ospf interface igp-instance metro-edge
Interface      State  Area      DR ID      BDR ID      Nbrs
ge-0/0/0.1    DR     0.0.0.0   10.1.255.2 10.1.255.1   1
ge-0/0/0.2    DR     0.0.0.0   10.1.255.2 10.1.255.1   1
ge-0/0/1.1    DR     0.0.0.0   10.1.255.2 10.1.255.3   1
ge-0/0/1.2    DR     0.0.0.0   10.1.255.2 10.1.255.3   1
lo0.0         DRother 0.0.0.0   0.0.0.0    0.0.0.0      0
lo0.0         DRother 0.0.0.0   0.0.0.0    0.0.0.0      0
```

Meaning

The output shows the subinterfaces of R2 mapped to the OSPF instances (metro-edge).

Verify the OSPF Neighbor

Purpose

Verify the adjacencies between the configured links.

Action

From the operational mode, run the `show ospf neighbor igp-instance igp-instance` command.

```
user@R2>show ospf neighbor igp-instance metro-edge
Address      Interface      State      ID          Pri  Dead
10.0.0.1     ge-0/0/0.1    Full      10.1.255.1  128  35
10.0.0.11    ge-0/0/0.2    Full      10.1.255.1  128  39
10.2.2.2     ge-0/0/1.1    Full      10.1.255.3  128  33
10.2.2.12    ge-0/0/1.2    Full      10.1.255.3  128  36
```

Meaning

Device R2 has established adjacency with Device R1 and Device R3 and as indicated by the State output field which is Full.

Appendix 1: Set Commands on All Devices

IN THIS SECTION

- R1 | 321
- R2 | 322
- R3 | 323

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

R1

```
set system host-name R1
set interfaces ge-0/0/0 unit 1 family inet address 10.0.0.1/24
set interfaces ge-0/0/0 unit 2 family inet address 10.0.0.11/24
set interfaces ge-0/0/0 unit 2 enable
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 2 vlan-id 2
set interfaces ge-0/0/0 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/0 unit 2 family mpls maximum-labels 5
set interfaces lo0 unit 0 family inet address 10.1.255.1/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
set policy-options policy-statement prefix-sid term 1 from route-filter 10.1.255.1/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1000
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 10.1.255.1
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range srgb-label-range 800000 879999
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
```

```

set protocols mpls interface fxp0.0 disable
set protocols ospf-instance metro-edge source-packet-routing prefix-segment prefix-sid
set protocols ospf-instance metro-edge source-packet-routing node-segment ipv4-index 0
set protocols ospf-instance metro-edge area 0.0.0.0 interface all
set protocols ospf-instance metro-edge area 0.0.0.0 interface lo0.0 passive
set protocols ospf-instance metro-edge area 0.0.0.0 interface fxp0.0 disable

```

R2

```

set system host-name R2
set interfaces ge-0/0/0 unit 1 family inet address 10.0.0.2/24
set interfaces ge-0/0/0 unit 2 family inet address 10.0.0.12/24
set interfaces ge-0/0/1 unit 1 family inet address 10.2.2.1/24
set interfaces ge-0/0/1 unit 2 family inet address 10.2.2.11/24
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 2 vlan-id 2
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 1 vlan-id 1
set interfaces ge-0/0/1 unit 2 vlan-id 2
set interfaces ge-0/0/0 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/0 unit 2 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 2 family mpls maximum-labels 5
set interfaces lo0 unit 0 family inet address 10.1.255.2/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
set policy-options policy-statement prefix-sid term 1 from route-filter 10.1.255.2/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1001
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 10.1.255.2
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range srgb-label-range 800000 879999
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf-instance metro-edge source-packet-routing prefix-segment prefix-sid
set protocols ospf-instance metro-edge source-packet-routing node-segment ipv4-index 1

```

```

set protocols ospf-instance metro-edge area 0.0.0.0 interface all
set protocols ospf-instance metro-edge area 0.0.0.0 interface lo0.0 passive
set protocols ospf-instance metro-edge area 0.0.0.0 interface fxp0.0 disable

```

R3

```

set system host-name R3
set interfaces ge-0/0/0 unit 1 family inet address 10.2.2.2/24
set interfaces ge-0/0/0 unit 2 family inet address 10.2.2.12/24
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 2 vlan-id 2
set interfaces ge-0/0/0 unit 1 family mpls maximum-labels 5
set interfaces ge-0/0/0 unit 2 family mpls maximum-labels 5
set interfaces lo0 unit 0 family inet address 10.1.255.3/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
set policy-options policy-statement prefix-sid term 1 from route-filter 10.1.255.3/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1002
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 10.1.255.3
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range srgb-label-range 800000 879999
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf-instance metro-edge source-packet-routing prefix-segment prefix-sid
set protocols ospf-instance metro-edge source-packet-routing node-segment ipv4-index 2
set protocols ospf-instance metro-edge area 0.0.0.0 interface lo0.0 passive
set protocols ospf-instance metro-edge area 0.0.0.0 interface all
set protocols ospf-instance metro-edge area 0.0.0.0 interface fxp0.0 disable

```

8

CHAPTER

Configure OSPF Timers

IN THIS CHAPTER

- [Configuring OSPF Timers | 325](#)
-

Configuring OSPF Timers

IN THIS SECTION

- [OSPF Timers Overview | 325](#)
- [Example: Configuring OSPF Timers | 326](#)

OSPF Timers Overview

OSPF routing devices constantly track the status of their neighbors, sending and receiving hello packets that indicate whether each neighbor still is functioning, and sending and receiving link-state advertisement (LSA) and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You configure OSPF timers on the interface of the routing device participating in OSPF. Depending on the timer, the configured interval must be the same on all routing devices on a shared network (area).

You can configure the following OSPF timers:

- Hello interval—Routing devices send hello packets at a fixed interval on all interfaces, including virtual links, to establish and maintain neighbor relationships. The hello interval specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. This interval must be the same on all routing devices on a shared network. By default, the routing device sends hello packets every 10 seconds (broadcast and point-to-point networks) and 30 seconds (nonbroadcast multiple access (NBMA) networks).



NOTE: For EX Series and QFX Series switches, as well as SRX Series devices the hello interval is 10 seconds or longer.

- Poll interval—(OSPFv2, Nonbroadcast networks only) Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The poll interval specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before establishing adjacency with a neighbor. By default, the routing device sends hello packets every 120 seconds until active neighbors are detected.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval.

- **LSA retransmission interval**—When a routing device sends LSAs to its neighbors, the routing device expects to receive an acknowledgment packet from each neighbor within a certain amount of time. The LSA retransmission interval specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting the LSA to an interface's neighbors. By default, the routing device waits 5 seconds for an acknowledgment before retransmitting the LSA.
- **Dead interval**—If a routing device does not receive a hello packet from a neighbor within a fixed amount of time, the routing device modifies its topology database to indicate that the neighbor is nonoperational. The dead interval specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. This interval must be the same on all routing devices on a shared network. By default, this interval is four times the default hello interval, which is 40 seconds (broadcast and point-to-point networks) and 120 seconds (NBMA networks).
- **Transit delay**—Before a link-state update packet is propagated out of an interface, the routing device must increase the age of the packet. The transit delay sets the estimated time required to transmit a link-state update on the interface. By default, the transit delay is 1 second. You should never have to modify the transit delay time.



NOTE: The default hello interval and default dead interval are the same across all devices. For lower interval requirements, we recommend using BFD.

SEE ALSO

| [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network](#) | 33

Example: Configuring OSPF Timers

IN THIS SECTION

- [Requirements](#) | 327
- [Overview](#) | 327
- [Configuration](#) | 328
- [Verification](#) | 334

This example shows how to configure the OSPF timers.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

The default OSPF timer settings are optimal for most networks. However, depending on your network requirements, you might need to modify the timer settings. This example explains why you might need to modify the following timers:

- Hello interval
- Dead interval
- LSA retransmission interval
- Transit delay

Hello Interval and Dead Interval

The hello interval and the dead interval optimize convergence times by efficiently tracking neighbor status. By lowering the values of the hello interval and the dead interval, you can increase the convergence of OSPF routes if a path fails. These intervals must be the same on all routing devices on a shared network. Otherwise, OSPF cannot establish the appropriate adjacencies.

In the first example, you lower the hello interval to 2 seconds and the dead interval to 8 seconds on point-to-point OSPF interfaces **fe-0/0/1** and **fe-1/0/1** in area 0.0.0.0 by configuring the following settings:

- **hello-interval**—Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds. The range is from 1 through 255 seconds.



NOTE: For EX Series and QFX Series switches, the hello interval is 10 seconds or longer.

- **dead-interval**—Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. By default, the routing device waits 40 seconds (four times the hello interval). The range is 1 through 65,535 seconds.

LSA Retransmission Interval

The link-state advertisement (LSA) retransmission interval optimizes the sending and receiving of LSA and acknowledgement packets. You must configure the LSA retransmission interval to be equal to or greater than 3 seconds to avoid triggering a retransmit trap because the Junos OS delays LSA acknowledgments by up to 2 seconds. If you have a virtual link, you might find increased performance by increasing the value of the LSA retransmission interval.

In the second example, you increase the LSA retransmission timer to 8 seconds on OSPF interface **fe-0/0/1** in area 0.0.0.1 by configuring the following setting:

- **retransmit-interval**—Specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting LSA to an interface's neighbors. By default, the routing device retransmits LSAs to its neighbors every 5 seconds. The range is from 1 through 65,535 seconds.

Transit Delay

The transit delay sets the time the routing device uses to age a link-state update packet. If you have a slow link (for example, one with an average propagation delay of multiple seconds), you should increase the age of the packet by a similar amount. Doing this ensures that you do not receive a packet back that is younger than the original copy.

In the final example, you increase the transit delay to 2 seconds on OSPF interface **fe-1/0/1** in area 0.0.0.1. By configuring the following setting, this causes the routing device to age the link-state update packet by 2 seconds:

- **transit-delay**—Sets the estimated time required to transmit a link-state update on the interface. You should never have to modify the transit delay time. By default, the routing device ages the packet by 1 second. The range is from 1 through 65,535 seconds.

Configuration

IN THIS SECTION

- [Configuring the Hello Interval and the Dead Interval | 329](#)

- [Controlling the LSA Retransmission Interval | 331](#)
- [Specifying the Transit Delay | 332](#)

Configuring the Hello Interval and the Dead Interval

CLI Quick Configuration

To quickly configure the hello and dead intervals, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-0/0/1 dead-interval 8
set protocols ospf area 0.0.0.0 interface fe-1/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-1/0/1 dead-interval 8
```

Step-by-Step Procedure

To configure the hello and dead intervals:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
user@host# set interface fe-1/0/1
```

3. Configure the hello interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 hello-interval 2
user@host# set interface fe-1/0/1 hello-interval 2
```

4. Configure the dead interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 dead-interval 8
user@host# set interface fe-1/0/1 dead-interval 8
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



NOTE: Repeat this entire configuration on all routing devices in a shared network.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
  interface fe-1/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Controlling the LSA Retransmission Interval

CLI Quick Configuration

To quickly configure the LSA retransmission interval, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-0/0/1 retransmit-interval 8
```

Step-by-Step Procedure

To configure the LSA retransmission interval:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1
```

3. Configure the LSA retransmission interval.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface fe-0/0/1 retransmit-interval 8
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-0/0/1.0 {
    retransmit-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Specifying the Transit Delay

CLI Quick Configuration

To quickly configure the transit delay, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-1/0/1 transit-delay 2
```

Step-by-Step Procedure

To configure the transit delay:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-1/0/1
```

3. Configure the transit delay.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface fe-1/0/1 transit-delay 2
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-1/0/1.0 {
    transit-delay 2;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Timer Configuration | 334](#)

Confirm that the configuration is working properly.

Verifying the Timer Configuration

Purpose

Verify that the interface for OSPF or OSPFv3 has been configured with the applicable timer values. Confirm that the Hello field, the Dead field, and the ReXmit field display the values that you configured.

Action

From operational mode, enter the `show ospf interface detail` for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

RELATED DOCUMENTATION

[About OSPF Interfaces | 16](#)

[Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network | 33](#)

9

CHAPTER

Configure OSPF Fault Detection using BFD

IN THIS CHAPTER

- [Configuring OSPF Fault Detection using BFD | 336](#)
-

Configuring OSPF Fault Detection using BFD

IN THIS SECTION

- [Understanding BFD for OSPF | 336](#)
- [Example: Configuring BFD for OSPF | 338](#)
- [Understanding BFD Authentication for OSPF | 344](#)
- [Configuring BFD Authentication for OSPF | 346](#)

Understanding BFD for OSPF

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

You can configure the following BFD protocol settings:

- `detection-time threshold`—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- `full-neighbors-only`—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.

- `minimum-interval`—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the `transmit-interval` `minimum-interval` and `minimum-receive-interval` statements.



NOTE: Depending on your network environment, the following may apply:

- The recommended minimum interval for centralised BFD is 300 ms with a multiplier of 3, and the recommended minimum interval for distributed BFD is 100 ms with a multiplier of 3.
- For BFD sessions to remain up during a Routing Engine switchover event when *nonstop active routing* (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms.
- For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- BFD is not distributed prior to Junos 21.2 (because for OSPFv3, BFD is based in the Routing Engine).

- `minimum-receive-interval`—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device expects to receive a hello packet from a neighbor with which it has established a BFD session. You can also specify the minimum receive interval using the `minimum-interval` statement.
- `multiplier`—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- `no-adaptation`—Disables BFD adaptation. This setting disables BFD sessions from adapting to changing network conditions.



NOTE: We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- `transmit-interval` `minimum-interval`—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the `minimum-interval` statement.

- `transmit-interval threshold`—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- `version`—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

Example: Configuring BFD for OSPF

IN THIS SECTION

- [Requirements | 338](#)
- [Overview | 339](#)
- [Configuration | 340](#)
- [Verification | 343](#)

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

Requirements

Before you begin:

- Configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62.

- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)

Overview

IN THIS SECTION

- [Topology | 340](#)

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.

BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the `bfd-liveness-detection` statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and `minimum-receive-interval` statements.



NOTE: Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of no less than 500 ms. An interval of 1000 ms is recommended to avoid any instability issues.



NOTE:

- For the `bfdd` process, the detection time interval set is lower than 300 ms. If there is a high priority process such as `ppmd` running on the system, the CPU might spend time on the `ppmd` process rather than the `bfdd` process.
- For branch SRX Series Firewalls, we recommend 1000 ms as the minimum keepalive time interval for BFD packets.
- For vSRX 3.0, we recommend 300 ms as the minimum keepalive time interval for BFD packets.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

Topology

Configuration

IN THIS SECTION

- Procedure | 341

Procedure

CLI Quick Configuration

To quickly configure the BFD protocol for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

Step-by-Step Procedure

To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospfv3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



NOTE: Repeat this entire configuration on the other neighboring interface.

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    bfd-liveness-detection {
      minimum-interval 300;
      multiplier 4;
      full-neighbors-only;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the BFD Sessions | 343](#)

Confirm that the configuration is working properly.

Verifying the BFD Sessions

Purpose

Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

Action

From operational mode, enter the `show bfd session detail` command.

Meaning

The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.
- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

Understanding BFD Authentication for OSPF

IN THIS SECTION

- [BFD Authentication Algorithms | 344](#)
- [Security Authentication Keychains | 345](#)
- [Strict Versus Loose Authentication | 345](#)

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over OSPFv2. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.

- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: *Nonstop active routing* (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.



NOTE: QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you

can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Configuring BFD Authentication for OSPF

IN THIS SECTION

- [Configuring BFD Authentication Parameters | 346](#)
- [Viewing Authentication Information for BFD Sessions | 348](#)

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over OSPFv2. Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on OSPF:

Configuring BFD Authentication Parameters

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the OSPFv2 protocol.
2. Associate the authentication keychain with the OSPFv2 protocol.
3. Configure the related security authentication keychain.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on an OSPF route or routing instance.

[edit]

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection  
authentication algorithm keyed-sha-1
```



NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified OSPF route or routing instance with the unique security authentication keychain attributes.

This keychain should match the keychain name configured at the [edit security authentication key-chains] hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection
authentication keychain bfd-ospf
```



NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys enables multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# authentication-key-chains key-chain bfd-ospf key 53 secret $ABC123$ABC123 start-
time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set protocols ospf interface if2-ospf bfd-liveness-detection authentication loose-
check
```

5. (Optional) View your configuration using the `show bfd session detail` or `show bfd session extensive` command.
6. Repeat the steps in this procedure to configure the other end of the BFD session.



NOTE: BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the `show bfd session detail` and `show bfd session extensive` commands.

The following example shows BFD authentication configured for the **if2-ospf** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-ospf**. The authentication keychain is configured with two keys. Key **1** contains the secret data “**\$ABC123\$ABC123**” and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data “**\$ABC123\$ABC123**” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols ospf]
area 0.0.0.1 {
  interface if2-ospf {
    bfd-liveness-detection {
      authentication {
        algorithm keyed-sha-1;
        key-chain bfd-ospf;
      }
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-ospf {
    key 1 {
      secret "$ABC123$ABC123"; ## SECRET-DATA
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$ABC123$ABC123";
      start-time "2009-6-1.15:29:20 -0700"; ## SECRET-DATA
    }
  }
}
```

```

    }
}

```

If you commit these updates to your configuration, you see output similar to the following. In the output for the `show bfd session detail` command, **Authenticate** is displayed to indicate that BFD authentication is configured.

show bfd session detail

```

user@host# show bfd session detail

Address          State   Interface      Detect   Transmit
                  Time    Interval Multiplier
10.9.1.33        Up      so-7/1/0.0     0.600   0.200   3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate
Session up time 3d 00:34
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated

1 sessions, 1 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

For more information about the configuration, use the `show bfd session extensive` command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session extensive

```

user@host# show bfd session extensive

Address          State   Interface      Detect   Transmit
                  Time    Interval Multiplier
10.9.1.33        Up      so-7/1/0.0     0.600   0.200   3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate
keychain bfd-ospf, algo keyed-md5, mode loose

Session up time 3d 00:34
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated
Min async interval 0.200, min slow interval 1.000

```

```
Adaptive async tx interval 0.200, rx interval 0.200
Local min tx interval 0.200, min rx interval 0.200, multiplier 3
Remote min tx interval 0.100, min rx interval 0.100, multiplier 3
Threshold transmission interval 0.000, Threshold for detection time 0.000
Local discriminator 11, remote discriminator 80
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-ospf, algo keyed-sha-1, mode strict
1 sessions, 1 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

RELATED DOCUMENTATION

[*bfd-liveness-detection*](#)

[Junos OS Administration Library for Routing Devices](#)

[CLI Explorer](#)

[Example: Configuring BFD Authentication for OSPF](#)

10

CHAPTER

Configure Graceful Restart for OSPF

IN THIS CHAPTER

- [Configuring Graceful Restart for OSPF | 352](#)
-

Configuring Graceful Restart for OSPF

IN THIS SECTION

- [Graceful Restart for OSPF Overview | 352](#)
- [Example: Configuring Graceful Restart for OSPF | 354](#)
- [Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart | 361](#)
- [Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart | 367](#)
- [Example: Disabling Strict LSA Checking for OSPF Graceful Restart | 372](#)

Graceful Restart for OSPF Overview

IN THIS SECTION

- [Helper Mode for Graceful Restart | 353](#)
- [Planned and Unplanned Graceful Restart | 353](#)

Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting device and its neighbors continue forwarding packets without disrupting network performance. Because neighboring devices assist in the restart (these neighbors are called *helpers*), the restarting device can quickly resume full operation without recalculating algorithms.



NOTE: On a broadcast link with a single neighbor, when the neighbor initiates an OSPFv3 graceful restart operation, the restart might be terminated at the point when the local routing device assumes the role of a helper. A change in the LSA is considered a topology change, which terminates the neighbor's restart operation.

Graceful restart is disabled by default. You can either globally enable graceful restart for all routing protocols, or you can enable graceful restart specifically for OSPF.

This topic describes the following information:

Helper Mode for Graceful Restart

When a device enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The device does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This device continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting device must send a grace LSA to all neighbors. In response, the helper routers enter helper mode (the ability to assist a neighboring device attempting a graceful restart) and send an acknowledgment back to the restarting device. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting device had remained in continuous OSPF operation.



NOTE: Helper mode is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode specifically for OSPF.

When the restarting device receives replies from all the helper routers, the restarting device selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting device receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting device or when the topology of the network changes, the helper routers also resume normal operation.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The Junos OS implementation is based on RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*. In restart signaling-based helper mode implementations, the restarting device informs its restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting device sends hello messages to its helper routers with the restart signal (RS) bit set in the hello packet header. When a helper router receives a hello packet with the RS bit set in the header, the helper router returns a hello message to the restarting device. The reply hello message from the helper router contains the ResyncState flag and the ResyncTimeout timer that enable the restarting device to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting device exits the restart mode.



NOTE: Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

Planned and Unplanned Graceful Restart

OSPF supports two types of graceful restart: planned and unplanned. During a planned restart, the restarting routing device informs the neighbors before restarting. The neighbors act as if the routing device is still within the network topology, and continue forwarding traffic to the restarting routing

device. A grace period is set to specify when the neighbors should consider the restarting routing device as part of the topology. During an unplanned restart, the routing device restarts without warning.

Example: Configuring Graceful Restart for OSPF

IN THIS SECTION

- Requirements | 354
- Overview | 354
- Configuration | 355
- Verification | 359

This example shows how to configure graceful restart specifically for OSPF.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54.](#)
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57.](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62.](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)

Overview

IN THIS SECTION

- Topology | 355

Graceful restart enables a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting routing device and its neighbors continue forwarding packets without disrupting network performance. By default, graceful restart is disabled. You can globally enable graceful restart for all routing protocols by including the `graceful-restart` statement at the `[edit routing-options]` hierarchy level, or you can enable graceful restart specifically for OSPF by including the `graceful-restart` statement at the `[edit protocols (ospf|ospf3)]` hierarchy level.

The first example shows how to enable graceful restart and configure the optional settings for the grace period interval. In this example, interfaces `fe-1/1/1` and `fe-1/1/2` are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. The grace period interval for OSPF graceful restart is determined as equal to or less than the sum of the **notify-duration** time interval and the **restart-duration** time interval. The grace period is the number of seconds that the routing device's neighbors continue to advertise the routing device as fully adjacent, regardless of the connection state between the routing device and its neighbors.

The `notify-duration` statement configures how long (in seconds) the routing device notifies helper routers that it has completed graceful restart by sending purged grace link-state advertisements (LSAs) over all interfaces. By default, the routing device sends grace LSAs for 30 seconds. The range is from 1 through 3600 seconds.

The `restart-duration` statement configures the amount of time the routing device waits (in seconds) to complete reacquisition of OSPF neighbors from each area. By default, the routing device allows 180 seconds. The range is from 1 through 3600 seconds.

The second example shows how to disable graceful restart for OSPF by including the `disable` statement.

Topology

Configuration

IN THIS SECTION

- [Enabling Graceful Restart for OSPF | 356](#)
- [Disabling Graceful Restart for OSPF | 358](#)

Enabling Graceful Restart for OSPF

CLI Quick Configuration

To quickly enable graceful restart for OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set routing-options graceful-restart
set protocols ospf graceful-restart restart-duration 190
set protocols ospf graceful-restart notify-duration 40
```

Step-by-Step Procedure

To enable graceful restart for OSPF:

1. Configure the interfaces.

 **NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces.

 **NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Configure graceful restart globally

```
[edit]
user@host#edit routing-options graceful-restart
```

4. Configure OSPF graceful restart.

```
[edit]
user@host# edit protocols ospf graceful-restart
```

5. (Optional) Configure the restart duration time.

```
[edit protocols ospf graceful-restart]
user@host# set restart-duration 190
```

6. (Optional) Configure the notify duration time.

```
[edit protocols ospf graceful-restart]
user@host# set notify-duration 40
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf graceful-restart]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
```

```

    }
}
fe-1/1/2 {
    unit 0 {
        family inet {
            address 10.0.0.5/32;
        }
    }
}
}
user@host# show protocols ospf
graceful-restart {
    restart-duration 190;
    notify-duration 40;
}
area 0.0.0.0 {
    interface fe-1/1/1.0;
    interface fe-1/1/2.0;
}

```

To confirm an OSPFv3 configuration, enter the `show interfaces` and the `show protocols ospf3` commands.

Disabling Graceful Restart for OSPF

CLI Quick Configuration

To quickly disable graceful restart for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```

[edit]
user@host# set protocols ospf graceful-restart disable

```

Step-by-Step Procedure

To disable graceful restart for OSPF:

1. Disable graceful restart for the OSPF protocol only.

This command does not affect the global graceful restart configuration setting.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
graceful-restart disable;
```

To confirm an OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the OSPF Graceful Restart Configuration | 360](#)
- [Verifying Graceful Restart Status | 360](#)

Confirm that the configuration is working properly.

Verifying the OSPF Graceful Restart Configuration

Purpose

Verify information about your OSPF graceful restart configuration.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2. Enter the `show ospf3 overview` command for OSPFv3.

Meaning

The Restart field displays the status of graceful restart as either enabled or disabled. The Restart duration field displays how much time the restarted routing device requires to complete reacquisition of OSPF neighbors. The Restart grace period field displays how much time the neighbors should consider the restarted routing device as part of the topology.

Verifying Graceful Restart Status

Purpose

Verify the status of graceful restart.

Action

From operational mode, enter the `show route instance detail` command.

Meaning

The Restart State field displays Pending if the restart has not been completed or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have or have not yet completed graceful restart for the specified routing table.

Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart

IN THIS SECTION

- [Requirements | 361](#)
- [Overview | 361](#)
- [Configuration | 362](#)
- [Verification | 366](#)

This example shows how to disable and reenable the helper mode capability for OSPFv2 graceful restart.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62](#).
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

IN THIS SECTION

- [Topology | 362](#)

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv2 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the standard OSPFv2 graceful restart helper capability by including the `helper-disable standard` statement. This configuration is useful if you have an environment that contains other vendor equipment that is configured for restart signaling-based graceful restart.



NOTE: The `helper-disable` statement and the `no-strict-lsa-checking` statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the `show protocols ospf` command.

The second example shows how to reenabling the standard OSPFv2 restart helper capability that you disabled in the first example.

Topology

Configuration

IN THIS SECTION

- [Disabling Helper Mode for OSPFv2 | 363](#)
- [Reenabling Helper Mode for OSPFv2 | 365](#)

Disabling Helper Mode for OSPFv2

CLI Quick Configuration

To quickly enable graceful restart for OSPFv2 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart helper-disable standard
```

Step-by-Step Procedure

To enable graceful restart for OSPFv2 with helper mode disabled:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPFv2 on the interfaces

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable the OSPFv2 graceful restart helper capability.

If you disable the OSPFv2 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf graceful-restart helper-disable standard
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  helper-disable {
    standard;
  }
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

Reenabling Helper Mode for OSPFv2

CLI Quick Configuration

To quickly reenable standard helper-mode for OSPFv2, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
delete protocols ospf graceful-restart helper-disable standard
```



NOTE: To reenable restart signaling-based helper mode, include the `restart-signaling` statement. To reenable both standard and restart signaling-based helper mode, include the `both` statement.

Step-by-Step Procedure

To reenable standard helper mode for OSPFv2:

1. Delete the standard helper-mode statement from the OSPFv2 configuration.

```
[edit]
user@host# delete protocols ospf graceful-restart helper-disable standard
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

After you reenable standard helper mode, the `show protocols ospf` command no longer displays the graceful restart configuration.

Verification

IN THIS SECTION

- [Verifying the OSPFv2 Graceful Restart Configuration | 366](#)
- [Verifying Graceful Restart Status | 366](#)

Confirm that the configuration is working properly.

Verifying the OSPFv2 Graceful Restart Configuration

Purpose

Verify information about your OSPFv2 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, the Graceful restart helper mode field displays the status of the standard helper mode capability as enabled or disabled, and the Restart-signaling helper mode field displays the status of the restart signaling-based helper mode as enabled or disabled. By default, both standard and restart signaling-based helper modes are enabled.

Action

From operational mode, enter the `show ospf overview` command.

Verifying Graceful Restart Status

Purpose

Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action

From operational mode, enter the `show route instance detail` command.

Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart

IN THIS SECTION

- [Requirements | 367](#)
- [Overview | 367](#)
- [Configuration | 368](#)
- [Verification | 371](#)

This example shows how to disable and reenable the helper mode capability for OSPFv3 graceful restart.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62](#).
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

IN THIS SECTION

- [Topology | 368](#)

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv3 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the OSPFv3 graceful restart helper capability by including the `helper-disable` statement.



NOTE: The `helper-disable` statement and the `no-strict-lsa-checking` statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the `show protocols ospf` command.

The second example shows how to reenabling the OSPFv3 restart helper capability that you disabled in the first example.

Topology

Configuration

IN THIS SECTION

- [Disabling Helper Mode for OSPFv3 | 368](#)
- [Reenabling Helper Mode for OSPFv3 | 370](#)

Disabling Helper Mode for OSPFv3

CLI Quick Configuration

To quickly enable graceful restart for OSPFv3 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet6 address 2001:0a00:0004::
set interfaces fe-1/1/2 unit 0 family inet6 address 2001:0a00:0005::
```

```
set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
set protocols ospf3 graceful-restart helper-disable
```

Step-by-Step Procedure

To enable graceful restart for OSPFv3 with helper mode disabled:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet6 address 2001:0a00:0004::
user@host# set interfaces fe-1/1/1 unit 0 family inet address 2001:0a00:0005::
```

2. Configure OSPFv3 on the interfaces

```
[edit]
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
```

3. Disable the OSPFv3 graceful restart helper capability.

If you disable the OSPFv3 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf3 graceful-restart helper-disable
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf3` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet6 {
      address 2001:0a00:0004::/128;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet6 {
      address 2001:0a00:0005::/128;
    }
  }
}
user@host# show protocols ospf3
graceful-restart {
  helper-disable;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

Reenabling Helper Mode for OSPFv3

CLI Quick Configuration

To quickly reenable helper-mode for OSPFv3, copy the following command and paste it into the CLI.

```
[edit]
delete protocols ospf3 graceful-restart helper-disable
```

Step-by-Step Procedure

To reenable helper mode for OSPFv3:

1. Delete the standard helper-mode statement from the OSPFv3 configuration.

```
[edit]  
user@host# delete protocols ospf3 graceful-restart helper-disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

Results

After you reenable standard helper mode, the `show protocols ospfs` command no longer displays the graceful restart configuration.

Verification

IN THIS SECTION

- [Verifying the OSPFv3 Graceful Restart Configuration | 371](#)
- [Verifying Graceful Restart Status | 372](#)

Confirm that the configuration is working properly.

Verifying the OSPFv3 Graceful Restart Configuration

Purpose

Verify information about your OSPFv3 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, and the Helper mode field displays the status of the helper mode capability as either enabled or disabled.

Action

From operational mode, enter the `show ospf3 overview` command.

Verifying Graceful Restart Status

Purpose

Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action

From operational mode, enter the `show route instance detail` command.

Example: Disabling Strict LSA Checking for OSPF Graceful Restart

IN THIS SECTION

- [Requirements | 372](#)
- [Overview | 373](#)
- [Configuration | 373](#)
- [Verification | 376](#)

This example shows how to disable strict link-state advertisement (LSA) checking for OSPF graceful restart.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.

- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62.](#)
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65.](#)

Overview

IN THIS SECTION

- [Topology | 373](#)

You can disable strict LSA checking to prevent the termination of graceful restart by a helping router. You might configure this option for interoperability with other vendor devices. The OSPF graceful restart helper capability must be enabled if you disable strict LSA checking. By default, LSA checking is enabled.

In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable strict LSA checking by including the `no-strict-lsa-checking` statement.



NOTE: The `helper-disable` statement and the `no-strict-lsa-checking` statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the `show protocols ospf` command.

Topology

Configuration

IN THIS SECTION

- [Procedure | 374](#)

Procedure

CLI Quick Configuration

To quickly enable graceful restart for OSPF with strict LSA checking disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart no-strict-lsa-checking
```

Step-by-Step Procedure

To enable graceful restart for OSPF with strict LSA checking disabled:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable strict LSA checking.

If you disable the strict LSA checking, OSPF graceful restart helper capability must be enabled (which is the default behavior).

```
[edit]
user@host# set protocols ospf graceful-restart no-strict-lsa-checking
```

4. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  no-strict-lsa-checking;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
```

```
interface fe-1/1/2.0;  
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces` and the `show protocols ospf3` commands.

Verification

IN THIS SECTION

- [Verifying the OSPF Graceful Restart Configuration | 376](#)
- [Verifying Graceful Restart Status | 376](#)

Confirm that the configuration is working properly.

Verifying the OSPF Graceful Restart Configuration

Purpose

Verify information about your OSPF graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` command for OSPFv3.

Verifying Graceful Restart Status

Purpose

Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

Action

From operational mode, enter the `show route instance detail` command.

RELATED DOCUMENTATION

| [Understanding Graceful Restart](#)

11

CHAPTER

Configure Loop-Free Alternate Routes for OSPF

IN THIS CHAPTER

- [Configuring Loop-Free Alternate Routes for OSPF | 379](#)
-

Configuring Loop-Free Alternate Routes for OSPF

IN THIS SECTION

- [Per Prefix Loop Free Alternates for OSPF | 379](#)
- [Configuring Per-Prefix LFA for OSPF | 380](#)
- [Loop-Free Alternate Routes for OSPF Overview | 381](#)
- [Example: Configuring Loop-Free Alternate Routes for OSPF | 382](#)
- [Configuring Link Protection for OSPF | 421](#)
- [Configuring Node-Link Protection for OSPF | 422](#)
- [Configuring Node to Link Protection Fallback for OSPF | 423](#)
- [Excluding an OSPF Interface as a Backup for a Protected Interface | 424](#)
- [Configuring Backup SPF Options for Protected OSPF Interfaces | 425](#)
- [Configuring RSVP Label-Switched Paths as Backup Paths for OSPF | 427](#)
- [Remote LFA over LDP Tunnels in OSPF Networks Overview | 428](#)
- [Configuring Remote LFA Backup over LDP Tunnels in an OSPF Network | 429](#)
- [Example: Configuring Remote LFA Over LDP Tunnels in OSPF Networks | 431](#)


Per Prefix Loop Free Alternates for OSPF

In certain topologies and usage scenarios, when multiple destinations originate the same prefix and there is no viable LFA to the best prefix originator, whilst a non-best prefix originator has one. *Per-prefix LFA* is a technology by which, the LFA to a non-best prefix originator can be used in lieu of the LFA to the best prefix originator to provide local repair. This can be used to increase the local repair coverage for the OSPF protocol also.

Per-Prefix Loop Free Alternates (LFA)—Loop Free Alternates (LFA) is a technology by which a neighbor can be used as a backup next hop to provide a local repair path for the traffic to flow temporarily in case of failures in the primary next hop (node or link). For this, the basic requirement is that the selected backup neighbor provides a loop free path with respect to primary next hop towards a destination, originating a set of interior gateway protocol (IGP) prefixes.

The following topology explains the deployment case where per prefix LFA feature is applicable.

Figure 24: Per-Prefix LFA Usage Scenario



ABR1 and ABR2 are area boundary routers (ABRs), dual homed to an IPv6 core network, which advertises the summary LSA for the prefix 10.0.1.0/24 with a metric of 10. Also, from PE router's perspective, ABR1 is the best prefix originator for 10.0.1.0/24. In this case, P2 is not a valid LFA for ABR1 because of the equal cost multi paths (ECMP) {P2, PE, P1, ABR1} and {P2, ABR2, ABR1} causing some of the traffic to be looped back through the router PE (no valid LFA). However for ABR2, which is also a prefix originator for 10.0.1.0/24, P2 is a valid LFA because the only path is {P2, ABR2}.

Configuring Per-Prefix LFA for OSPF

Per prefix LFA is a mechanism by which LFA to a non-best prefix originator can be used in lieu of the LFA to the best prefix originator to provide local repair. In such cases, per prefix LFA can be used to increase the local repair coverage for the OSPF protocol.

Loop Free Alternates (LFA) is a mechanism by which a neighbor can be used as a backup next hop to provide a local repair path for the traffic to flow temporarily in case of failures in the primary next hop (node or link). For this the basic requirement is that the selected backup neighbor provides a loop free path with respect to primary next hop towards a destination originating a set of IGP prefixes. In certain topologies and usage scenarios, it may be possible that multiple destinations are originating the same prefix and there is no viable LFA to the best prefix originator, whilst a non-best prefix originator has one.

Per prefix LFA is a mechanism by which LFA to a non-best prefix originator can be used in lieu of the LFA to the best prefix originator to provide local repair. In such cases, per prefix LFA can be used to increase the local repair coverage for the OSPF protocol.

To configure per prefix LFA for an OSPF interface:

- Configure the per-prefix-calculation configuration statement at the [edit protocols (ospf | ospf3) backup-spf-options] hierarchy level.

Loop-Free Alternate Routes for OSPF Overview

Support for OSPF loop-free alternate routes essentially adds IP fast-reroute capability for OSPF. Junos OS precomputes loop-free backup routes for all OSPF routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives precomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path first to the destination traverses the routing device that is not used as a backup route to that destination. To determine loop-free alternate paths for OSPF routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free routes on any OSPF interface. Because it is common practice to enable LDP on an interface for which OSPF is already enabled, this feature also provides support for LDP label-switched paths (LSPs.)



NOTE: If you enable support for alternate loop-free routes on an interface configured for both LDP and OSPF, you can use the `traceroute` command to trace the active path to the primary next hop.

The level of backup coverage available through OSPF routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSP paths.

Junos OS provides three mechanisms for route redundancy for OSPF through alternate loop-free routes:

- Link protection—Offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

- Node-link protection—Establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device.
- Per-prefix loop-free alternates (LFAs)—It is a technology by which a neighbor can be used as a backup next hop to provide a local repair path for the traffic to flow temporarily in case of failures in the primary next hop (node or link). For this, the basic requirement is that the selected backup neighbor provides a loop-free path with respect to a primary next hop towards a destination, originating a set of interior gateway protocol (IGP) prefixes.

In certain topologies and usage scenarios, it may be possible that multiple destinations are originating the same prefix and there is no viable LFA to the best prefix originator, while a non-best prefix originator has a viable LFA. *Per-prefix LFA* is a mechanism by which LFA to a non-best prefix originator can be used in lieu of the LFA to the best prefix originator to provide local repair. In such cases, per prefix LFA can be used to increase the local repair coverage for the OSPF protocol.

When you enable link protection or node-link protection on an OSPF interface, Junos OS creates an alternate path to the primary next hop for all destination routes that traverse a protected interface.

Example: Configuring Loop-Free Alternate Routes for OSPF

IN THIS SECTION

- [Requirements | 383](#)
- [Overview | 383](#)
- [Configuration | 384](#)
- [Verification | 406](#)

This example demonstrates the use of link protection for interfaces that have OSPF enabled.

When you enable link protection, Junos OS creates an alternate path to the primary next hop for all destination routes that traverse a protected interface. Use link protection when you assume that only a single link might become unavailable but that the neighboring node would still be available through another interface.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

IN THIS SECTION

- [Topology | 383](#)

In this example, six OSPF neighbors are configured with link protection. This causes Junos OS to create an alternate path to the primary next hop for all destination routes that traverse each protected interface. Link protection is used here because even if a link becomes unavailable, the neighboring node would still be available through another interface.


The example shows two topologies. One is the default topology, and the other is the voice topology. For more information about multitopology routing, see the [Multitopology Routing User Guide](#).

The example also includes RSVP LSPs configured as backup LSPs for protected OSPF interfaces.

Topology

[Figure 25 on page 384](#) shows the sample network.

Figure 25: OSPF Link Protection



"CLI Quick Configuration" on page 384 shows the configuration for all of the devices in Figure 25 on page 384.

The section "No Link Title" on page 401 describes the steps on Device R1.

Configuration

IN THIS SECTION

-  CLI Quick Configuration | 384
-  Procedure | 400

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R1

```
set interfaces so-0/2/2 unit 0 description to-R2

set interfaces so-0/2/2 unit 0 family inet address 192.168.242.1/30

set interfaces so-0/2/2 unit 0 family mpls

set interfaces t1-0/1/2 unit 0 description to-R2

set interfaces t1-0/1/2 unit 0 family inet address 192.168.241.1/30

set interfaces t1-0/1/2 unit 0 family mpls

set interfaces t1-0/1/0 unit 0 description to-R4

set interfaces t1-0/1/0 unit 0 family inet address 192.168.241.17/30

set interfaces t1-0/1/0 unit 0 family mpls

set interfaces so-0/2/0 unit 0 description to-R4

set interfaces so-0/2/0 unit 0 family inet address 192.168.242.17/30

set interfaces so-0/2/0 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.1/32 primary
```

```
set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable

set protocols mpls label-switched-path path1 backup

set protocols mpls label-switched-path path1 to 10.255.164.3

set protocols mpls label-switched-path path2 backup

set protocols mpls label-switched-path path2 to 10.255.164.3

set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf topology voice topology-id 32

set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 metric 10
```

```
set protocols ospf area 0.0.0.0 interface so-0/2/2.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 metric 10

set protocols ldp interface all

set protocols ldp interface fxp0.0 disable

set policy-options policy-statement pplb then load-balance per-packet

set routing-options forwarding-table export pplb

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Device R2

```
set interfaces so-0/2/2 unit 0 description to-R1

set interfaces so-0/2/2 unit 0 family inet address 192.168.242.2/30

set interfaces so-0/2/2 unit 0 family mpls

set interfaces t1-0/1/2 unit 0 description to-R1

set interfaces t1-0/1/2 unit 0 family inet address 192.168.241.2/30

set interfaces t1-0/1/2 unit 0 family mpls

set interfaces so-0/2/0 unit 0 description to-R5

set interfaces so-0/2/0 unit 0 family inet address 192.168.242.21/30

set interfaces so-0/2/0 unit 0 family mpls

set interfaces so-0/2/1 unit 0 description to-R3

set interfaces so-0/2/1 unit 0 family inet address 192.168.242.5/30

set interfaces so-0/2/1 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.2/32 primary

set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable
```

```
set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf topology voice topology-id 32

set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 metric 10

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 metric 10

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 metric 10
```

```
set protocols ldp interface all

set protocols ldp interface fxp0.0 disable

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Device R3

```
set interfaces t1-0/1/2 unit 0 description to-R6

set interfaces t1-0/1/2 unit 0 family inet address 192.168.241.25/30

set interfaces t1-0/1/2 unit 0 family mpls

set interfaces so-0/2/1 unit 0 description to-R2

set interfaces so-0/2/1 unit 0 family inet address 192.168.242.6/30

set interfaces so-0/2/1 unit 0 family mpls

set interfaces so-0/2/0 unit 0 description to-R6

set interfaces so-0/2/0 unit 0 family inet address 192.168.242.25/30
```

```
set interfaces so-0/2/0 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.3/32 primary

set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable

set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf traceoptions file ospf

set protocols ospf traceoptions file size 5m

set protocols ospf traceoptions file world-readable

set protocols ospf traceoptions flag error

set protocols ospf topology voice topology-id 32

set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

```
set protocols ospf area 0.0.0.0 interface so-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 metric 5

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 metric 10

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 metric 10

set protocols ldp interface all

set protocols ldp interface fxp0.0 disable

set routing-options static route 11.3.1.0/24 discard

set routing-options static route 11.3.2.0/24 discard

set routing-options static route 11.3.3.0/24 discard

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Device R4

```
set interfaces t1-0/1/0 unit 0 description to-R1

set interfaces t1-0/1/0 unit 0 family inet address 192.168.241.18/30

set interfaces t1-0/1/0 unit 0 family mpls

set interfaces so-0/2/0 unit 0 description to-R1

set interfaces so-0/2/0 unit 0 family inet address 192.168.242.18/30

set interfaces so-0/2/0 unit 0 family mpls

set interfaces t1-0/1/2 unit 0 description to-R5

set interfaces t1-0/1/2 unit 0 family inet address 192.168.241.9/30

set interfaces t1-0/1/2 unit 0 family mpls

set interfaces so-0/2/2 unit 0 description to-R5

set interfaces so-0/2/2 unit 0 family inet address 192.168.242.9/30

set interfaces so-0/2/2 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.4/32 primary
```

```
set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable

set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf topology voice topology-id 32

set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 metric 10

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 metric 10
```

```
set protocols ospf area 0.0.0.0 interface so-0/2/2.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 metric 10

set protocols ldp interface all

set protocols ldp interface fxp0.0 disable

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Device R5

```
set interfaces t1-0/1/2 unit 0 description to-R4

set interfaces t1-0/1/2 unit 0 family inet address 192.168.241.10/30

set interfaces t1-0/1/2 unit 0 family mpls

set interfaces s0-0/2/0 unit 0 description to-R2

set interfaces s0-0/2/0 unit 0 family inet address 192.168.242.22/30

set interfaces s0-0/2/0 unit 0 family mpls
```

```
set interfaces so-0/2/2 unit 0 description to-R4

set interfaces so-0/2/2 unit 0 family inet address 192.168.242.10/30

set interfaces so-0/2/2 unit 0 family mpls

set interfaces so-0/2/1 unit 0 description to-R6

set interfaces so-0/2/1 unit 0 family inet address 192.168.242.13/30

set interfaces so-0/2/1 unit 0 family mpls

set interfaces t1-0/2/1 unit 0 description to-R6

set interfaces t1-0/2/1 unit 0 family inet address 192.168.241.13/30

set interfaces t1-0/2/1 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.5/32 primary

set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable

set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf topology voice topology-id 32
```

```
set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 metric 5

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/2.0 metric 10

set protocols ospf area 0.0.0.0 interface s0-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface s0-0/2/0.0 metric 10

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/2.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/2/1.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/2/1.0 metric 10

set protocols ldp interface all
```

```
set protocols ldp interface fxp0.0 disable

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Device R6

```
set interfaces so-0/2/0 unit 0 description to-R5

set interfaces so-0/2/0 unit 0 family inet address 192.168.242.14/30

set interfaces so-0/2/0 unit 0 family mpls

set interfaces t1-0/1/0 unit 0 description to-R5

set interfaces t1-0/1/0 unit 0 family inet address 192.168.241.14/30

set interfaces t1-0/1/0 unit 0 family mpls

set interfaces t1-0/1/1 unit 0 description to-R3

set interfaces t1-0/1/1 unit 0 family inet address 192.168.241.26/30

set interfaces t1-0/1/1 unit 0 family mpls
```

```
set interfaces so-0/2/1 unit 0 description to-R3

set interfaces so-0/2/1 unit 0 family inet address 192.168.242.26/30

set interfaces so-0/2/1 unit 0 family mpls

set interfaces lo0 unit 0 family inet address 10.255.164.6/32 primary

set protocols rsvp interface all link-protection

set protocols rsvp interface fxp0.0 disable

set protocols mpls interface all

set protocols mpls interface fxp0.0 disable

set protocols ospf topology voice topology-id 32

set protocols ospf traffic-engineering

set protocols ospf area 0.0.0.0 interface fxp0.0 disable

set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/1.0 metric 5
```

```
set protocols ospf area 0.0.0.0 interface so-0/2/0.0 link-protection

set protocols ospf area 0.0.0.0 interface so-0/2/0.0 metric 5

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/0.0 metric 10

set protocols ospf area 0.0.0.0 interface t1-0/1/1.0 link-protection

set protocols ospf area 0.0.0.0 interface t1-0/1/1.0 metric 10

set protocols ldp interface all

set protocols ldp interface fxp0.0 disable

set routing-options topologies family inet topology voice

set routing-options forwarding-table indirect-next-hop-change-acknowledgements
```

Procedure

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device R1:

1. Configure the device interfaces.

```
[edit interfaces]
user@R1# set so-0/2/2 unit 0 description to-R2
user@R1# set so-0/2/2 unit 0 family inet address 192.168.242.1/30
user@R1# set so-0/2/2 unit 0 family mpls
user@R1# set t1-0/1/2 unit 0 description to-R2
user@R1# set t1-0/1/2 unit 0 family inet address 192.168.241.1/30
user@R1# set t1-0/1/2 unit 0 family mpls
user@R1# set t1-0/1/0 unit 0 description to-R4
user@R1# set t1-0/1/0 unit 0 family inet address 192.168.241.17/30
user@R1# set t1-0/1/0 unit 0 family mpls
user@R1# set so-0/2/0 unit 0 description to-R4
user@R1# set so-0/2/0 unit 0 family inet address 192.168.242.17/30
user@R1# set so-0/2/0 unit 0 family mpls
user@R1# set lo0 unit 0 family inet address 10.255.164.1/32 primary
```

2. Extend backup coverage to include RSVP LSP paths.

```
[edit protocols rsvp]
user@R1# set interface all link-protection
user@R1# set interface fxp0.0 disable
```

3. Enable MPLS on the interfaces, and configure backup LSPs to Device R3.

```
[edit protocols mpls]
user@R1# set interface all
user@R1# set interface fxp0.0 disable
user@R1# set label-switched-path path1 backup
user@R1# set label-switched-path path1 to 10.255.164.3
user@R1# set label-switched-path path2 backup
user@R1# set label-switched-path path2 to 10.255.164.3
```

4. Configure OSPF connections, link metrics, and link protection.

```
[edit protocols ospf]
user@R1# set traffic-engineering
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fxp0.0 disable
user@R1# set interface lo0.0 passive
user@R1# set interface so-0/2/0.0 link-protection
user@R1# set interface so-0/2/0.0 metric 10
user@R1# set interface so-0/2/2.0 link-protection
user@R1# set interface so-0/2/2.0 metric 10
user@R1# set interface t1-0/1/0.0 link-protection
user@R1# set interface t1-0/1/0.0 metric 10
user@R1# set interface t1-0/1/2.0 link-protection
user@R1# set interface t1-0/1/2.0 metric 10
```

5. (Optional) Configure a specific OSPF topology for voice traffic.

```
[edit protocols ospf]
user@R1# set topology voice topology-id 32
[edit routing-options topologies family inet]
user@R1# set topology voice
```

6. Enable LDP on the interfaces.

```
[edit protocols ldp]
user@R1# set interface all
user@R1# set interface fxp0.0 disable
```

7. (Optional) Configure per-packet load balancing.

```
[edit policy-options policy-statement pplb]
user@R1# set then load-balance per-packet
[edit routing-options forwarding-table]
user@R1# set export pplb
```

8. Configure the routing protocol process (rpd) to request an acknowledgement when creating a new forwarding next hop.

We recommend that the `indirect-next-hop-change-acknowledgements` statement be configured when protection mechanisms are being used. This includes MPLS RSVP protection such as fast reroute (FRR) as well as interior gateway protocol (IGP) loop-free alternate (LFA) link or node protection.

```
[edit routing-options forwarding-table]
user@R1# set indirect-next-hop-change-acknowledgements
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
so-0/2/2 {
  unit 0 {
    description to-R2;
    family inet {
      address 192.168.242.1/30;
    }
    family mpls;
```

```
    }  
  }  
  t1-0/1/2 {  
    unit 0 {  
      description to-R2;  
      family inet {  
        address 192.168.241.1/30;  
      }  
      family mpls;  
    }  
  }  
  t1-0/1/0 {  
    unit 05 {  
      description to-R4;  
      family inet {  
        address 192.168.241.17/30;  
      }  
      family mpls;  
    }  
  }  
  so-0/2/0 {  
    unit 0 {  
      description to-R4;  
      family inet {  
        address 192.168.242.17/30;  
      }  
      family mpls;  
    }  
  }  
  lo0 {  
    unit 0 {  
      family inet {  
        address 10.255.164.1/32 {  
          primary;  
        }  
      }  
    }  
  }  
}
```

```
user@R1# show protocols  
rsvp {
```

```
interface all {
    link-protection;
}
interface fxp0.0 {
    disable;
}
}
mpls {
    label-switched-path path1 {
        backup;
        to 10.255.164.3;
    }
    label-switched-path path2 {
        backup;
        to 10.255.164.3;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
ospf {
    topology voice topology-id 32;
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
        interface so-0/2/0.0 {
            link-protection;
            metric 10;
        }
        interface so-0/2/2.0 {
            link-protection;
            metric 10;
        }
        interface t1-0/1/0.0 {
            link-protection;
            metric 10;
        }
    }
}
```

```
        interface t1-0/1/2.0 {
            link-protection;
            metric 10;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
```

```
user@R1# show policy-options
policy-statement pplb {
    then {
        load-balance per-packet;
    }
}
```

```
user@R1# show routing-options
forwarding-table {
    export pplb;
    indirect-next-hop-change-acknowledgements;
}
topologies {
    family inet {
        topology voice;
    }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

 [Verifying the Routes on Device R1 | 407](#)


```

                to 192.168.242.18 via so-0/2/0.0
                to 192.168.241.18 via t1-0/1/0.0
192.168.241.8/30  *[OSPF/10] 1d 23:34:00, metric 20
                > to 192.168.242.18 via so-0/2/0.0
                to 192.168.241.18 via t1-0/1/0.0
192.168.241.12/30 *[OSPF/10] 1d 23:34:00, metric 30
                to 192.168.242.2 via so-0/2/2.0
                to 192.168.241.2 via t1-0/1/2.0
                to 192.168.242.18 via so-0/2/0.0
                > to 192.168.241.18 via t1-0/1/0.0
192.168.241.24/30 *[OSPF/10] 1d 23:34:00, metric 30
                to 192.168.242.2 via so-0/2/2.0
                > to 192.168.241.2 via t1-0/1/2.0
192.168.242.4/30  *[OSPF/10] 1d 23:34:00, metric 20
                to 192.168.242.2 via so-0/2/2.0
                > to 192.168.241.2 via t1-0/1/2.0
192.168.242.8/30  *[OSPF/10] 1d 23:34:00, metric 20
                > to 192.168.242.18 via so-0/2/0.0
                to 192.168.241.18 via t1-0/1/0.0
192.168.242.12/30 *[OSPF/10] 1d 23:34:00, metric 25
                to 192.168.242.2 via so-0/2/2.0
                > to 192.168.241.2 via t1-0/1/2.0
                to 192.168.242.18 via so-0/2/0.0
                to 192.168.241.18 via t1-0/1/0.0
192.168.242.20/30 *[OSPF/10] 1d 23:34:00, metric 20
                > to 192.168.242.2 via so-0/2/2.0
                to 192.168.241.2 via t1-0/1/2.0
192.168.242.24/30 *[OSPF/10] 1d 23:34:00, metric 25
                to 192.168.242.2 via so-0/2/2.0
                > to 192.168.241.2 via t1-0/1/2.0
224.0.0.5/32     *[OSPF/10] 1w1d 02:46:58, metric 1
                MultiRecv

inet.3: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)

:voice.inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.164.2/32  *[OSPF/10] 1d 23:34:00, metric 10
                > to 192.168.242.2 via so-0/2/2.0
                to 192.168.241.2 via t1-0/1/2.0
10.255.164.3/32  *[OSPF/10] 1d 23:34:00, metric 20
                > to 192.168.242.2 via so-0/2/2.0

```

```

    to 192.168.241.2 via t1-0/1/2.0
10.255.164.4/32  * [OSPF/10] 1d 23:34:00, metric 10
    to 192.168.242.18 via so-0/2/0.0
    > to 192.168.241.18 via t1-0/1/0.0
10.255.164.5/32  * [OSPF/10] 1d 23:34:00, metric 20
    to 192.168.242.2 via so-0/2/2.0
    to 192.168.241.2 via t1-0/1/2.0
    > to 192.168.242.18 via so-0/2/0.0
    to 192.168.241.18 via t1-0/1/0.0
10.255.164.6/32  * [OSPF/10] 1d 23:34:00, metric 25
    to 192.168.242.2 via so-0/2/2.0
    to 192.168.241.2 via t1-0/1/2.0
    > to 192.168.242.18 via so-0/2/0.0
    to 192.168.241.18 via t1-0/1/0.0
192.168.241.8/30 * [OSPF/10] 1d 23:34:00, metric 20
    > to 192.168.242.18 via so-0/2/0.0
    to 192.168.241.18 via t1-0/1/0.0
192.168.241.12/30 * [OSPF/10] 1d 23:34:00, metric 30
    > to 192.168.242.2 via so-0/2/2.0
    to 192.168.241.2 via t1-0/1/2.0
    to 192.168.242.18 via so-0/2/0.0
    to 192.168.241.18 via t1-0/1/0.0
192.168.241.24/30 * [OSPF/10] 1d 23:34:00, metric 30
    to 192.168.242.2 via so-0/2/2.0
    > to 192.168.241.2 via t1-0/1/2.0
192.168.242.4/30 * [OSPF/10] 1d 23:34:00, metric 20
    to 192.168.242.2 via so-0/2/2.0
    > to 192.168.241.2 via t1-0/1/2.0
192.168.242.8/30 * [OSPF/10] 1d 23:34:00, metric 20
    to 192.168.242.18 via so-0/2/0.0
    > to 192.168.241.18 via t1-0/1/0.0
192.168.242.12/30 * [OSPF/10] 1d 23:34:00, metric 25
    to 192.168.242.2 via so-0/2/2.0
    to 192.168.241.2 via t1-0/1/2.0
    > to 192.168.242.18 via so-0/2/0.0
    to 192.168.241.18 via t1-0/1/0.0
192.168.242.20/30 * [OSPF/10] 1d 23:34:00, metric 20
    to 192.168.242.2 via so-0/2/2.0
    > to 192.168.241.2 via t1-0/1/2.0
192.168.242.24/30 * [OSPF/10] 1d 23:34:00, metric 25
    > to 192.168.242.2 via so-0/2/2.0
    to 192.168.241.2 via t1-0/1/2.0

```

```
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
```

Meaning

As expected, Device R1 has multiple potential routes to each destination.

Checking the Backup Coverage

Purpose

On Device R1, use the `show (ospf | ospf3) backup coverage` command to check the level of backup coverage available for all the nodes and prefixes in the network.

Action

```
user@R1> show ospf backup coverage
```

```
Topology default coverage:
```

```
Node Coverage:
```

Area	Covered Nodes	Total Nodes	Percent Covered
0.0.0.0	5	5	100.00%

```
Route Coverage:
```

Path Type	Covered Routes	Total Routes	Percent Covered
Intra	17	18	94.44%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	0	0	100.00%
All	17	18	94.44%

```
Topology voice coverage:
```

```
Node Coverage:
```

Area	Covered Nodes	Total Nodes	Percent Covered
0.0.0.0	5	5	100.00%

Route Coverage:

Path Type	Covered Routes	Total Routes	Percent Covered
Intra	17	18	94.44%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	0	0	100.00%
All	17	18	94.44%

Checking the Backup LSPs

Purpose

On Device R1, use the `show (ospf | ospf3) backup lsp` command to check LSPs designated as backup routes for OSPF routes.

Action

```
user@R1> show ospf backup lsp

path1
  Egress: 10.255.164.3, Status: up, Last change: 01:13:48
  TE-metric: 19, Metric: 0
path2
  Egress: 10.255.164.3, Status: up, Last change: 01:13:48
  TE-metric: 19, Metric: 0
```

Checking the Backup Neighbors

Purpose

On Device R1, use the `show (ospf | ospf3) backup neighbor` command to check the neighbors through which direct next hops for the backup paths are available.

Action

```
user@R1> show ospf backup neighbor

Topology default backup neighbors:

Area 0.0.0.0 backup neighbors:

10.255.164.4
  Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10
  Direct next-hop: so-0/2/0.0 via 192.168.242.18
  Direct next-hop: t1-0/1/0.0 via 192.168.241.18

10.255.164.2
  Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10
  Direct next-hop: so-0/2/2.0 via 192.168.242.2
  Direct next-hop: t1-0/1/2.0 via 192.168.241.2

10.255.164.3 (LSP endpoint)
  Neighbor to Self Metric: 20
  Self to Neighbor Metric: 20
  Direct next-hop: path1
  Direct next-hop: path2

Topology voice backup neighbors:

Area 0.0.0.0 backup neighbors:

10.255.164.4
  Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10
  Direct next-hop: so-0/2/0.0 via 192.168.242.18
  Direct next-hop: t1-0/1/0.0 via 192.168.241.18

10.255.164.2
  Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10
  Direct next-hop: so-0/2/2.0 via 192.168.242.2
  Direct next-hop: t1-0/1/2.0 via 192.168.241.2
```

```

10.255.164.3 (LSP endpoint)
  Neighbor to Self Metric: 20
  Self to Neighbor Metric: 20
  Direct next-hop: path1
  Direct next-hop: path2

```

Checking the SPF Calculations

Purpose

On Device R1, use the `show (ospf | ospf3) backup spf detail` command to check OSPF shortest-path-first (SPF) calculations for backup paths. To limit the output, the voice topology is specified in the command.

Action

```

user@R1> show ospf backup spf detail topology voice

Topology voice results:

Area 0.0.0.0 results:

192.168.241.2
  Self to Destination Metric: 10
  Parent Node: 10.255.164.1
  Primary next-hop: t1-0/1/2.0
  Backup next-hop: path1
  Backup Neighbor: 10.255.164.3 (LSP endpoint)
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
    Self to Neighbor Metric: 20, Backup preference: 0x0
    Track Item: 10.255.164.2
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.164.2
    Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Not evaluated, Reason: Interface is already covered
  Backup Neighbor: 10.255.164.4
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Track Item: 10.255.164.1
    Not evaluated, Reason: Interface is already covered

```

192.168.241.18

Self to Destination Metric: 10
Parent Node: 10.255.164.1
Primary next-hop: t1-0/1/0.0
Backup next-hop: so-0/2/0.0 via 192.168.242.18
Backup Neighbor: 10.255.164.3 (LSP endpoint)
Neighbor to Destination Metric: 30, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Track Item: 10.255.164.1
Track Item: 10.255.164.2
Track Item: 10.255.164.4
Not eligible, Reason: Path loops
Backup Neighbor: 10.255.164.4
Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.164.2
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Track Item: 10.255.164.1
Not evaluated, Reason: Interface is already covered

192.168.242.2

Self to Destination Metric: 10
Parent Node: 10.255.164.1
Primary next-hop: so-0/2/2.0
Backup next-hop: path2
Backup Neighbor: 10.255.164.3 (LSP endpoint)
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Track Item: 10.255.164.2
Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.164.2
Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Interface is already covered
Backup Neighbor: 10.255.164.4
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Track Item: 10.255.164.1
Not evaluated, Reason: Interface is already covered

192.168.242.18

Self to Destination Metric: 10
 Parent Node: 10.255.164.1
 Primary next-hop: so-0/2/0.0
 Backup next-hop: t1-0/1/0.0 via 192.168.241.18
 Backup Neighbor: 10.255.164.3 (LSP endpoint)
 Neighbor to Destination Metric: 30, Neighbor to Self Metric: 20
 Self to Neighbor Metric: 20, Backup preference: 0x0
 Track Item: 10.255.164.1
 Track Item: 10.255.164.2
 Track Item: 10.255.164.4
 Not eligible, Reason: Path loops
 Backup Neighbor: 10.255.164.4
 Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
 Self to Neighbor Metric: 10, Backup preference: 0x0
 Eligible, Reason: Contributes backup next-hop
 Backup Neighbor: 10.255.164.2
 Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
 Self to Neighbor Metric: 10, Backup preference: 0x0
 Track Item: 10.255.164.1
 Not evaluated, Reason: Interface is already covered

10.255.164.2

Self to Destination Metric: 10
 Parent Node: 192.168.241.2
 Parent Node: 192.168.242.2
 Primary next-hop: so-0/2/2.0 via 192.168.242.2
 Primary next-hop: t1-0/1/2.0 via 192.168.241.2
 Backup Neighbor: 10.255.164.3 (LSP endpoint)
 Neighbor to Destination Metric: 10, Neighbor to Self Metric: 20
 Self to Neighbor Metric: 20, Backup preference: 0x0
 Track Item: 10.255.164.2
 Not evaluated, Reason: Primary next-hop multipath
 Backup Neighbor: 10.255.164.2
 Neighbor to Destination Metric: 0, Neighbor to Self Metric: 10
 Self to Neighbor Metric: 10, Backup preference: 0x0
 Not evaluated, Reason: Primary next-hop multipath
 Backup Neighbor: 10.255.164.4
 Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
 Self to Neighbor Metric: 10, Backup preference: 0x0
 Track Item: 10.255.164.1
 Track Item: 10.255.164.2
 Not evaluated, Reason: Primary next-hop multipath

10.255.164.4

Self to Destination Metric: 10
Parent Node: 192.168.241.18
Parent Node: 192.168.242.18
Primary next-hop: so-0/2/0.0 via 192.168.242.18
Primary next-hop: t1-0/1/0.0 via 192.168.241.18
Backup Neighbor: 10.255.164.3 (LSP endpoint)
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Track Item: 10.255.164.4
Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.4
Neighbor to Destination Metric: 0, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.2
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Track Item: 10.255.164.1
Track Item: 10.255.164.4
Not evaluated, Reason: Primary next-hop multipath

192.168.241.10

Self to Destination Metric: 20
Parent Node: 10.255.164.4
Primary next-hop: so-0/2/0.0 via 192.168.242.18
Primary next-hop: t1-0/1/0.0 via 192.168.241.18
Backup Neighbor: 10.255.164.3 (LSP endpoint)
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.4
Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.2
Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

192.168.242.6

Self to Destination Metric: 20
Parent Node: 10.255.164.2

```
Primary next-hop: so-0/2/2.0 via 192.168.242.2
Primary next-hop: t1-0/1/2.0 via 192.168.241.2
Backup Neighbor: 10.255.164.3 (LSP endpoint)
  Neighbor to Destination Metric: 10, Neighbor to Self Metric: 20
  Self to Neighbor Metric: 20, Backup preference: 0x0
  Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.2
  Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10, Backup preference: 0x0
  Not evaluated, Reason: Primary next-hop multipath
Backup Neighbor: 10.255.164.4
  Neighbor to Destination Metric: 30, Neighbor to Self Metric: 10
  Self to Neighbor Metric: 10, Backup preference: 0x0
  Track Item: 10.255.164.1
  Track Item: 10.255.164.2
  Not evaluated, Reason: Primary next-hop multipath

192.168.242.10
  Self to Destination Metric: 20
  Parent Node: 10.255.164.4
  Primary next-hop: so-0/2/0.0 via 192.168.242.18
  Primary next-hop: t1-0/1/0.0 via 192.168.241.18
  Backup Neighbor: 10.255.164.3 (LSP endpoint)
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
    Self to Neighbor Metric: 20, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath
  Backup Neighbor: 10.255.164.4
    Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath
  Backup Neighbor: 10.255.164.2
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath

192.168.242.22
  Self to Destination Metric: 20
  Parent Node: 10.255.164.2
  Primary next-hop: so-0/2/2.0 via 192.168.242.2
  Primary next-hop: t1-0/1/2.0 via 192.168.241.2
  Backup Neighbor: 10.255.164.3 (LSP endpoint)
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 20
    Self to Neighbor Metric: 20, Backup preference: 0x0
```

Track Item: 10.255.164.2

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

10.255.164.3

Self to Destination Metric: 20

Parent Node: 192.168.242.6

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 0, Neighbor to Self Metric: 20

Self to Neighbor Metric: 20, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

10.255.164.5

Self to Destination Metric: 20

Parent Node: 192.168.241.10

Parent Node: 192.168.242.10

Parent Node: 192.168.242.22

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Primary next-hop: so-0/2/0.0 via 192.168.242.18

Primary next-hop: t1-0/1/0.0 via 192.168.241.18

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 20

Self to Neighbor Metric: 20, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

192.168.242.14

Self to Destination Metric: 25

Parent Node: 10.255.164.5

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Primary next-hop: so-0/2/0.0 via 192.168.242.18

Primary next-hop: t1-0/1/0.0 via 192.168.241.18

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 10, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

192.168.242.26

Self to Destination Metric: 25

Parent Node: 10.255.164.3

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 5, Neighbor to Self Metric: 20
Self to Neighbor Metric: 20, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 10
Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0
Not evaluated, Reason: Primary next-hop multipath

10.255.164.6

Self to Destination Metric: 25

Parent Node: 192.168.242.14

Parent Node: 192.168.242.26

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Primary next-hop: so-0/2/0.0 via 192.168.242.18

Primary next-hop: t1-0/1/0.0 via 192.168.241.18

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 5, Neighbor to Self Metric: 20

Self to Neighbor Metric: 20, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

192.168.241.14

Self to Destination Metric: 30

Parent Node: 10.255.164.5

Primary next-hop: so-0/2/2.0 via 192.168.242.2

Primary next-hop: t1-0/1/2.0 via 192.168.241.2

Primary next-hop: so-0/2/0.0 via 192.168.242.18

Primary next-hop: t1-0/1/0.0 via 192.168.241.18

Backup Neighbor: 10.255.164.3 (LSP endpoint)

Neighbor to Destination Metric: 15, Neighbor to Self Metric: 20

Self to Neighbor Metric: 20, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.2

Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

Backup Neighbor: 10.255.164.4

Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10

Self to Neighbor Metric: 10, Backup preference: 0x0

Not evaluated, Reason: Primary next-hop multipath

```

192.168.241.26
  Self to Destination Metric: 30
  Parent Node: 10.255.164.3
  Primary next-hop: so-0/2/2.0 via 192.168.242.2
  Primary next-hop: t1-0/1/2.0 via 192.168.241.2
  Backup Neighbor: 10.255.164.3 (LSP endpoint)
    Neighbor to Destination Metric: 10, Neighbor to Self Metric: 20
    Self to Neighbor Metric: 20, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath
  Backup Neighbor: 10.255.164.2
    Neighbor to Destination Metric: 20, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath
  Backup Neighbor: 10.255.164.4
    Neighbor to Destination Metric: 25, Neighbor to Self Metric: 10
    Self to Neighbor Metric: 10, Backup preference: 0x0
    Not evaluated, Reason: Primary next-hop multipath

```

SEE ALSO

[Example: Configuring Multitopology Routing for Class-Based Forwarding of Voice, Video, and Data Traffic](#)

Configuring Link Protection for OSPF

You can configure link protection for any interface for which OSPF is enabled. When you enable link protection, Junos OS creates an alternate path to the primary next hop for all destination routes that traverse a protected interface. Use link protection when you assume that only a single link might become unavailable but that the neighboring node would still be available through another interface.

Link protection is supported on:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 unicast realms
- OSPFv2 unicast topologies, except for multicast topologies
- All routing instances supported by OSPFv2 and OSPFv3

- Logical systems

To configure link protection for an OSPF interface:

- Include the `link-protection` statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-name]` hierarchy level.



BEST PRACTICE: When you configure link protection for OSPF, you must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

In the following example, the OSPF interface **so-0/0/0.0** in area 0.0.0.0 is configured for link protection. If a link for a destination route that traverses this interface becomes unavailable, Junos OS creates a loop-free backup path through another interface on the neighboring node, thus avoiding the link that is no longer available.

```
[edit]
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/0.0 {
        link-protection;
      }
    }
  }
}
```

SEE ALSO

| [link-protection](#)

Configuring Node-Link Protection for OSPF

You can configure node-link protection on any interface for which OSPF is enabled. Node-link protection establishes an alternative path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.

Node-link protection is supported on:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 unicast realms
- OSPFv2 unicast topologies
- All routing instances supported by OSPFv2 and OSPFv3
- Logical systems

To configure node-link protection for an OSPF interface:

- Include the `node-link-protection` statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-name]` hierarchy level.



BEST PRACTICE: You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

In the following example, the OSPF interface **so-0/0/0.0** in area 0.0.0.0 is configured for node-link protection. If a link for a destination route that traverses this interface becomes unavailable, Junos OS creates a loop-free backup path through a different routing device altogether, thus avoiding the primary next-hop routing device.

```
[edit]
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/0.0 {
        node-link-protection;
      }
    }
  }
}
```

Configuring Node to Link Protection Fallback for OSPF

You can configure link protection for any interface for which OSPF is enabled. When you enable link protection, Junos OS creates an alternate path to the primary next hop for all destination routes that

traverse a protected interface. Use link protection when you assume that only a single link might become unavailable but that the neighboring node would still be available through another interface.

You can configure node-link protection on any interface for which OSPF is enabled. Node-link protection establishes an alternative path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.

In certain topologies it may be desirable to have local repair protection to node failures in the primary next hop, which may not be available. In that case, to ensure that some level of local repair capabilities exist, a fallback mechanism is required. Since the link protection is less stringent than node protection, it may be possible that link protection exists and provide the same to those destination (and hence the prefixes originated by it).

To configure node to link protection fallback for an OSPF interface:

- Include the `node-link-degradation` statement at the `[edit protocols (ospf | ospf3) backup-spf-options]` hierarchy level.

Excluding an OSPF Interface as a Backup for a Protected Interface

By default, all OSPF interfaces that belong to the default instance or to a specific routing instance are eligible as a backup interface for interfaces configured with link-protection or node-link protection. You can specify that any OSPF interface be excluded from functioning as a backup interface to protected interfaces.

To exclude an OSPF interface as a backup interface for a protected interface:

- Include the `no-eligible-backup` statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-name]` hierarchy level.

In the following example, interface `so-0/0/0.0` has been configured to prohibit backup traffic for traffic destined for a protected interface. This means that if a neighboring next-hop path or node for a protected interface fails, interface `so-0/0/0.0` cannot be used to transmit traffic to a backup path.

```
[edit]
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/0.0 {
        no-eligible-backup;
      }
    }
  }
}
```

```

    }
  }
}

```

Configuring Backup SPF Options for Protected OSPF Interfaces

By default, if at least one OSPF interface is configured for link-protection or node-link protection, Junos OS calculates backup next hops for all the topologies in an OSPF instance. You can configure the following backup shortest-path-first (SPF) options to override the default behavior:

- Disable the calculation of backup next hops for an OSPF instance or a specific topology in an instance.
- Prevent the installation of backup next hops in the routing table or the forwarding table for an OSPF instance or a specific topology in an instance.
- Limit the calculation of backup next hops to a subset of paths as defined in RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*.

You can disable the backup SPF algorithm for an OSPF instance or specific topology in an instance. Doing so prevents the calculation of backup next hops for that OSPF instance or topology.

To disable the calculation of backup next hops for an OSPF instance or topology:

- Include the `disable` statement at the `[edit protocols (ospf | ospf3) backup-spf-options]` or `[edit protocols ospf backup-spf-options topology topology-name]` hierarchy level.

In the following example, the calculation of backup next hops is disabled for the OSPF topology **voice**:

```

[edit]
protocols {
  ospf {
    topology voice {
      backup-spf-options {
        disable;
      }
    }
  }
}

```

You can configure the routing device to prevent the installation of backup next hops in the routing table or the forwarding table for an OSPF instance, or a specific topology in an OSPF instance. The SPF algorithm continues to calculate backup next hops, but they are not installed.

To prevent the routing device from installing backup next hops in the routing table or the forwarding table:

- Include the `no-install` statement at the `[edit protocols (ospf | ospf3) backup-spf-options]` or the `[edit protocols ospf topology topology-name]` hierarchy level.

In the following example, backup next hops for the OSPF topology **voice** are not installed in the routing table or forwarding table. Any calculated backup next hops for other OSPF instances or topologies continue to be installed.

```
[edit]
protocols {
  ospf {
    topology voice {
      backup-spf-options {
        no-install;
      }
    }
  }
}
```

You can limit the calculation of backup next hops to *downstream paths*, as defined in RFC 5286. You can specify for Junos OS to use only downstream paths as backup next hops for protected interfaces for an OSPF instance or a specific topology in an OSPF instance. In a downstream path, the distance from the backup neighbor to the destination must be smaller than the distance from the calculating routing device to the destination. Using only downstream paths as loop-free alternate paths for protected interfaces ensures that these paths do not result in microloops. However, you might experience less than optimal backup coverage for your network.

To limit the calculation of backup next hops to downstream paths:

- Include the `downstream-paths-only` statement at the `[edit protocols (ospf | ospf3) backup-spf-options]` or `[edit protocols ospf backup-spf-options topology topology-name]` hierarchy level.

In the following example, only downstream paths are calculated as backup next hops for the topology **voice**:

```
[edit]
protocols {
```

```

ospf {
  topology voice {
    backup-spf-options {
      downstream-paths-only;
    }
  }
}

```

SEE ALSO

| [backup-spf-options](#)

Configuring RSVP Label-Switched Paths as Backup Paths for OSPF

When configuring an OSPF interface for link protection or node-link protection, relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of OSPF and LDP label-switched-paths (LSPs) by configuring RSVP LSPs as backup paths.

When configuring an LSP, you must specify the IP address of the egress router.



NOTE: RSVP LSPs can be used as backup paths only for the default topology for OSPFv2 and not for a configured topology. Additionally, RSVP LSP cannot be used as backup paths for non-default instances for OSPFv2 or OSPFv3.

To configure a specific RSVP LSP as a backup path:

1. Include the backup statement at the [edit protocols mpls labeled-switched-path *lsp-name*] hierarchy level.
2. Specify the address of the egress router by including the to *ip-address* statement at the [edit protocols mpls label-switched-path] hierarchy level.

In the following example, the RSVP LSP **f-to-g** is configured as a backup LSP for protected OSPF interfaces. The egress router is configured with the IP address **192.168.1.4**.

```

[edit]
protocols {
  mpls {
    label-switched-path f-to-g {

```

```

        to 192.168.1.4;
        backup;
    }
}
}

```

Remote LFA over LDP Tunnels in OSPF Networks Overview

In an OSPF network, a loop free alternate (LFA) is a directly connected neighbor that provides precomputed backup paths to the destinations reachable through the protected link on the point of local repair (PLR). A remote LFA is not directly connected to the PLR and provides precomputed backup paths using dynamically created LDP tunnels to the remote LFA node. The PLR uses this remote LFA backup path when the primary link fails. The primary goal of the remote LFA is to increase backup coverage for the OSPF networks and provide protection for Layer 1 metro-rings.

LFAs do not provide full backup coverage for OSPF networks. This is a major setback for metro Ethernet networks that are often shaped as ring topologies. To overcome this setback, Resource Reservation Protocol - Traffic Engineering (RSVP-TE) backup tunnels are commonly used to extend the backup coverage. However, a majority of network providers have already implemented LDP as the MPLS tunnel setup protocol and do not want to implement the RSVP-TE protocol merely for backup coverage. LDP automatically brings up transport tunnels to all potential destinations in an OSPF network and hence is the preferred protocol. The existing LDP implemented for the MPLS tunnel setup can be reused for protection of OSPF networks and subsequent LDP destinations, thereby eliminating the need for RSVP-TE backup tunnels for backup coverage.

To calculate the remote LFA backup path, the OSPF protocol determines the remote LFA node in the following manner:

1. Calculates the reverse shortest path first from the adjacent router across the protected link of a PLR. The reverse shortest path first uses the incoming link metric instead of the outgoing link metric to reach a neighboring node.

The result is a set of links and nodes, which is the shortest path from each leaf node to the root node.

2. Calculates the shortest path first (SPF) on the remaining adjacent routers to find the list of nodes that can be reached without traversing the link being protected.

The result is another set of links and nodes on the shortest path from the root node to all leaf nodes.

3. Determines the common nodes from the above results. These nodes are the remote LFAs.

OSPF listens to the advertised labels for the LDP routes. For each advertised LDP route, OSPF checks whether it contains an LDP supplied next hop. If the corresponding OSPF route does have a backup next hop, then OSPF runs the backup policy and adds an additional tracking route with the corresponding LDP label-switched path next hop as the backup next hop. If there are no backup next hops, LDP builds a dynamic LDP tunnel to the remote LFA, and LDP establishes a targeted adjacency between the remote LFA node and the PLR node. This backup route has two LDP labels. The top label is the OSPF route, which denotes the backup path from the PLR to the remote LFA route. The bottom label is the LDP MPLS label-switched path that denotes the route for reaching the ultimate destination from the remote LFA. When an LDP session goes down and a remote tunnel is no longer available, OSPF changes all the routes that have been using this backup LDP tunnel.



NOTE: Currently, Junos OS supports only IPv4 transport LSPs. If you need to reuse IPv4 transport LSPs for IPv6 IGP networks, add an IPv6 explicit NULL label to the label stack of the tracking route. The system automatically converts the IPv4 LSP to an IPv6 LSP.

LDP might be vulnerable by an automatically targeted adjacency, and these threats can be mitigated using all or some of the following mechanisms:

- Remote LFAs that are several hops away use extended hello messages to indicate willingness to establish a targeted LDP session. A remote LFA can reduce the threat of spoofed extended hello messages by filtering them and accepting only those originating at sources permitted by an access or filter list.
- There is a need to authenticate with TCP-MD5 all auto-targeted LDP sessions in the given IGP/LDP domain using apply groups or LDP global-level authentication.
- As an added security measure, the repair or remote tunnel endpoint routers should be assigned from a set of addresses that are not reachable from outside of the routing domain.

SEE ALSO

[auto-targeted-session](#)

Configuring Remote LFA Backup over LDP Tunnels in an OSPF Network

The primary goal of a remote loop free alternate (LFA) is to increase backup coverage for OSPF routes and provide protection especially for Layer 1 metro-rings. The existing LDP implemented for the MPLS tunnel setup can be reused for protection of OSPF networks and subsequent LDP destinations. The OSPF protocol creates a dynamic LDP tunnel to reach the remote LFA node from the point of local repair (PLR). The PLR uses this remote LFA backup path when the primary link fails.

Before you configure remote LFA over LDP tunnels in an OSPF network, you must do the following:

1. Enable LDP on the loopback interface.

Configure a loopback interface because an LDP targeted adjacency cannot be formed without a loopback interface. LDP targeted adjacency is essential for determining remote LFA backup paths.

2. Make sure that remote LFA allows asymmetric remote neighbor discovery—that is, it must send periodic targeted hello messages to the router that initiated the remote neighbor for LDP auto-targeted adjacency.
3. Configure link protection or node-link protection on the PLR.

To configure remote LFA backup over LDP tunnels in an OSPF network:

1. Enable remote LFA backup to determine the backup next hop using dynamic LDP label-switched path.

```
[edit protocols ospf backup-spf-options]
user@host# set remote-backup-calculation
```

2. Enable automatically targeted LDP sessions using the loopback addresses between the PLR and the remote LFA node.

```
[edit protocols ldp]
user@host# set auto-targeted-session
```

3. Specify a time interval for which the targeted LDP sessions are kept up even after the remote LFA node goes down.

```
[edit protocols ldp auto-targeted-session]
user@host# set teardown-delay seconds
```

For example, to set a teardown delay value of 60 seconds:

```
[edit protocols ldp auto-targeted-session]
user@host# set teardown-delay 60
```

4. Specify the maximum number of automatically targeted LDP sessions to optimize memory usage.

```
[edit protocols ldp auto-targeted-session]
user@host# set maximum-sessions number of sessions
```

For example, to set a maximum sessions allowed to 20:

```
[edit protocols ldp auto-targeted-session]
user@host# set maximum-sessions 20
```

SEE ALSO

[auto-targeted-session](#)

[backup-spf-options](#)

Example: Configuring Remote LFA Over LDP Tunnels in OSPF Networks

IN THIS SECTION

- [Requirements | 432](#)
- [Overview | 432](#)
- [Configuration | 433](#)
- [Verification | 443](#)

In an OSPF network, a loop free alternate(LFA) is a directly connected neighbor that provides precomputed backup paths to the destinations reachable via the protected link on the point of local repair (PLR). A remote LFA is not directly connected to the PLR and provides precomputed backup paths using dynamically created LDP tunnels to the remote LFA node. The PLR uses this remote LFA backup path when the primary link fails. The primary goal of the remote LFA is to increase backup coverage for the OSPF networks and provide protection for Layer 1 metro-rings. This example shows how to configure remote LFA for LDP tunnels in an OSPF network for extending backup protection.

Requirements

This example uses the following hardware and software components:

- Nine MX Series routers with OSPF protocol and LDP enabled on the connected interfaces.
- Junos OS Release 15.1 or later running on all devices.

Before you configure remote LFA over LDP tunnels in an OSPF networks, make sure of the following:

- LDP is enabled on the loopback interface. Without a loopback interface, LDP targeted adjacency cannot be formed. Remote LFA cannot be configured without LDP targeted adjacency.
- Remote LFA must allow asymmetric remote neighbor discovery, that is, it must send periodic targeted hellos to the router that initiated the remote neighbor for LDP auto targeted adjacency.
- Link protection or node-link protection must be configured on the point of local repair (PLR).

Overview

IN THIS SECTION


- [Topology | 432](#)

The example includes nine routers in a ring topology. Configure the OSPF protocol on the directly connected interfaces. Device R6 is the PLR. This example verifies that Junos OS updates the routing table of Device R6 with LDP next-hop routes as the backup route.

Topology

In the topology [Figure 26 on page 433](#) shows the remote LFA over LDP tunnels in OSPF networks is configured on Device R6.

Figure 26: Example Remote LFA over LDP Tunnels



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 433](#)
- [Configuring Device R6 | 439](#)
- [Results | 441](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter `commit` from configuration mode.

R0

```

set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.90.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.110.1.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.1.1.1
set routing-options forwarding-table export per-packet
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface lo0.0
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf export static
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp egress-policy static
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept
set policy-options policy-statement static from protocol static
set policy-options policy-statement static then accept

```

R1

```

set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.20.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.80.1.1/24

```

```
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 10.100.1.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.2.2.2/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.2.2.2
set routing-options forwarding-table export per-packet
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface ge-0/0/3.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept
```

R2

```
set interfaces ge-0/0/0 unit 0 family inet address 10.20.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.30.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.110.1.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.3.3.3/32
set interfaces lo0 unit 0 family mpls
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept
```

R3

```
set interfaces ge-0/0/0 unit 0 family inet address 10.30.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.40.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.4.4.4/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.4.4.4
set routing-options forwarding-table export per-packet
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept
```

R4

```
set interfaces ge-0/0/0 unit 0 family inet address 10.40.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.50.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.5.5.5/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.5.5.5
set routing-options forwarding-table export per-packet
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 60
set protocols ldp auto-targeted-session maximum-sessions 20
set protocols ldp interface ge-0/0/0.0
```

```

set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept

```

R5

```

set interfaces ge-0/0/0 unit 0 family inet address 10.50.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.60.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.6.6.6/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.6.6.6
set routing-options forwarding-table export per-packet
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept

```

R6

```

set interfaces ge-0/0/0 unit 0 family inet address 10.60.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.70.1.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.80.1.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.7.7.7/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.7.7.7
set routing-options forwarding-table export per-packet
set protocols ospf topology default backup-spf-options remote-backup-calculation

```

```

set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 link-protection
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept

```

R7

```

set interfaces ge-0/0/0 unit 0 family inet address 10.70.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.8.8.8/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.8.8.8
set routing-options forwarding-table export per-packet
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept

```

R8

```

set interfaces ge-0/0/0 unit 0 family inet address 10.90.1.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.100.1.2/24
set interfaces ge-0/0/1 unit 0 family mpls

```

```

set interfaces lo0 unit 0 family inet address 10.9.9.9/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.9.9.9
set routing-options forwarding-table export per-packet
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface lo0.0
set protocols ospf backup-spf-options remote-backup-calculation
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp auto-targeted-session teardown-delay 20
set protocols ldp auto-targeted-session maximum-sessions 60
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement per-packet then load-balance per-packet
set policy-options policy-statement per-packet then accept

```

Configuring Device R6

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device R6:

1. Configure the interfaces.

```

[edit interfaces]
user@R6# set ge-0/0/0 unit 0 family inet address 10.60.1.2/24
user@R6# set ge-0/0/0 unit 0 family mpls
user@R6# set ge-0/0/1 unit 0 family inet address 10.70.1.1/24
user@R6# set ge-0/0/1 unit 0 family mpls
user@R6# set ge-0/0/2 unit 0 family inet address 10.80.1.2/24
user@R6# set ge-0/0/2 unit 0 family mpls

```

2. Assign the loopback addresses to the device.

```
[edit lo0 unit 0 family]
user@R6# set address 10.7.7.7/32
user@R6# set mpls
```

3. Configure the router ID. Apply the policy to the forwarding table of the local router with the export statement.

```
[edit routing-options]
user@R6# set router-id 10.7.7.7
user@R6# set forwarding-table export per-packet
```

4. Enable remote LFA backup which calculates the backup next hop using dynamic LDP label-switched path.

```
[edit protocols ospf]
user@R6# set topology default backup-spf-options remote-backup-calculation
user@R6# set backup-spf-options remote-backup-calculation
```

5. Configure the traffic engineering and the link protection for the interfaces in the OSPF area.

```
[edit protocols ospf]
user@R6# set traffic-engineering
user@R6# set area 0.0.0.0 interface ge-0/0/0.0 link-protection
user@R6# set area 0.0.0.0 interface ge-0/0/1.0 link-protection
user@R6# set area 0.0.0.0 interface ge-0/0/2.0 link-protection
user@R6# set area 0.0.0.0 interface lo0.0
```

6. Specify a time interval for which the targeted LDP sessions are kept up when the remote LFA goes down, and specify a maximum number of automatically, targeted LDP sessions to optimize the use of memory.

```
[edit protocols ldp]
user@R6# set auto-targeted-session teardown-delay 20
user@R6# set auto-targeted-session maximum-sessions 60
```

7. Configure the LDP protocols on the interfaces.

```
[edit protocols ldp]
user@R6# set interface ge-0/0/0.0
user@R6# set interface ge-0/0/1.0
user@R6# set interface ge-0/0/2.0
user@R6# set interface lo0.0
```

8. Configure the policy options to load balance the per-packet of the policy-statement routing policy.

```
[edit policy-options policy-statement]
user@R6# set per-packet then load-balance per-packet
user@R6# set per-packet then accept
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R6# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.60.1.2/24;
    }
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.70.1.1/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
```

```
        address 10.80.1.2/24;
    }
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 10.7.7.7/32;
        }
        family mpls;
    }
}
```

```
user@R6# show protocols
ospf {
    topology default {
        backup-spf-options {
            remote-backup-calculation;
        }
    }
    backup-spf-options {
        remote-backup-calculation;
        inactive: per-prefix-calculation all;
    }
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/0/0.0 {
            link-protection;
        }
        interface ge-0/0/1.0 {
            link-protection;
        }
        interface ge-0/0/2.0 {
            link-protection;
        }
        interface lo0.0;
    }
}
ldp {
    auto-targeted-session {
```

```
teardown-delay 20;
maximum-sessions 60;
}
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface lo0.0;
}
```

```
user@R6# show policy-options
policy-statement per-packet {
  then {
    load-balance per-packet;
    accept;
  }
}
```

```
user@R6# show routing-options
router-id 10.7.7.7;
forwarding-table {
  export per-packet;
}
```

If you are done configuring the device, enter `commit` from the configuration mode.

Verification

IN THIS SECTION

- [Verifying the Routes | 444](#)
- [Verifying the LDP Routes | 446](#)
- [Verifying the OSPF Routes | 447](#)
- [Verifying the Designated Backup Path Node | 448](#)
- [Verifying the Backup Neighbors | 450](#)

Confirm that the configuration is working properly.

Verifying the Routes

Purpose

Verify that the expected routes are learned.

Action

On Device R6, from operational mode, run the `show route 10.6.6.6/24` command to display the routes in the routing table.

```

user@R6> show route 10.6.6.6/24

inet.0: 75 destinations, 75 routes (75 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.6.6.6/32          *[OSPF/10] 02:21:07, metric 1
                    > to 10.60.1.1 via ge-0/0/0.0
                    to 10.80.1.1 via ge-0/0/2.0, Push 299872

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.6.6.6/32          *[LDP/9] 02:21:07, metric 1
                    > to 10.60.1.1 via ge-0/0/0.0
                    to 10.80.1.1 via ge-0/0/2.0, Push 299792, Push 299872(top)

```

```

inet.0: 75 destinations, 75 routes (75 active, 0 holddown, 0 hidden)
10.6.6.6/32 (1 entry, 1 announced)
  State: <FlashAll>
  *OSPF  Preference: 10
    Next hop type: Router, Next hop index: 1048585
    Address: 0x9df2690
    Next-hop reference count: 10
    Next hop: 10.60.1.1 via ge-0/0/0.0 weight 0x1, selected
    Session Id: 0x141
    Next hop: 10.80.1.1 via ge-0/0/2.0 weight 0x101 uflags Remote neighbor path
    Label operation: Push 299872
    Label TTL action: prop-ttl

```

```

Load balance label: Label 299872: None;
Label element ptr: 0x9dc27a0
Label parent element ptr: 0x0
Label element references: 6
Label element child references: 4
Label element lsp id: 0
Session Id: 0x142
State: <Active Int>
Age: 2:22:40    Metric: 1
Validation State: unverified
Area: 0.0.0.0
Task: OSPF
Announcement bits (2): 0-KRT 4-LDP
AS path: I

```

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

10.6.6.6/32 (1 entry, 1 announced)

State: <FlashAll>

```

*LDP    Preference: 9
        Next hop type: Router, Next hop index: 0
        Address: 0x9df2a90
        Next-hop reference count: 1
        Next hop: 10.60.1.1 via ge-0/0/0.0 weight 0x1, selected
        Label element ptr: 0x9dc0dc0
        Label parent element ptr: 0x0
        Label element references: 1
        Label element child references: 0
        Label element lsp id: 0
        Session Id: 0x0
        Next hop: 10.80.1.1 via ge-0/0/2.0 weight 0x101 uflags Remote neighbor path
        Label operation: Push 299792, Push 299872(top)
        Label TTL action: prop-ttl, prop-ttl(top)
        Load balance label: Label 299792: None; Label 299872: None;
        Label element ptr: 0x9dc1ba0
        Label parent element ptr: 0x9dc27a0
        Label element references: 1
        Label element child references: 0
        Label element lsp id: 0
        Session Id: 0x0
        State: <Active Int>
        Age: 2:22:40    Metric: 1
        Validation State: unverified

```

```
Task: LDP
Announcement bits (1): 0-Resolve tree 1
AS path: I
```

Meaning

The output shows all the routes in the routing table of Device R6.

Verifying the LDP Routes

Purpose

Verify the automatically targeted LDP routes.

Action

From operational mode, enter the show ldp session auto-targeted detail command.

```
user@R6>show ldp session auto-targeted detail

Address: 10.4.4.4, State: Operational, Connection: Open, Hold time: 28
Session ID: 10.7.7.7:0--10.4.4.4:0
Next keepalive in 8 seconds
Active, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: auto-targeted
Keepalive interval: 10, Connect retry interval: 1
Local address: 10.7.7.7, Remote address: 10.4.4.4
Up for 02:28:28
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Session flags: none
Local - Restart: disabled, Helper mode: enabled
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
MTU discovery: disabled
Nonstop routing state: Not in sync
```

Next-hop addresses received:

```
10.4.4.4
10.30.1.2
10.40.1.1
```

Verifying the OSPF Routes

Purpose

Display all the LDP backup routes in the OSPF routing table of Device R6.

Action

On Device R6, from operational mode, run the `show ospf route` command to display the routes in the OSPF routing table.

```
user@R6> show ospf route
```

Topology default Route Table:

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	Nexthop Address/LSP
10.1.1.1	Intra AS	BR	IP	2	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.2.2.2	Intra Router		IP	1	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.4.4.4	Intra Router		IP	3	ge-0/0/0.0	10.60.1.1
					ge-0/0/2.0	10.80.1.1
10.5.5.5	Intra Router		IP	2	ge-0/0/0.0	10.60.1.1
			Bkup LSP			LDP->10.4.4.4
10.6.6.6	Intra Router		IP	1	ge-0/0/0.0	10.60.1.1
			Bkup LSP			LDP->10.4.4.4
10.8.8.8	Intra Router		IP	1	ge-0/0/1.0	10.70.1.2
10.9.9.9	Intra Router		IP	2	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.1.1.1/32	Intra Network		IP	2	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.2.2.2/32	Intra Network		IP	1	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.3.3.3/32	Intra Network		IP	2	ge-0/0/2.0	10.80.1.1
			Bkup LSP			LDP->10.4.4.4
10.4.4.4/32	Intra Network		IP	3	ge-0/0/0.0	10.60.1.1

			ge-0/0/2.0	10.80.1.1
10.5.5.5/32	Intra Network	IP	2 ge-0/0/0.0	10.60.1.1
		Bkup LSP		LDP->10.4.4.4
10.6.6.6/32	Intra Network	IP	1 ge-0/0/0.0	10.60.1.1
		Bkup LSP		LDP->10.4.4.4
10.7.7.7/32	Intra Network	IP	0 lo0.0	
10.8.8.8/32	Intra Network	IP	1 ge-0/0/1.0	10.70.1.2
10.9.9.9/32	Intra Network	IP	2 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4
10.1.1.0/24	Intra Network	IP	2 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4
10.20.1.0/24	Intra Network	IP	2 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4
10.30.1.0/24	Intra Network	IP	3 ge-0/0/2.0	10.80.1.1
		Bkup IP	ge-0/0/0.0	10.60.1.1
10.40.1.0/24	Intra Network	IP	3 ge-0/0/0.0	10.60.1.1
		Bkup IP	ge-0/0/2.0	10.80.1.1
10.50.1.0/24	Intra Network	IP	2 ge-0/0/0.0	10.60.1.1
		Bkup LSP		LDP->10.4.4.4
10.60.1.0/24	Intra Network	IP	1 ge-0/0/0.0	
10.70.1.0/24	Intra Network	IP	1 ge-0/0/1.0	
10.80.1.0/24	Intra Network	IP	1 ge-0/0/2.0	
90.1.1.0/24	Intra Network	IP	3 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4
10.100.1.0/24	Intra Network	IP	2 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4
10.110.1.0/24	Intra Network	IP	3 ge-0/0/2.0	10.80.1.1
		Bkup LSP		LDP->10.4.4.4

Meaning

The output shows all the LDP backup routes in the OSPF routing table of Device R6.

Verifying the Designated Backup Path Node

Purpose

Display the remote LFA next hop determined for a given destination.

Action

From operational mode, enter the `show ospf backup spf results` command.

```

user@R6> show ospf backup spf results

Topology default results:

Area 0.0.0.0 results:

10.6.6.6
  Self to Destination Metric: 1
  Parent Node: 10.60.1.2
  Primary next-hop: ge-0/0/0.0 via 60.1.1.1
  Backup next-hop: LDP->10.4.4.4 via ge-0/0/2.0
  Backup Neighbor: 10.6.6.6 via: Direct
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1, Backup preference: 0x0
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.2.2.2 via: Direct
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1, Backup preference: 0x0
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.8.8.8 via: Direct
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1, Backup preference: 0x0
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.4.4.4 via: LDP (LSP endpoint)
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 3
    Self to Neighbor Metric: 3, Backup preference: 0x0
    Eligible, Reason: Contributes backup next-hop

```

Meaning

The output indicates whether a specific interface or node has been designated as a remote backup path and why.

Verifying the Backup Neighbors

Purpose

Display the backup neighbors for the Device R6

Action

From operational mode, enter the show ospf backup neighbor command.

```
user@R6>show ospf backup neighbor

Topology default backup neighbors:

Area 0.0.0.0 backup neighbors:

10.6.6.6 via: Direct
  Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Direct next-hop: ge-0/0/0.0 via 10.60.1.1

10.8.8.8 via: Direct
  Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Direct next-hop: ge-0/0/1.0 via 10.70.1.2

10.2.2.2 via: Direct
  Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Direct next-hop: ge-0/0/2.0 via 10.80.1.1

10.4.4.4 via: LDP (LSP endpoint)
  Neighbor to Self Metric: 3
  Self to Neighbor Metric: 3
  Direct next-hop: LDP->10.4.4.4 via ge-0/0/2.0
  Direct next-hop: LDP->10.4.4.4 via ge-0/0/0.0
  Neighbors Protected: 2
```

Meaning

The output displays the backup neighbors available for area 0.0.0.0.

SEE ALSO

| [auto-targeted-session](#)

12

CHAPTER

Configure OSPF Support for Traffic Engineering

IN THIS CHAPTER

- [Configuring OSPF Support for Traffic Engineering | 453](#)
 - [How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering | 496](#)
 - [Configuring Application-Specific Link Attribute on an OSPF Interface | 540](#)
 - [How to Enable Link Delay Measurement and Advertising in OSPF | 546](#)
 - [How to Configure Microloop Avoidance in OSPFv2 Segment Routing Networks | 554](#)
-

Configuring OSPF Support for Traffic Engineering

IN THIS SECTION

- [OSPF Support for Traffic Engineering | 453](#)
- [Example: Enabling OSPF Traffic Engineering Support | 455](#)
- [Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface | 464](#)
- [OSPF Passive Traffic Engineering Mode | 467](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode | 467](#)
- [Advertising Label-Switched Paths into OSPFv2 | 471](#)
- [Example: Advertising Label-Switched Paths into OSPFv2 | 471](#)
- [Static Adjacency Segment Identifier for OSPF | 491](#)
- [Overview of Segment Routing Statistics | 494](#)

OSPF Support for Traffic Engineering

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path.

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of OSPF. When traffic engineering is enabled on the routing device, you can enable OSPF traffic engineering support. When you enable traffic engineering for OSPF, the shortest-path-first (SPF) algorithm takes into account the various label-switched paths (LSPs) configured under MPLS and configures OSPF to generate opaque link-state advertisements (LSAs) that carry traffic engineering parameters. The parameters are used to populate the traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. The Constrained Shortest Path First (CSPF) algorithm uses the traffic engineering database to compute the paths that MPLS LSPs take. RSVP uses this path information to set up LSPs and to reserve bandwidth for them.

By default, traffic engineering support is disabled. To enable traffic engineering, include the **traffic-engineering** statement. You can also configure the following OSPF traffic engineering extensions:

- **advertise-unnnumbered-interfaces**—(OSPFv2 only) Advertises the link-local identifier in the link-local traffic engineering LSA packet. You do not need to include this statement if RSVP is able to signal unnumbered interfaces as defined in RFC 3477, *Signalling Unnumbered Links in Resource Reservation Protocol - Traffic Engineering (RSVP-TE)*.
- **credibility-protocol-preference**—(OSPFv2 only) Assigns a credibility value to OSPF routes in the traffic engineering database. By default, Junos OS prefers OSPF routes in the traffic engineering database over other interior gateway protocol (IGP) routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure OSPF to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, OSPF routes are not automatically preferred by the traffic engineering database, depending on your configuration.
- **ignore-lsp-metrics**—Ignores RSVP LSP metrics in OSPF traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs. This option avoids mutual dependency between OSPF and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date. In addition, if you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.
- **multicast-rpf-routes**—(OSPFv2 only) Installs unicast IPv4 routes (not LSPs) in the multicast routing table (**inet.2**) for multicast reverse-path forwarding (RPF) checks. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check if the packet is coming in on an interface that is also sending data back to the packet source.
- **no-topology**—(OSPFv2 only) To disable the dissemination of link-state topology information. If disabled, traffic engineering topology information is no longer distributed within the OSPF area.
- **shortcuts**—Configures IGP shortcuts, which allows OSPF to use an LSP as the next hop as if it were a *logical interface* from the ingress routing device to the egress routing device. The address specified in the **to** statement at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level on the ingress routing device must match the router ID of the egress routing device for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. When used in this way, LSPs are no different from Asynchronous Transfer Mode (ATM) and Frame Relay virtual circuits (VCs), except that LSPs carry only IPv4 traffic.

OSPFv2 installs the prefix for IPv4 routes in the **inet.0** routing table, and the LSPs are installed by default in the **inet.3** routing table.

OSPFv3 LSPs used for shortcuts continue to be signaled using IPv4. However, by default, shortcut IPv6 routes calculated through OSPFv3 are added to the **inet6.3** routing table. The default behavior is for BGP only to use LSPs in its calculations. If you configure MPLS so that both BGP and IGP use

LSPs for forwarding traffic, IPv6 shortcut routes calculated through OSPFv3 are added to the **inet6.0** routing table.



NOTE: Whenever possible, use OSPF IGP shortcuts instead of traffic engineering shortcuts.

- **lsp-metric-info-summary**—Advertises the LSP metric in summary LSAs to treat the LSP as a link. This configuration allows other routing devices in the network to use this LSP. To accomplish this, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

When you enable traffic engineering on the routing device, you can also configure an OSPF metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the traffic engineering database. Its value does not affect normal OSPF forwarding.

Example: Enabling OSPF Traffic Engineering Support

IN THIS SECTION

- [Requirements | 455](#)
- [Overview | 456](#)
- [Configuration | 456](#)
- [Verification | 462](#)

This example shows how to enable OSPF traffic engineering support to advertise the label-switched path (LSP) metric in summary link-state advertisements (LSAs).

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure BGP per your network requirements. See the [BGP User Guide](#)
- Configure MPLS per your network requirements. See the [MPLS Applications User Guide](#).

Overview

You can configure OSPF to treat an LSP as a link and have other routing devices in the network use this LSP. To accomplish this, you configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

In this example, there are four routing devices in area 0.0.0.0, and you want OSPF to treat the LSP named R1-to-R4 that goes from the ingress Device R1 to the egress Device R4 as a link.

For OSPF, you enable traffic engineering on all four routing devices in the area by including the `traffic-engineering` statement. This configuration ensures that the shortest-path-first (SPF) algorithm takes into account the LSPs configured under MPLS and configures OSPF to generate LSAs that carry traffic engineering parameters. You further ensure that OSPF uses the MPLS LSP as the next hop and advertises the LSP metric in summary LSAs, by including the optional shortcuts `lsp-metric-into-summary` statement on the ingress Device R1.

For MPLS, you enable traffic engineering so that MPLS performs traffic engineering on both BGP and IGP destinations by including the `traffic-engineering bgp-igp` statement, and you include the LSP named R1-to-R4 by including the `label-switched-path lsp-path-name to address` statement on the ingress Device R1. The address specified in the `to` statement on the ingress Device R1 must match the router ID of the egress Device R4 for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. In this example, the router ID of the egress Device R4 is 10.0.0.4.

Configuration

IN THIS SECTION

- [Procedure | 456](#)

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

Procedure

CLI Quick Configuration

To quickly enable OSPF traffic engineering support to advertise the LSP metric in summary LSAs, copy the following commands and paste them into the CLI.

Configuration on R1:

```
[edit]
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

Configuration on R2:

```
[edit]
set routing-options router-id 10.0.0.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R3:

```
[edit]
set routing-options router-id 10.0.0.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R4:

```
[edit]
set routing-options router-id 10.0.0.4
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Step-by-Step Procedure

To enable OSPF traffic engineering support to advertise LSP metrics in summary LSAs:

1. Configure the router ID.

```
[edit]  
user@R1# set routing-options router-id 10.0.0.1
```

```
[edit]  
user@R2# set routing-options router-id 10.0.0.2
```

```
[edit]  
user@R3# set routing-options router-id 10.0.0.3
```

```
[edit]  
user@R4# set routing-options router-id 10.0.0.4
```

2. Configure the OSPF area and add the interfaces.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface all
user@R1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface all
user@R2# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface all
user@R3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface all
user@R4# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

3. Enable OSPF traffic engineering.

```
[edit]
user@R1# set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
```

```
[edit]
user@R2# set protocols ospf traffic-engineering
```

```
[edit]
user@R3# set protocols ospf traffic-engineering
```

```
[edit]
user@R4# set protocols ospf traffic-engineering
```

4. On Device R1, configure MPLS traffic engineering.

```
[edit ]
user@R1# set protocols mpls traffic-engineering bgp-igp
user@R1# set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options`, `show protocols ospf`, and `show protocols mpls` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@host# show routing-options
router-id 10.0.0.1;
```

```
user@host# show protocols ospf
  traffic-engineering {
    shortcuts lsp-metric-into-summary;
  }
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

```
user@host# show protocols mpls
  traffic-engineering bgp-igp;
  label-switched-path R1-to-R4 {
    to 10.0.0.4;
  }
}
```

Output for R2:

```
user@host# show routing-options
router-id 10.0.0.2;
```

```
user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

Output for R3:

```
user@host# show routing-options
router-id 10.0.0.3;
```

```
user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

Output for R4:

```
user@host# show routing-options
router-id 10.0.0.4;
```

```
user@host# show protocols ospf
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show routing-options`, `show protocols ospf3`, and `show protocols mpls` commands.

Verification

IN THIS SECTION

- [Verifying the Traffic Engineering Capability for OSPF | 463](#)

- [Verifying OSPF Entries in the Traffic Engineering Database | 463](#)
- [Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF | 463](#)

Confirm that the configuration is working properly.

Verifying the Traffic Engineering Capability for OSPF

Purpose

Verify that traffic engineering has been enabled for OSPF. By default, traffic engineering is disabled.

Action

From operational mode, enter the `show ospf overview` command for OSPFv2, and enter the `show ospf3 overview` for OSPFv3.

Verifying OSPF Entries in the Traffic Engineering Database

Purpose

Verify the OSPF information in the traffic engineering database. The Protocol field displays OSPF and the area from which the information was learned.

Action

From operational mode, enter the `show ted database` command.

Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF

Purpose

Verify that OSPF is reporting node information. The Protocol name field displays OSPF and the area from which the information was learned.

Action

From operational mode, enter the `show ted protocol` command.

Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface

IN THIS SECTION

- Requirements | 464
- Overview | 464
- Configuration | 464
- Verification | 466

This example shows how to configure the OSPF metric value used for traffic engineering.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure OSPF for traffic engineering. See "[Example: Enabling OSPF Traffic Engineering Support](#)" on [page 455](#)

Overview

You can configure an OSPF metric that is used exclusively for traffic engineering. To modify the default value of the traffic engineering metric, include the `te-metric` statement. The OSPF traffic engineering metric does not affect normal OSPF forwarding. By default, the traffic engineering metric is the same value as the OSPF metric. The range is 1 through 65,535.

In this example, you configure the OSPF traffic engineering metric on OSPF interface **fe-0/1/1** in area 0.0.0.0.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 465
- Procedure | 465

CLI Quick Configuration

To quickly configure the OSPF traffic engineering metric for a specific interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/1/1 te-metric 10
```

Procedure

Step-by-Step Procedure

To configure an OSPF traffic engineering metric for a specific interface used only for traffic engineering:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the traffic engineering metric of the OSPF network segments.

```
[edit protocols ospf area 0.0.0.0]
user@host set interface fe-0/1/1 te-metric 10
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/1/1.0 {
    te-metric 10;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Configured Traffic Engineering Metric | 466](#)

Confirm that the configuration is working properly.

Verifying the Configured Traffic Engineering Metric

Purpose

Verify the traffic engineering metric value. Confirm that Metric field displays the configured traffic engineering metric.

Action

From operational mode, enter the `show ted database extensive` command.

OSPF Passive Traffic Engineering Mode

Ordinarily, interior routing protocols such as OSPF are not run on links between autonomous systems. However, for inter-AS traffic engineering to function properly, information about the inter-AS link—in particular, the address on the remote interface—must be made available inside the autonomous system (AS). This information is not normally included either in the external BGP (EBGP) reachability messages or in the OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include it in the traffic engineering database. OSPF traffic engineering mode allows MPLS label-switched paths (LSPs) to dynamically discover OSPF AS boundary routers and to allow routers to establish a traffic engineering LSP across multiple autonomous systems.

Example: Configuring OSPF Passive Traffic Engineering Mode

IN THIS SECTION

- [Requirements | 467](#)
- [Overview | 468](#)
- [Configuration | 468](#)
- [Verification | 470](#)

This example shows how to configure OSPF passive mode for traffic engineering on an inter-AS interface. The AS boundary router link between the EBGP peers must be a directly connected link and must be configured as a passive traffic engineering link.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure BGP per your network requirements. See the [BGP User Guide](#).
- Configure the LSP per your network requirements. See the [MPLS Applications User Guide](#).

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)
- Configure a single-area OSPF network. See ["Example: Configuring a Single-Area OSPF Network" on page 62](#).
- Configure a multiarea OSPF network. See ["Example: Configuring a Multiarea OSPF Network" on page 65](#).

Overview

You can configure OSPF passive mode for traffic engineering on an inter-AS interface. The address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGP link. In this example, you configure interface **so-1/1/0** in area 0.0.0.1 as the inter-AS link to distribute traffic engineering information with OSPF within the AS and include the following settings:

- **passive**—Advertises the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.
- **traffic-engineering**—Configures an interface in OSPF passive traffic-engineering mode to enable dynamic discovery of OSPF AS boundary routers. By default, OSPF passive traffic-engineering mode is disabled.
- **remote-node-id**—Specifies the IP address at the far end of the inter-AS link. In this example, the remote IP address is 192.168.207.2.

Configuration

IN THIS SECTION

- [Procedure | 469](#)

To quickly configure OSPF passive mode for traffic engineering, copy the following command, remove any line breaks, and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

Procedure

Step-by-Step Procedure

To configure OSPF passive traffic engineering mode:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1
```

2. Configure interface **so-1/1/0** as a passive interface configured for traffic engineering, and specify the IP address at the far end of the inter-AS link.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-1/1/0 passive traffic-engineering remote-node-id 192.168.207.2
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface so-1/1/0.0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

Verification

IN THIS SECTION

- [Verifying the Status of OSPF Interfaces | 470](#)

Confirm that the configuration is working properly.

Verifying the Status of OSPF Interfaces

Purpose

Verify the status of OSPF interfaces. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

Action

From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

Advertising Label-Switched Paths into OSPFv2

One main reason to configure label-switched paths (LSPs) in your network is to control the shortest path between two points on the network. You can advertise LSPs into OSPFv2 as point-to-point links so that all participating routing devices can take the LSP into account when performing SPF calculations. The advertisement contains a local address (the **from** address of the LSP), a remote address (the **to** address of the LSP), and a metric with the following precedence:

1. Use the LSP metric defined under OSPFv2.
2. Use the LSP metric configured for the label-switched path under MPLS.
3. If you do not configure any of the above, use the default OSPFv2 metric of 1.



NOTE: If you want an LSP that is announced into OSPFv2 to be used in SPF calculations, there must be a reverse link (that is, a link from the tail end of the LSP to the head end). You can accomplish this by configuring an LSP in the reverse direction and also announcing it in OSPFv2.

Example: Advertising Label-Switched Paths into OSPFv2

IN THIS SECTION

- Requirements | 471
- Overview | 472
- Configuration | 474
- Verification | 490

This example shows how to advertise LSPs into OSPFv2.

Requirements

Before you begin, configure the device interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).

Overview

IN THIS SECTION

- [Topology | 473](#)

To advertise an LSP into OSPFv2, you define the LSP and configure OSPFv2 to route traffic using the LSP. By doing this, you can use the LSP to control the shortest path between two points on the network. You might choose to do this if you want to have OSPF traffic routed along the LSP instead of having OSPF use the default best-effort routing.

In this example, you configure the following to advertise an LSP into OSPFv2:

- BGP

For all routing devices, configure the local AS number 65000 and define the IBGP group that recognizes the specified BGP systems as peers. All members are internal to the local AS, so you configure an internal group with a full list of peers. You also include the peer AS group, which is the same as the local AS number that you configure.

- MPLS

For all routing devices, configure the protocol family on each transit logical interface and enable MPLS on all interfaces, except for the management interface (**fxp0.0**). Specify the **mpls** protocol family type.

- RSVP

For all routing devices, enable RSVP on all interfaces, except for the management interface (**fxp0.0**). You enable RSVP on the devices in this network to ensure that the interfaces can signal the LSP.

- OSPFv2

For all routing devices, use the loopback address to assign the router ID, administratively group all of the devices into OSPF area 0.0.0.0, add all of the interfaces participating in OSPF to area 0.0.0.0, and disable OSPF on the management interface (**fxp0.0**).

- Label-switched path

On the ingress routing device R1, which is the beginning (or head end) of the LSP, configure an LSP with an explicit path. The explicit path indicates that the LSP must go to the next specified IP address in the path without traversing other nodes. In this example, you create an LSP named R1-to-R6, and you specify the IP address of the egress routing device R6.

- Advertise the LSP in OSPFv2


On the ingress routing device R1, you advertise the LSP as a point-to-point link into OSPFv2. You can optionally assign a metric to have the LSP be the more or less preferred path to the destination.

Topology

[Figure 27 on page 474](#) shows a sample network topology that consists of the following:

- BGP is configured on all routing devices, with one local autonomous system (AS) 65000 that contains three routing devices:
 - R1—Device R1 is the ingress device with a router ID of 10.0.0.1. Interface **so-0/0/2** connects to Device R3.
 - R3—Device R3 is the transit device with a router ID of 10.0.0.3. Interface **so-0/0/2** connects to Device R1, and interface **so-0/0/3** connects to Device R6.
 - R6—Device R6 is the egress device with a router ID of 10.0.0.6. Interface **so-0/0/3** connects to Device R3.
- OSPFv2 is configured on all routing devices.
- MPLS and RSVP are enabled on all routing devices.
- One RSVP-signaled LSP is configured on Device R1.

Figure 27: Advertising an LSP into OSPFv2



Configuration

IN THIS SECTION

- [Configuring BGP | 475](#)
- [Configuring MPLS | 478](#)
- [Configuring RSVP | 482](#)
- [Configuring OSPF | 484](#)
- [Configuring the LSP | 488](#)
- [Advertising the LSP into OSPFv2 | 489](#)

The following examples require you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in [CLI User Guide](#).

To configure the devices to advertise an LSP into OSPFv2, perform the following tasks:

Configuring BGP

CLI Quick Configuration

To quickly configure BGP on each routing device, copy the following commands and paste them into the CLI.

Configuration on Device R1:

```
[edit]
set routing-options autonomous-system 65000
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 10.0.0.1
set protocols bgp group internal-peers neighbor 10.0.0.3
set protocols bgp group internal-peers neighbor 10.0.0.6
set protocols bgp group internal-peers peer-as 65000
```

Configuration on Device R3:

```
[edit]
set routing-options autonomous-system 65000
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 10.0.0.3
set protocols bgp group internal-peers neighbor 10.0.0.1
set protocols bgp group internal-peers neighbor 10.0.0.6
set protocols bgp group internal-peers peer-as 65000
```

Configuration on Device R6:

```
[edit]
set routing-options autonomous-system 65000
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 10.0.0.6
set protocols bgp group internal-peers neighbor 10.0.0.1
set protocols bgp group internal-peers neighbor 10.0.0.3
set protocols bgp group internal-peers peer-as 65000
```

Step-by-Step Procedure

To configure BGP:

1. On each routing device, configure the local AS number.

```
[edit]
user@R1# set routing-options autonomous-system 65000
```

```
[edit]
user@R3# set routing-options autonomous-system 65000
```

```
[edit]
user@R6# set routing-options autonomous-system 65000
```

2. On each routing device, configure the internal BGP neighbor connections.

```
[edit]
user@R1# set protocols bgp group internal-peers type internal
user@R1# set protocols bgp group internal-peers local-address 10.0.0.1
user@R1# set protocols bgp group internal-peers neighbor 10.0.0.3
user@R1# set protocols bgp group internal-peers neighbor 10.0.0.6
user@R1# set protocols bgp group internal-peers peer-as 65000
```

```
[edit]
user@R3# set protocols bgp group internal-peers type internal
user@R3# set protocols bgp group internal-peers local-address 10.0.0.3
user@R3# set protocols bgp group internal-peers neighbor 10.0.0.1
user@R3# set protocols bgp group internal-peers neighbor 10.0.0.6
user@R3# set protocols bgp group internal-peers peer-as 65000
```

```
[edit]
user@R6# set protocols bgp group internal-peers type internal
user@R6# set protocols bgp group internal-peers local-address 10.0.0.6
user@R6# set protocols bgp group internal-peers neighbor 10.0.0.1
user@R6# set protocols bgp group internal-peers neighbor 10.0.0.3
user@R6# set protocols bgp group internal-peers peer-as 65000
```

3. If you are done configuring the devices, commit the configuration.

```
[edit]  
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options` and `show protocols bgp` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on R1:

```
user@R1# show routing-options  
autonomous-system 65000;
```

```
user@R1# show protocols bgp  
group internal-peers {  
    type internal;  
    local-address 10.0.0.1;  
    peer-as 65000;  
    neighbor 10.0.0.3;  
    neighbor 10.0.0.6;  
}
```

Configuration on R3:

```
user@R3# show routing-options  
autonomous-system 65000;
```

```
user@R3# show protocols bgp  
group internal-peers {  
    type internal;  
    local-address 10.0.0.3;  
    peer-as 65000;  
    neighbor 10.0.0.1;
```

```
neighbor 10.0.0.6;
}
```

Configuration on R6:

```
user@R6# show routing-options
autonomous-system 65000;
```

```
user@R6# show protocols bgp
group internal-peers {
  type internal;
  local-address 10.0.0.6;
  peer-as 65000;
  neighbor 10.0.0.1;
  neighbor 10.0.0.3;
}
```

Configuring MPLS

CLI Quick Configuration

To quickly configure MPLS on all of the routing devices in AS 65000, copy the following commands and paste them into the CLI.

Configuration on Device R1:

```
[edit]
set interfaces so-0/0/2 unit 0 family mpls
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```

Configuration on Device R3:

```
[edit]
set interfaces so-0/0/2 unit 0 family mpls
set interfaces so-0/0/3 unit 0 family mpls
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```

Configuration on Device R6:

```
[edit]
set interfaces so-0/0/3 unit 0 family mpls
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```

Step-by-Step Procedure

To configure MPLS:

1. Configure the transit interfaces for MPLS.

```
[edit ]
user@R1# set interfaces so-0/0/2 unit 0 family mpls
```

```
[edit ]
user@R3# set interfaces so-0/0/2 unit 0 family mpls
user@R3# set interfaces so-0/0/3 unit 0 family mpls
```

```
[edit ]
user@R6# set interfaces so-0/0/3 unit 0 family mpls
```

2. Enable MPLS.

```
[edit ]
user@R1# set protocols mpls interface all
```

```
[edit ]
user@R3# set protocols mpls interface all
```

```
[edit ]
user@R6# set protocols mpls interface all
```

3. Disable MPLS on the management interface (fxp0.0).

```
[edit ]
user@R1# set protocols mpls interface fxp0.0 disable
```

```
[edit ]
user@R3# set protocols mpls interface fxp0.0 disable
```

```
[edit ]
user@R6# set protocols mpls interface fxp0.0 disable
```

4. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces` and `show protocols mpls` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on Device R1:

```
user@R1# show interfaces
so-0/0/2 {
  unit 0 {
    family mpls;
  }
}
```

```
user@R1# show protocols mpls
interface all;
interface fxp0.0 {
```

```
    disable;  
}
```

Configuration on Device R3:

```
user@R3# show interfaces  
so-0/0/2 {  
  unit 0 {  
    family mpls;  
  }  
}  
so-0/0/3 {  
  unit 0 {  
    family mpls;  
  }  
}
```

```
user@R3# show protocols mpls  
interface all;  
interface fxp0.0 {  
  disable;  
}
```

Configuration on Device R6:

```
user@R6# show interfaces  
so-0/0/3 {  
  unit 0 {  
    family mpls;  
  }  
}
```

```
user@R6# show protocols mpls  
interface all;  
interface fxp0.0 {  
  disable;  
}
```

Configuring RSVP

CLI Quick Configuration

To quickly configure RSVP on all of the routing devices in AS 65000, copy the following commands and paste them into the CLI.

Configuration on Device R1:

```
[edit]
set protocols rsvp interface so-0/0/2
set protocols rsvp interface fxp0.0 disable
```

Configuration on Device R3:

```
[edit]
set protocols rsvp interface so-0/0/2
set protocols rsvp interface so-0/0/3
set protocols rsvp interface fxp0.0 disable
```

Configuration on Device R6:

```
[edit]
set protocols rsvp interface so-0/0/3
set protocols rsvp interface fxp0.0 disable
```

Step-by-Step Procedure

To configure RSVP:

1. Enable RSVP.

```
[edit ]
user@R1# set protocols rsvp interface so-0/0/2
```

```
[edit ]
user@R3# set protocols rsvp interface so-0/0/2
user@R3# set protocols rsvp interface so-0/0/3
```

```
[edit ]
user@R6# set protocols rsvp interface so-0/0/3
```

2. Disable RSVP on the management interface (fxp0.0).

```
[edit ]
user@R1# set protocols rsvp interface fxp0.0 disable
```

```
[edit ]
user@R3# set protocols rsvp interface fxp0.0 disable
```

```
[edit ]
user@R6# set protocols rsvp interface fxp0.0 disable
```

3. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols rsvp` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on Device R1:

```
user@R1# show protocols rsvp
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
```

Configuration on Device R3:

```
user@R3# show protocols rsvp
interface so-0/0/2.0;
interface so-0/0/3.0;
interface fxp0.0 {
    disable;
}
```

Configuration on Device R6:

```
user@R3# show protocols rsvp
interface so-0/0/3.0;
interface fxp0.0 {
    disable;
}
```

Configuring OSPF

CLI Quick Configuration

To quickly configure OSPF, copy the following commands and paste them into the CLI.

Configuration on Device R1:

```
[edit]
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

Configuration on Device R3:

```
[edit]
set routing-options router-id 10.0.0.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

Configuration on Device R6:

```
[edit]
set routing-options router-id 10.0.0.6
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

Step-by-Step Procedure

To configure OSPF:

1. Configure the router ID.

```
[edit]
user@R1# set routing-options router-id 10.0.0.1
```

```
[edit]
user@R3# set routing-options router-id 10.0.0.3
```

```
[edit]
user@R6# set routing-options router-id 10.0.0.6
```

2. Configure the OSPF area and the interfaces.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface all
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface all
```

```
[edit]
user@R6# set protocols ospf area 0.0.0.0 interface all
```

3. Disable OSPF on the management interface (fxp0.0).

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R6# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

4. If you are done configuring the devices, commit the configuration.

```
[edit ]
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on Device R1:

```
user@R1# show routing-options
router-id 10.0.0.1;
```

```
user@R1# show protocols ospf
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

Configuration on Device R3:

```
user@R3# show routing-options
router-id 10.0.0.3;
```

```
user@R3# show protocols ospf
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

Configuration on Device R6:

```
user@R6# show routing-options
router-id 10.0.0.6;
```

```
user@R6# show protocols ospf
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

```
}  
}
```

Configuring the LSP

CLI Quick Configuration

To quickly configure the LSP on the ingress routing device Router R1, copy the following command and paste it into the CLI.

```
[edit]  
set protocols mpls label-switched-path R1-to-R6 to 10.0.0.6
```

Step-by-Step Procedure

To configure the LSP on Device R1:

1. Enter MPLS configuration mode.

```
[edit]  
user@R1# edit protocols mpls
```

2. Create the LSP.

```
[edit protocols mpls]  
user@R1# set label-switched-path R1-to-R6 to 10.0.0.6
```

3. If you are done configuring the device, commit the configuration.

```
[edit ]  
user@R1# commit
```

Results

Confirm your configuration by entering the `show protocols mpls` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols mpls
label-switched-path R1-to-R6 {
  to 10.0.0.6;
}
```

Advertising the LSP into OSPFv2

CLI Quick Configuration

To quickly advertise the LSP into OSPFv2 and optionally include a metric for the LSP on Device R1, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 label-switched-path R1-to-R6
set protocols ospf area 0.0.0.0 label-switched-path R1-to-R6 metric 2
```

Step-by-Step Procedure

To advertise the LSP into OSPFv2 on Router R1:

1. Enter OSPF configuration mode.

```
[edit]
user@R1# edit protocols ospf
```

2. Include the `label-switched-path` statement, and specify the LSP R1-to-R6 that you created.

```
[edit protocols ospf]
user@R1# set area 0.0.0.0 label-switched-path R1-to-R6
```

3. (Optional) Specify a metric for the LSP.

```
[edit protocols ospf]
user@R1# set protocols ospf area 0.0.0.0 label-switched-path R1-to-R6 metric 2
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@R1# commit
```

Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols ospf
area 0.0.0.0 {
  label-switched-path R1-to-R6 {
    metric 2;
  }
}
```

Verification

IN THIS SECTION

- [Verifying the OSPF Neighbor | 490](#)

Confirm that the configuration is working properly.

Verifying the OSPF Neighbor

Purpose

Verify that another neighbor is listed and is reachable over the LSP. The interface field indicates the name of the LSP.

Action

From operational mode, enter the `show ospf neighbor` command.

Static Adjacency Segment Identifier for OSPF

Adjacency segment is a strict forwarded single-hop tunnel that carries packets over a specific link between two nodes, irrespective of the link cost. You can configure static adjacency segment identifier (SID) labels for an interface.

Configuring a static adjacency SID on an interface causes the existing dynamically allocated adjacency SID to be removed along with the transit route for the same.

For static adjacency SIDs, the labels are picked from either a static reserved label pool or from an OSPF segment routing global block (SRGB).

You can reserve a label range to be used for static allocation of labels using the following configuration:

```
user@host# set protocols mpls label-range static-label-range start-value end-value
```

The static pool can be used by any protocol to allocate a label in this range. You need to ensure that no two protocols use the same static label. OSPF adjacency SIDs can be allocated from this label block through the configuration using keyword `label`. The label value for the specific adjacency SIDs need to be explicitly configured. The following is a sample configuration:

```
user@host# set protocols mpls label-range static-label-range 700000 799999;
user@host# set protocols ospf source-packet-routing srgb start-label 800000 index-range 4000;
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0.1 ipv4-adjacency-segment unprotected label
700001;
```



NOTE: When you use `ipv4-adjacency-segment` command, the underlying interface must be point-to-point.

SRGB is a global label space that is allocated for the protocol based on configuration. The labels in the entire SRGB is available for OSPF to use and are not allocated to other applications/protocols. Prefix SIDs (and Node SIDs) are indexed from this SRGB.

OSPF Adj-SIDs can be allocated from OSPF SRGB using keyword 'index' in the configuration. In such cases, it should be ensured that the Adj-SID index does not conflict with any other prefix SID in the domain. Like Prefix-SIDs, Adj-SIDs will also be configured by mentioning the index with respect to the

SRGB. However, the Adj-SID subtlv will still have the SID as a value and the L and V flags are set. The following is a sample configuration:

```
user@host# set protocols ospf source-packet-routing srgb start-label 800000 index-range 4000;
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0.1 ipv4-adjacency-segment unprotected index 1;
```

Static adjacency SIDs can be configured per area and also based on whether the protection is required or not. Adjacency SIDs should be configured per interface at the [edit protocols ospf area *area* interface *interface-name*] hierarchy level.

- Protected—Ensures adjacency SID is eligible to have a backup path and a B-flag is set in an adjacency SID advertisement.
- Unprotected—Ensures no backup path is calculated for a specific adjacency SID and a B-flag is not set in an adjacency SID advertisement.

The following is a sample configuration:

```
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0.1 ipv4-adjacency-segment unprotected index 1;
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/1.1 ipv4-adjacency-segment protected index 2;
```

When segment routing is used in LAN subnetworks, each router in the LAN may advertise the adjacency SID of each of its neighbors. To configure adjacency SID for a LAN interface to a specific neighbor, you should configure the adjacency SIDs under the lan-neighbor configuration at the [edit protocols ospf area 0.0.0.0 interface *interface_name* lan-neighbor *neighbor-routerid*] hierarchy level. The following is a sample configuration:

```
user@host# set protocols mpls label-range static-label-range 700000 799999;
user@host# set protocols ospf source-packet-routing srgb start-label 800000 index-range 4000;
user@host# set protocols ospf area 0.0.0.0 interface ge-1/0/0.1 lan-neighbor 11.12.1.2 ipv4-adjacency-segment unprotected label 700001;
```

Use the following CLI hierarchy for configuring adjacency SID:

```
[edit ]
protocols {
  ospf {
    area 0.0.0.0 {
      interface <interface_name> {
        ipv4-adjacency-segment {
          protected {
```



```

Label      Flags      Adj-Sid-Type
90010      BVLP       Protected
1212      VLP        UnProtected
regress@10.49.129.231# run show route label 90010

mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

90010      *[L-OSPF/10/5] 00:00:21, metric 0
          > to 11.12.1.2 via ge-1/0/0.0, Pop
          to 11.12.2.2 via ge-1/0/2.0, Swap 16021
          to 11.12.3.2 via ge-1/0/3.0, Swap 16021

```

Overview of Segment Routing Statistics

Traffic statistics in a segment routing network can be recorded in an OpenConfig compliant format for the Layer 3 interfaces. The statistics is recorded for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per segment identifier (SID). To enable recording of segment routing statistics, include `sensor-based-stats` statement at the `[edit protocol isis source-packet-routing]` hierarchy level.

Earlier, sensors were available for collecting segment routing statistics for MPLS transit traffic only, which is MPLS-to-MPLS in nature. On MX Series routers with MPC and MIC interfaces and PTX Series routers, additional sensors are introduced to collect segment routing statistics for MPLS ingress traffic, which is IP-to-MPLS in nature. With this feature, you can enable sensors for label IS-IS segment routing traffic only, and stream the statistics to a gRPC client.

You can enable the segment routing statistics for MPLS ingress traffic using the `egress` option under the `per-sid` configuration statement. The resource name for the `per-sid egress` functionality is:

```
/junos/services/segment-routing/sid/egress/usage/
```

You can view the label IS-IS route association with the sensors using the `show isis spring sensor info` command output. This command does not display counter values of the actual sensors.

The segment routing statistics records are exported to a server. You can view segment routing statistics data from the following the OpenConfig paths:

- `/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='L-ISIS-10.1.1.1']/state/counters[name='oc-xxx']/out-pkts`
- `/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='L-ISIS-10.1.1.1']/state/counters[name='oc-xxx']/out-pkts`



NOTE:

- Graceful Routing Engine switchover (GRES) is not supported for segment routing statistics.

Nonstop active routing (NSR) is not supported for label IS-IS. During a Routing Engine switchover, a new sensor is created in the new primary Routing Engine, replacing the sensor created by the previous primary Routing Engine. As a result, at the time of a Routing Engine switchover, the segment routing statistics counter start from zero.

- Graceful restart is not supported for label IS-IS.

In case of graceful restart, the existing sensor is deleted and a new sensor is created during IS-IS initialization. The segment routing statistics counter restarts from zero.

- In-service software upgrade (ISSU) and nonstop software upgrade (NSSU) are not supported. In such cases, the segment routing statistics counter is restarted.
- Zero-statistics segment routing data is suppressed and does not get streamed to the gRPC clients.

SEE ALSO

[sensor-based-stats](#)

[sensor \(Junos Telemetry Interface\)](#)

[sensor-based-stats \(Junos Telemetry Interface\)](#)

RELATED DOCUMENTATION

[MPLS Applications User Guide](#)

How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering

SUMMARY

A flexible algorithm allows IGP's alone to compute constraint based paths over the network thereby providing simple traffic engineering without using a network controller. This is a light weight solution for networks that have not implemented a controller with full fledged segment routing but still want to reap the benefits of segment routing in their network.

IN THIS SECTION

- [Understanding OSPF Flexible Algorithm for Segment Routing | 496](#)
- [Example: OSPF Flexible Algorithm | 505](#)
- [| 539](#)
- [| 539](#)
- [| 539](#)

WHAT'S NEXT

For more information on configuring flexible algorithms, see the [OSPF User Guide](#)

Understanding OSPF Flexible Algorithm for Segment Routing

IN THIS SECTION

- [Benefits of Configuring Flexible Algorithm | 497](#)
- [What is Flexible Algorithm Definition \(FAD\)? | 497](#)
- [Participation in a Flexible Algorithm | 499](#)
- [Network Topology Configured with Flexible Algorithm Definitions | 499](#)
- [Flexible Algorithm RIBs | 504](#)
- [BGP Community and Flexible Algorithms | 504](#)
- [Supported and Unsupported Features | 505](#)

You can thin slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering using segment routing without actually implementing a network controller. You can use the prefix SIDs to steer packets along the constraint-based paths. You can configure the prefix SIDs for flexible algorithm through policy configurations.

IGP protocols use a link metric to calculate a best path. However, the best IGP path might not always be the best path for certain types of traffic. Therefore, the IGP computed best path based on the shortest IGP metric is often replaced with traffic engineered path due to the traffic requirements that are not reflected by the IGP metric. Typically RSVP-TE or SR TE is used for computing the path based on additional metrics and constraints to overcome this limitation. Junos installs such paths in the forwarding tables in addition to or as a replacement for the original path computed by the IGP.

Benefits of Configuring Flexible Algorithm

- A lightweight version of segment routing traffic engineering that can be used in the core of the network.
- Allows you to configure traffic engineering using segment routing even without installing a network controller.
- Utilize equal-cost multipath (ECMP) and TI-LFA per-slice without configuring BGP-LS or static path.
- Compute TI-LFA backup path using the same flexible algorithm definition and constraints computation.
- Take advantage of segment routing traffic engineering using only OSPFv2 without configuring RSVP or LDP.
- Ability to provision constrained primary path based on a single label.

What is Flexible Algorithm Definition (FAD)?

A flexible algorithm allows IGP to calculate additional best paths based on specified constraints thereby providing simple traffic engineering without using a network controller. This is a lightweight solution for networks that have not implemented a controller with full fledged segment routing but still want to reap the benefits of segment routing in their network. Every operator can define separate constraints or colors depending on their requirements.

To define a flexible algorithm, include `flex-algorithm id` statement at the `[edit routing-options]` hierarchy level. The flexible algorithm definition (FAD) is assigned with an identifier ranging from 128 through 255. This flexible algorithm can be defined on one or more routers in a network. A flexible algorithm computes a best path based on the following parameters:

- **Calculation type**—SPF or strict SPF are the two available calculation type options. You can specify one of these calculation types in your FAD. Select the SPF calculation type if you want to influence the SPF computation on your device based on a certain local policy such as traffic engineering shortcuts. If you select strict SPF then the local policy cannot influence the SPF path selection.
- **Metric type**- IGP metric or TE metric are the available metric type options. You can specify one of these metric types in your FAD depending on your network requirement. If you do not want to use the IGP metric for a specific link you can configure a TE metric that OSPFv2 can use for calculating the route.
- **Priority**- You can assign a priority to your FADs as per your requirement and OSPFv2 prioritizes a particular FAD advertisement over another FAD based on your assigned priority.
- **Set of Link constraints**- You can configure admin-groups for many protocols at the [edit protocols mpls admin-groups] hierarchy level to color an individual link. These admin-groups can then be defined as include any, include-all or exclude at the [edit routing-options flex-algorithm *definition* admin-groups] hierarchy level.

We recommend configuring flexible algorithm on only a few routers to provide redundancy and to avoid conflicts. Flexible algorithm definition is advertised in IGP as FAD sub-TLVs. In very large networks, we do not recommend configuring more than 8 flexible algorithm definitions as each flexible algorithm will compute its own path and might cause performance issues beyond that.

The default FAD has the following parameters:

- calculation type: spf
- metric type: igp-metric
- priority: 0
- Link constraints: none



NOTE: Modifying the flexible algorithm definition in a live network or on the fly could cause traffic disruptions until all the nodes converge on the new paths.

Starting in Junos OS 21.4R1, we support flexible Algorithm Definition (FAD) and “Flexible Algorithm Prefix Metric (FAPM)” in TED and implements two new corresponding TLVs “FAD TLV” and “FAPM TLV” in BGP-LS. The value of FAD TLV contains Flex-Algorithm, Metric-Type, Calculation-Type and Priority, all of which are one byte each. The TLV might have zero or more sub-TLVs included in it. The five sub-tlvs are Flex Algo Exclude Any Affinity, Flex Algo Include Any Affinity, Flex Algo Include All Affinity, Flex Algo Definition Flags and Flex Algo Exclude SRLG.

The FAD TLV can only be added to the BGP-LS Attribute of the Node NLRI if the corresponding node originates in the underlying IGP TLV or sub-TLV. The BGP-LS Attribute associated with a Node NLRI

might include one or more FAD TLVs corresponding to the Flexible Algorithm Definition for each algorithm that the node is advertising.

The value of FAPM TLV contains Flex-Algorithm (1 byte), Reserved (3 bytes) and Metric (4 bytes). The FAPM TLV can be added to the BGP-LS Attribute of the Prefix NLRI originated by a node, only if the corresponding node originates from the Prefix.

Starting in Junos OS Release 22.4R1, we've defined the Flexible Algorithm Prefix Metric (FAPM) to allow optimal end-to-end path for an interarea prefix. The area border router (ABR) must include the FAPM when advertising the prefix between areas that is reachable in that given Flexible Algorithm (flex algo). When a prefix is unreachable, the ABR must not include that prefix in that flex algo when advertising between areas. The defined FAPM provides inter-area support.

Participation in a Flexible Algorithm

You can configure specific routers to participate in a particular flexible algorithm as per your requirement. Paths computed based on a flexible algorithm definition is used by various applications each potentially using its own specific data plane for forwarding the data over such paths. The participating device must explicitly advertise its participation in a particular flexible algorithm to every application in the segment routing flexible algorithm sub TLV for OSPFv2. You can configure a node to participate in a certain flexible algorithm provided it can support the constraints specified in that FAD.

To configure participation in a flexible algorithm include the `flex-algorithm` statement at the `[edit protocols isis source-packet-routing]` hierarchy level. The same device can advertise a FAD and also participate in a flexible algorithm.

You can control path selection by configuring the preference for OSPF Flexible Algorithm routes in `inetcolor.0` and `mpls.0` routing tables.

Network Topology Configured with Flexible Algorithm Definitions

[Figure 28 on page 500](#) shows the sample topology, there are 8 routers R0, R1, R2, R3, R4, R5, R6, and R7. Four flexible algorithms, 128, 129, 130, and 135 are defined and configured with admin-groups as listed in the following table:

Flex Algorithm Definition (FAD)	Color
128	Include any Red
129	Include any Green

(Continued)

Flex Algorithm Definition (FAD)	Color
130	Include any Green and Blue
135	Exclude Red

Figure 28: Flexible Algorithm Topology




Figure 29 on page 501 shows how FAD 128 routes traffic on any interface that is configured with admin group red.

Figure 29: Traffic Flow for Flexible Algorithm Definition 128




Figure 30 on page 502 shows how FAD 129 routes traffic on any interface that is configured with admin group green.

Figure 30: Traffic Flow for Flexible Algorithm Definition 129




Figure 31 on page 503 shows how FAD 130 routes traffic on any interface that is configured with admin group green and blue.

Figure 31: Traffic flow for Flexible Algorithm Definition 130





Figure 32 on page 504 shows how FAD 135 routes traffic on any interface that is not configured with admin group red.

Figure 32: Traffic Flow for Flexible Algorithm Definition 135



Flexible Algorithm RIBs

For every flexible algorithm that a router participates in the corresponding flexible algorithm routes are installed in the corresponding flexible algorithm RIB groups also known as routing tables. By default, labeled OSPFv2 flexible algorithm routes are installed in the `inet.color`, `inet(6)color.0` and `mpls.0` RIBs.

BGP Community and Flexible Algorithms

A flexible algorithm can have an associated BGP color community to resolve routes of other services such as VPN service. By default, the associated BGP color community is the same as the flexible algorithm ID. The flexible algorithm ingress routes that are installed in the `inet(6)color.0` tables will have this color community in the route. However, you can configure a different BGP color community at the `[edit routing-options flex-algorithm id color desired_color_community value]` hierarchy level.



NOTE: Changing the BGP color community for a flexible algorithm might result in traffic disruption. If you modify a BGP color community for a flexible algorithm then all routes pertaining to that flexible algorithm are removed from the RIB and added again with new colors.

Supported and Unsupported Features

Junos OS supports flexible algorithms in the following scenarios:

- Support for configuring and advertising prefix SIDs for different flexible algorithms.
- Partially supports Internet Draft draft-ietf-lsr-flex-algo-05.txt *IGP Flexible Algorithm*
- The current implementation for flexible algorithms is supported for only OSPFv2 only as only OSPFv2 supports segment routing.

Junos OS does not support the following features in conjunction with flexible algorithms:

- Flexible algorithm is applicable only for default unicast topology, OSPFv2 multi-topology is not supported.
- OSPFv2 shortcuts and other OSPFv2 traffic engineering configuration options are not applicable for flexible algorithm computation. .
- The current implementation for flexible algorithms is not supported for OSPFv3.
- Inter-area (OSPFv2) leaking of flexible algorithm prefix SIDs is not supported.
- Prefix and SID conflict resolution is not supported.
- Remote loop free alternate functionality is not supported because TI-LFA is the preferred FRR computation.
- Advertising flexible algorithm definition in the absence of flexible algorithm participation is not supported.
- Transport class RIB is not supported.

Example: OSPF Flexible Algorithm

IN THIS SECTION

- [Overview | 506](#)
- [Requirements | 507](#)
- [Configuration | 507](#)
- [Verification | 528](#)

Overview

IN THIS SECTION

- [Topology | 506](#)


This example shows how to configure flexible algorithm in an OSPFv2 network. The flexible algorithm allows networks without a controller to configure traffic engineering using segment routing without actually implementing a network controller.

You can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. The set consisting of calculation-type, metric-type, and a set of constraints is referred to as a flexible algorithm definition (FAD). You can define FADs and advertise the same in an OSPFv2 network. A device can also be configured to participate in a certain flexible algorithm provided it supports the constraints for that specific FAD.

Topology

Figure 6 shows a flexible algorithm topology in which there are 6 devices R0, R1, R2, R3, R4, and R5. Two flexible algorithms 128 and 129 are defined on each of these devices. The admin-groups red, blue, and green are configured on the devices. The FADs with different parameters such as metric-types, calculation-types, and link constraints are defined on each of the devices.

Figure 33: Flexible Algorithm Topology



Requirements

This example uses the following hardware and software components:

- Six MX Series routers.
- Junos OS Release 21.1R1 or later running on all devices.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 508](#)
- [Configuring Device R0 | 518](#)
- [Results | 524](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R0

```
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R0_to_R1_1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.1.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R0_to_R1_2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.1.5/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R0_to_R3_1
set interfaces ge-0/0/2 unit 0 family inet address 10.10.3.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R0_to_R3_2
set interfaces ge-0/0/3 unit 0 family inet address 10.10.3.5/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.10/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement ex-bgp term 1 from route-filter 10.1.1.0/24 exact
set policy-options policy-statement ex-bgp term 1 then community add blue
set policy-options policy-statement ex-bgp term 1 then accept
set policy-options policy-statement ex-bgp term 0 from route-filter 10.1.0.0/24 exact
set policy-options policy-statement ex-bgp term 0 then community add red
set policy-options policy-statement ex-bgp term 0 then accept
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.10/32
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1280
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1290
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 130 index
1300
```

```
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 130 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 131 index
1310
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 131 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 132 index
1320
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 132 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 133 index
1330
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 133 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 134 index
1340
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 134 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 135 index
1350
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 135 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1000
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1001 then accept
set policy-options community blue members color:1:129
set policy-options community red members color:0:128
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group RED
set protocols mpls interface ge-0/0/1.0 admin-group GREEN
set protocols mpls interface ge-0/0/2.0 admin-group RED
set protocols mpls interface ge-0/0/3.0 admin-group GREEN
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
```

```

set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
set routing-options flex-algorithm 128 definition metric-type igp-metric
set routing-options flex-algorithm 128 definition spf
set routing-options flex-algorithm 128 definition admin-group include-any RED
set routing-options flex-algorithm 129 definition metric-type te-metric
set routing-options flex-algorithm 129 definition spf
set routing-options flex-algorithm 129 definition admin-group include-all BLUE
set routing-options router-id 192.168.255.10
set routing-options autonomous-system 65000
set routing-options static route 10.1.1.0/24 receive
set routing-options static route 10.1.0.0/24 receive
set routing-options forwarding-table export pplb
set routing-options flex-algorithm 128 use-transport-class
set routing-options transport-class auto-create

```

Device R1

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R1_to_R0_1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.1.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R1_to_R0_2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.1.6/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R1_to_R2_1
set interfaces ge-0/0/2 unit 0 family inet address 10.10.12.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R1_to_R2_2
set interfaces ge-0/0/3 unit 0 family inet address 10.10.12.5/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 description R1_to_R3
set interfaces ge-0/0/4 unit 0 family inet address 10.10.13.1/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 description R1_to_R4
set interfaces ge-0/0/5 unit 0 family inet address 10.10.14.1/30

```

```
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.11/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.11/32
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1281
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1291
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1001
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1001 then accept
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group RED
set protocols mpls interface ge-0/0/1.0 admin-group GREEN
set protocols mpls interface ge-0/0/2.0 admin-group RED
set protocols mpls interface ge-0/0/3.0 admin-group GREEN
set protocols mpls interface ge-0/0/4.0 admin-group BLUE
set protocols mpls interface ge-0/0/5.0 admin-group GREEN
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
```

```

set routing-options router-id 192.168.255.11
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb

```

Device R2

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R2_to_R1_1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R2_to_R1_2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.12.6/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R2_to_R3

set interfaces ge-0/0/2 unit 0 family inet address 10.10.23.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R2_to_R4
set interfaces ge-0/0/3 unit 0 family inet address 10.10.24.1/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 description R2_to_R5_1
set interfaces ge-0/0/4 unit 0 family inet address 10.10.25.1/30

set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 description R2_to_R5_2
set interfaces ge-0/0/5 unit 0 family inet address 10.10.25.5/30
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.12/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.12/32
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1282
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1292
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1002
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment

```

```

set policy-options policy-statement prefix-sid term 1001 then accept
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group RED
set protocols mpls interface ge-0/0/1.0 admin-group GREEN
set protocols mpls interface ge-0/0/2.0 admin-group RED
set protocols mpls interface ge-0/0/3.0 admin-group BLUE
set protocols mpls interface ge-0/0/4.0 admin-group RED
set protocols mpls interface ge-0/0/5.0 admin-group GREEN
set protocols mpls interface ge-0/0/5.0 admin-group BLUE
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 post-convergence-lfa node-protection
set routing-options router-id 192.168.255.12
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb

```

Device R3

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R3_to_R0_1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.3.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R3_to_R0_2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.3.6/30

```

```
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R3_to_R1
set interfaces ge-0/0/2 unit 0 family inet address 10.10.13.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R3_to_R2
set interfaces ge-0/0/3 unit 0 family inet address 10.10.23.2/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 description R3_to_R4_1
set interfaces ge-0/0/4 unit 0 family inet address 10.10.34.1/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 description R3_to_R4_2
set interfaces ge-0/0/5 unit 0 family inet address 10.10.34.5/30
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.13/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.13/32
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1284
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1294
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1003
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1001 then accept
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group RED
set protocols mpls interface ge-0/0/1.0 admin-group GREEN
set protocols mpls interface ge-0/0/2.0 admin-group BLUE
set protocols mpls interface ge-0/0/3.0 admin-group RED
set protocols mpls interface ge-0/0/4.0 admin-group RED
set protocols mpls interface ge-0/0/5.0 admin-group BLUE
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
```

```

set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 post-convergence-lfa node-protection
set routing-options router-id 192.168.255.13
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb

```

Device R4

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R4_to_R1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.14.2/30

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R4_to_R2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.24.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R4_to_R3_1
set interfaces ge-0/0/2 unit 0 family inet address 10.10.34.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R4_to_R3_2
set interfaces ge-0/0/3 unit 0 family inet address 10.10.34.6/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 description R4_to_R5_1
set interfaces ge-0/0/4 unit 0 family inet address 10.10.45.1/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 description R4_to_R5_2
set interfaces ge-0/0/5 unit 0 family inet address 10.10.45.5/30
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.14/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.14/32

```

```
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1284
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1294
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1004
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1001 then accept
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/2.0 admin-group RED
set protocols mpls interface ge-0/0/3.0 admin-group BLUE
set protocols mpls interface ge-0/0/0.0 admin-group GREEN
set protocols mpls interface ge-0/0/1.0 admin-group BLUE
set protocols mpls interface ge-0/0/4.0 admin-group GREEN
set protocols mpls interface ge-0/0/5.0 admin-group RED
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 post-convergence-lfa node-protection
set routing-options router-id 192.168.255.14
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
```

Device R5

```
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description R5_to_R2_1
set interfaces ge-0/0/0 unit 0 family inet address 10.10.25.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description R5_to_R2_2
set interfaces ge-0/0/1 unit 0 family inet address 10.10.25.6/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description R5_to_R4_1
set interfaces ge-0/0/2 unit 0 family inet address 10.10.45.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 description R5_to_R4_2
set interfaces ge-0/0/3 unit 0 family inet address 10.10.45.6/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.15/32
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1001 from route-filter 192.168.255.15/32
exact
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 index
1285
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 128 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 index
1295
set policy-options policy-statement prefix-sid term 1001 then prefix-segment algorithm 129 node-
segment
set policy-options policy-statement prefix-sid term 1001 then prefix-segment index 1005
set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1001 then accept
set protocols mpls admin-groups RED 0
set protocols mpls admin-groups BLUE 1
set protocols mpls admin-groups GREEN 2
set protocols mpls label-range static-label-range 1000 8000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group RED
set protocols mpls interface ge-0/0/1.0 admin-group GREEN
set protocols mpls interface ge-0/0/1.0 admin-group BLUE
set protocols mpls interface ge-0/0/2.0 admin-group GREEN
set protocols mpls interface ge-0/0/3.0 admin-group RED
```

```

set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf traffic-engineering advertisement always
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing srgb start-label 80000
set protocols ospf source-packet-routing srgb index-range 5000
set protocols ospf source-packet-routing flex-algorithm 128
set protocols ospf source-packet-routing flex-algorithm 129
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-protection
set routing-options router-id 192.168.255.15
set routing-options autonomous-system 65000
set routing-options static route 10.1.15.0/24 reject
set routing-options forwarding-table export pplb

```

Configuring Device R0

To configure flexible algorithm for OSPFv2, perform the following steps on the device R0:

1. Configure the device interfaces to enable IP transport.

```

[edit]
user@R0set interfaces ge-0/0/0 description R0_to_R1_1
user@R0set interfaces ge-0/0/0 unit 0 family inet address 10.10.1.1/30
user@R0set interfaces ge-0/0/0 unit 0 family mpls
user@R0set interfaces ge-0/0/1 description R0_to_R1_2
user@R0set interfaces ge-0/0/1 unit 0 family inet address 10.10.1.5/30
user@R0set interfaces ge-0/0/1 unit 0 family mpls
user@R0set interfaces ge-0/0/2 description R0_to_R3_1
user@R0set interfaces ge-0/0/2 unit 0 family inet address 10.10.3.1/30
user@R0set interfaces ge-0/0/2 unit 0 family mpls
user@R0set interfaces ge-0/0/3 description R0_to_R3_2
user@R0set interfaces ge-0/0/3 unit 0 family inet address 10.10.3.5/30
user@R0set interfaces ge-0/0/3 unit 0 family mpls

```

2. Configure the loopback interface (lo0) address that is used as router ID for OSPF sessions.

```
[edit]
user@R0set interfaces lo0 unit 0 family inet address 192.168.255.10/32
```

3. Configure the router ID and autonomous system (AS) number to propagate routing information within a set of routing devices that belong to the same AS.

```
[edit]
user@R0set routing-options router-id 192.168.255.10
user@R0set routing-options autonomous-system 65000
```

4. Define a policy to load balance packets and apply the per-packet policy to enable load balancing of traffic.

```
[edit]
user@R0set policy-options policy-statement pplb then load-balance per-packet
user@R0set routing-options forwarding-table export pplb
```

5. Configure the route filter for the routing policy term that enables the Device R0 to reach the 192.168.255.10/32 network.

```
[edit]
user@R0set policy-options policy-statement prefix-sid term 1001 from route-filter
192.168.255.10/32 exact
```

6. Configure MPLS on all interfaces excluding the management interface.

```
[edit]
user@R0set protocols mpls interface all
user@R0set protocols mpls interface fxp0.0 disable
```

7. Configure the MPLS label range to assign static labels for the links.

```
[edit]
user@R0set protocols mpls label-range static-label-range 1000 8000
```

8. Configure TI-LFA to enable protection against link and node failures. SR using TI-LFA provides faster restoration of network connectivity by routing the traffic instantly to a backup or an alternate path if the primary path fails or becomes unavailable.

```
[edit]
user@R0set protocols ospf backup-spf-options use-source-packet-routing
```

9. Configure the maximum number of labels for segment routing routed paths for protection of backup shortest-path-first attributes.

```
[edit]
user@R0set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
```

10. Configure prefix segment attributes, the start label and the index range for segment routing global blocks (SRGBs) in SPRING for the OSPF protocol.

```
[edit]
user@R0set protocols ospf source-packet-routing prefix-segment prefix-sid
user@R0set protocols ospf source-packet-routing srgb start-label 80000
user@R0set protocols ospf source-packet-routing srgb index-range 5000
```

11. Enable node-link protection on the OSPF interfaces that follow post-convergence path.

```
[edit]
user@R0set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 post-convergence-lfa node-
protection
user@R0set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 post-convergence-lfa node-
protection
```

```

user@R0set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 post-convergence-lfa node-
protection
user@R0set protocols ospf area 0.0.0.0 interface ge-0/0/3.0 post-convergence-lfa node-
protection

```

12. Configure the loopback interface as passive to ensure the protocols do not run over the loopback interface and that the loopback interface is advertised correctly throughout the network.

```

[edit]
user@R0set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

13. Define flexible algorithms on the device R0. Assign a name for each of the FADs ranging from 128 through 255.

```

[edit]
user@R0set routing-options flex-algorithm 128
user@R0set routing-options flex-algorithm 129

```

Specify the parameters of the definition. OSPFv2 calculates the path based on these specified parameters of the FAD.

- a. Specify the calculation type based on which the OSPFv2 protocol calculates the path.

```

[edit]
user@R0set routing-options flex-algorithm 128 definition spf
user@R0set routing-options flex-algorithm 128 definition spf

```

- b. Specify the metric type based on which OSPFv2 calculates the path.

```

[edit]
user@R0set routing-options flex-algorithm 128 definition metric-type igp-metric
user@R0set routing-options flex-algorithm 129 definition metric-type te-metric

```

- c. If you have enabled RSVP traffic engineering, you can configure admin-groups for many protocols to color an individual link.

```
[edit]
user@R0set protocols mpls admin-groups RED 0
user@R0set protocols mpls admin-groups BLUE 1
user@R0set protocols mpls admin-groups GREEN 2
```

- d. Assign the configured admin-groups policies to the device R0 interfaces.

```
[edit]
user@R0set protocols mpls interface ge-0/0/0.0 admin-group RED
user@R0set protocols mpls interface ge-0/0/1.0 admin-group GREEN
user@R0set protocols mpls interface ge-0/0/2.0 admin-group RED
user@R0set protocols mpls interface ge-0/0/3.0 admin-group GREEN
```

- e. Define the admin-groups as per your requirement.

```
[edit]
user@R0set routing-options flex-algorithm 128 definition admin-group include-any RED
user@R0set routing-options flex-algorithm 129 definition admin-group include-all GREEN
user@R0set routing-options flex-algorithm 129 definition admin-group include-all BLUE
```



NOTE: For FADs with link-constraints to work, all relevant links should advertise the admin-colors in OSPFv2. You must either enable RSVP on the interfaces or if you have not configured RSVP for traffic engineering, make sure you configure `set traffic-engineering advertise always` at the `[edit protocols ospf]` hierarchy level.

```
[edit]
user@R0set protocols ospf traffic-engineering advertisement always
```

14. Configure the flexible algorithm participation on the device R0. The same device can advertise a FAD and also participate in a flexible algorithm.

```
[edit]
user@R0set protocols ospf source-packet-routing flex-algorithm 128
user@R0set protocols ospf source-packet-routing flex-algorithm 129
```

15. Advertise prefix segments through policy configuration.

```
[edit]
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 128 index 1280
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 128 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 129 index 1290
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 129 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 130 index 1300
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 130 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 131 index 1310
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 131 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 132 index 1320
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 132 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 133 index 1330
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 133 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 134 index 1340
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 134 node-segment
user@R0set policy-options policy-statement prefix-sid term 1001 then prefix-segment
```

```

algorithm 135 index 1350
user@R0#set policy-options policy-statement prefix-sid term 1001 then prefix-segment
algorithm 135 node-segment
user@R0#set policy-options policy-statement prefix-sid term 1001 then prefix-segment index
1000
user@R0#set policy-options policy-statement prefix-sid term 1001 then prefix-segment node-
segment
user@R0#set policy-options policy-statement prefix-sid term 1001 then accept

```

Results

Check the results of the configuration:

From configuration mode, confirm your configuration by entering the, show interfaces, show routing-options, show protocols, and show policy-options commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

interfaces {
  ge-0/0/0 {
    description R0_to_R1_1;
    unit 0 {
      family inet {
        address 10.10.1.1/30;
      }
      family mpls;
    }
  }
  ge-0/0/1 {
    description R0_to_R1_2
    unit 0 {
      family inet {
        address 10.10.1.5/30;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    description R0_to_R3_1
    unit 0 {
      family inet {
        address 10.10.3.1/30;
      }
    }
  }
}

```

```
        family mpls;
    }
}
ge-0/0/3 {
    description R0_to_R3_2
    unit 0 {
        family inet {
            address 10.10.3.5/30;
        }
        family mpls;
    }
}

lo0 {
    unit 0 {
        family inet {
            address 192.168.255.10/32;
        }
    }
}

}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
    policy-statement prefix-sid {
        term 1001 {
            from {
                route-filter 192.168.255.10/32 exact;
            }
            then {
                prefix-segment {
                    algorithm 128 index 1280 node-segment;
                    algorithm 129 index 1290 node-segment;
                    algorithm 130 index 1300 node-segment;
                    algorithm 131 index 1310 node-segment;
                    algorithm 132 index 1320 node-segment;
                    algorithm 133 index 1330 node-segment;
                    algorithm 134 index 1340 node-segment;
                    algorithm 135 index 1350 node-segment;
                    index 1000;
                }
            }
        }
    }
}
```



```
        admin-group include-all [ GREEN BLUE ];
    }
}
router-id 192.168.255.10;
autonomous-system 65000;
forwarding-table {
    export pplb;
}
}
```

Verification

IN THIS SECTION

- [Verifying the OSPF Database | 528](#)
- [Action | 528](#)
- [Verifying the Flexible Algorithm Details | 530](#)
- [Action | 530](#)
- [Verifying Flexible Algorithm Specific OSPF Internal Routes | 531](#)
- [Action | 531](#)
- [Verifying Flex Colored routes | 534](#)
- [Action | 534](#)
- [Verifying OSPF Logs | 535](#)
- [Action | 535](#)

To confirm that the configuration is working properly, perform the following tasks:

Verifying the OSPF Database

Purpose

Verifying that the flexible algorithm signaling is displayed in the OSPF database.

Action

From operational mode, run the `show ospf database opaque-area extensive` command.

On R0

```
user@R0>show ospf database opaque-area extensive

OpaqArea 4.0.0.0          192.168.255.10  0x800000ad  503  0x22 0xb85d  76
Opaque LSA
  SR-Algorithm (8), length 3:
    Algo (1), length 1:
      0
    Algo (2), length 1:
      128
    Algo (3), length 1:
      129
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
      5000
    SID/Label (1), length 3:
      Label (1), length 3:
        80000
  Flex-Algorithm Definition (16), length 12:
    Flex-Algo (1), length 1:
      128
    Metric-Type (2), length 1:
      0
    Calc-Type (3), length 1:
      0
    Priority (4), length 1:
      0
    FAD AG Include Any (2), length 4:
      Include Any AG (1), length 4:
        0x1
  Flex-Algorithm Definition (16), length 12:
    Flex-Algo (1), length 1:
      129
    Metric-Type (2), length 1:
      2
    Calc-Type (3), length 1:
      0
    Priority (4), length 1:
      0
    FAD AG Include All (3), length 4:
      Include All AG (1), length 4:
```

```

0x6
Aging timer 00:51:37
Installed 00:08:20 ago, expires in 00:51:37, sent 00:08:18 ago
Last changed 5d 13:35:52 ago, Change count:

```

Meaning

This output on R0 illustrates that:

Three segment-routing algorithms (including two flexible algorithms) are advertised by this device.

Two FADs are advertised by this device.

Verifying the Flexible Algorithm Details

Purpose

Verifying that the flexible algorithm details are displayed.

Action

From operational mode, run the `show ospf spring flex-algorithm <flex-algorithm-id>` command.

On R0

```

user@R0>show ospf spring flex-algorithm 128
Flex Algo: 128, Area: 0.0.0.0
Color: 128, Participating, FAD supported
  Winner: 192.168.255.10, Metric: 0, Calc: 0, Prio: 0, inc-any: 0x1, FAD supported
  Include-Any: 0x1 RED
SPF Version: 296
Participation toggles: 1
Topo refresh count: 0
Full SPFs: 296, Partial SPFs: 0

```

Meaning

The flexible algorithm details that are configured on R0 are displayed.

Verifying Flexible Algorithm Specific OSPF Internal Routes

Purpose

Verifying that the flexible algorithm specific OSPF internal routes are displayed.

Action

From operational mode, run the `show ospf route flex-algorithm <flex-algorithm-id>` command.

On R0

```

user@R0>show ospf route flex-algorithm 128
Prefix          Path Route   NH      Metric NextHop      Nexthop
                Type Type    Type                    Interface      Address/LSP
192.168.255.11  Intra Router IP        1 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
192.168.255.12  Intra Router IP        2 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
                ge-0/0/2.0   10.10.3.2
                ge-0/0/3.0   10.10.3.6
192.168.255.13  Intra Router IP        1 ge-0/0/2.0   10.10.3.2
                ge-0/0/3.0   10.10.3.6
192.168.255.14  Intra Router IP        2 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
                ge-0/0/2.0   10.10.3.2
                ge-0/0/3.0   10.10.3.6
192.168.255.15  Intra Router IP        3 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
                ge-0/0/2.0   10.10.3.2
                ge-0/0/3.0   10.10.3.6
10.10.1.0/30    Intra Network IP        1 ge-0/0/0.0
10.10.1.4/30    Intra Network IP        1 ge-0/0/1.0
10.10.3.0/30    Intra Network IP        1 ge-0/0/2.0
10.10.3.4/30    Intra Network IP        1 ge-0/0/3.0
10.10.12.0/30   Intra Network IP        2 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
10.10.12.4/30   Intra Network IP        2 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
10.10.13.0/30   Intra Network IP        2 ge-0/0/0.0   10.10.1.2
                ge-0/0/1.0   10.10.1.6
                ge-0/0/2.0   10.10.3.2

```

			ge-0/0/3.0	10.10.3.6
10.10.14.0/30	Intra Network	IP	2 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
10.10.23.0/30	Intra Network	IP	2 ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.24.0/30	Intra Network	IP	3 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.25.0/30	Intra Network	IP	3 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.25.4/30	Intra Network	IP	3 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.34.0/30	Intra Network	IP	2 ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.34.4/30	Intra Network	IP	2 ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.45.0/30	Intra Network	IP	3 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
10.10.45.4/30	Intra Network	IP	3 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
128.49.106.245/32	Intra Network	IP	1 ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
128.49.107.40/32	Intra Network	IP	2 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
			ge-0/0/2.0	10.10.3.2
			ge-0/0/3.0	10.10.3.6
192.168.255.10/32	Intra Network	IP	0 lo0.0	
192.168.255.10/32	Intra Network	Spring	0 lo0.0	
192.168.255.11/32	Intra Network	IP	1 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
192.168.255.11/32	Intra Network	Spring	1 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6
192.168.255.12/32	Intra Network	IP	2 ge-0/0/0.0	10.10.1.2
			ge-0/0/1.0	10.10.1.6

				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.12/32	Intra Network	Spring	2	ge-0/0/0.0	10.10.1.2
				ge-0/0/1.0	10.10.1.6
				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.13/32	Intra Network	IP	1	ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.13/32	Intra Network	Spring	1	ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.14/32	Intra Network	IP	2	ge-0/0/0.0	10.10.1.2
				ge-0/0/1.0	10.10.1.6
				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.14/32	Intra Network	Spring	2	ge-0/0/0.0	10.10.1.2
				ge-0/0/1.0	10.10.1.6
				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.15/32	Intra Network	IP	3	ge-0/0/0.0	10.10.1.2
				ge-0/0/1.0	10.10.1.6
				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
192.168.255.15/32	Intra Network	Spring	3	ge-0/0/0.0	10.10.1.2
				ge-0/0/1.0	10.10.1.6
				ge-0/0/2.0	10.10.3.2
				ge-0/0/3.0	10.10.3.6
16	Intra Network	Mpls	0	ge-0/0/0.0	10.10.1.2
		Bkup MPLS		ge-0/0/1.0	10.10.1.6
16 (S=0)	Intra Network	Mpls	0	ge-0/0/0.0	10.10.1.2
		Bkup MPLS		ge-0/0/1.0	10.10.1.6
17	Intra Network	Mpls	0	ge-0/0/1.0	10.10.1.6
		Bkup MPLS		ge-0/0/0.0	10.10.1.2
17 (S=0)	Intra Network	Mpls	0	ge-0/0/1.0	10.10.1.6
		Bkup MPLS		ge-0/0/0.0	10.10.1.2
20	Intra Network	Mpls	0	ge-0/0/2.0	10.10.3.2
		Bkup MPLS		ge-0/0/3.0	10.10.3.6
20 (S=0)	Intra Network	Mpls	0	ge-0/0/2.0	10.10.3.2
		Bkup MPLS		ge-0/0/3.0	10.10.3.6
21	Intra Network	Mpls	0	ge-0/0/3.0	10.10.3.6
		Bkup MPLS		ge-0/0/2.0	10.10.3.2
21 (S=0)	Intra Network	Mpls	0	ge-0/0/3.0	10.10.3.6
		Bkup MPLS		ge-0/0/2.0	10.10.3.2
81001	Intra Network	Mpls	1	ge-0/0/0.0	10.10.1.2

```

      ge-0/0/1.0    10.10.1.6
81001 (S=0)      Intra Network    Mpls      1 ge-0/0/0.0    10.10.1.2
      ge-0/0/1.0    10.10.1.6
81003            Intra Network    Mpls      2 ge-0/0/0.0    10.10.1.2
      ge-0/0/1.0    10.10.1.6
      ge-0/0/2.0    10.10.3.2
      ge-0/0/3.0    10.10.3.6
81004            Intra Network    Mpls      1 ge-0/0/2.0    10.10.3.2
      ge-0/0/3.0    10.10.3.6
81004 (S=0)      Intra Network    Mpls      1 ge-0/0/2.0    10.10.3.2
      ge-0/0/3.0    10.10.3.6
81006            Intra Network    Mpls      2 ge-0/0/0.0    10.10.1.2
      ge-0/0/1.0    10.10.1.6
      ge-0/0/2.0    10.10.3.2
      ge-0/0/3.0    10.10.3.6
81007            Intra Network    Mpls      3 ge-0/0/0.0    10.10.1.2
      ge-0/0/1.0    10.10.1.6
      ge-0/0/2.0    10.10.3.2
      ge-0/0/3.0    10.10.3.6

```

Meaning

The `show ospf route` command is extended with `flex-algorithm` option to show flexible algorithm specific OSPF internal routes. Each route is prefixed with the *flex-algo-id*.

Verifying Flex Colored routes

Purpose

Verifying that the flexible algorithm specific OSPF internal routes are displayed.

Action

From operational mode, run the `show route protocol ospf` command.

On R0

```

user@R0>show route protocol ospf
junos-rti-tc-<color>.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

192.168.255.11-128<c>/64
    *[L-OSPF/10/5] 1w2d 01:23:04, metric 1
    > to 10.10.1.2 via ge-0/0/0.0
    to 10.10.3.2 via ge-0/0/2.0, Push 81281, Push 81283(top)
192.168.255.12-128<c>/64
    *[L-OSPF/10/5] 1w2d 01:23:04, metric 2
    to 10.10.1.2 via ge-0/0/0.0, Push 81283
    > to 10.10.3.2 via ge-0/0/2.0, Push 81283
192.168.255.13-128<c>/64
    *[L-OSPF/10/5] 1w2d 01:23:04, metric 1
    > to 10.10.3.2 via ge-0/0/2.0
    to 10.10.1.2 via ge-0/0/0.0, Push 81284, Push 81283(top)
192.168.255.14-128<c>/64
    *[L-OSPF/10/5] 1w2d 01:23:04, metric 2
    > to 10.10.3.2 via ge-0/0/2.0, Push 81286
    to 10.10.1.2 via ge-0/0/0.0, Push 81286, Push 81283(top)
192.168.255.15-128<c>/64
    *[L-OSPF/10/5] 1w2d 01:23:04, metric 3
    to 10.10.1.2 via ge-0/0/0.0, Push 81287
    > to 10.10.3.2 via ge-0/0/2.0, Push 81287

```

Meaning

The output displays all the colored flex routes programmed in `junos-rti-tc-<color>.inet.0` table in the following format: `prefix_address-flex-algo-id<c>/64`

Verifying OSPF Logs

Purpose

Verifying that the OSPF logs displays the flexible algorithm keyword.

Action

From operational mode, run the `show ospf log` command.

On R0

```

user@R0>show ospf log
Topology default SPF log:

```

Last instance of each event type

When	Type	Elapsed
1w2d 13:59:18	SPF	0.000316
1w2d 13:59:18	Stub	0.000233
1w2d 13:59:18	Interarea	0.000002
1w2d 13:59:18	External	0.000004
1w2d 13:59:18	NSSA	0.000001
1w2d 13:59:18	Cleanup	0.000551

Maximum length of each event type

When	Type	Elapsed
1w2d 14:34:27	SPF	0.000997
1w2d 15:59:35	Stub	0.000675
1w3d 07:08:27	Interarea	0.000010
1w3d 07:29:07	External	0.000013
1w3d 07:15:21	NSSA	0.000008
1w3d 08:38:05	Cleanup	0.001044

Last 100 events

When	Type	Elapsed
1w2d 14:08:36	FlexAlgo SPF	0.000680
1w2d 14:08:36	SPF	0.000204
1w2d 14:08:36	Stub	0.000025
1w2d 14:08:36	Interarea	0.000003
1w2d 14:08:36	External	0.000002
1w2d 14:08:36	NSSA	0.000001
1w2d 14:08:36	Prefix SID	0.000222
1w2d 14:08:36	Adj SID	0.000074
1w2d 14:08:36	Cleanup	0.000607
1w2d 14:08:36	Total	0.001209
1w2d 14:08:31	SPF	0.000188
1w2d 14:08:31	Stub	0.000054
1w2d 14:08:31	Interarea	0.000002
1w2d 14:08:31	External	0.000001
1w2d 14:08:31	NSSA	0.000001
1w2d 14:08:31	Prefix SID	0.000181
1w2d 14:08:31	Adj SID	0.000178
1w2d 14:08:31	Cleanup	0.000413
1w2d 14:08:31	Total	0.001656
1w2d 14:06:54	FlexAlgo SPF	0.001914
1w2d 14:06:54	FlexAlgo SPF	0.000081
1w2d 14:06:54	SPF	0.000215

1w2d 14:06:54	Stub	0.000030
1w2d 14:06:54	Interarea	0.000003
1w2d 14:06:54	External	0.000001
1w2d 14:06:54	NSSA	0.000001
1w2d 14:06:54	Prefix SID	0.000227
1w2d 14:06:54	Adj SID	0.000075
1w2d 14:06:54	Cleanup	0.000233
1w2d 14:06:54	Total	0.000859
1w2d 14:06:49	SPF	0.000234
1w2d 14:06:49	Stub	0.000072
1w2d 14:06:49	Interarea	0.000003
1w2d 14:06:49	External	0.000002
1w2d 14:06:49	NSSA	0
1w2d 14:06:49	Prefix SID	0.000262
1w2d 14:06:49	Adj SID	0.000254
1w2d 14:06:49	Cleanup	0.000495
1w2d 14:06:49	Total	0.001936
1w2d 14:06:30	FlexAlgo SPF	0.001356
1w2d 14:06:30	FlexAlgo SPF	0.000061
1w2d 14:06:30	SPF	0.000207
1w2d 14:06:30	Stub	0.000023
1w2d 14:06:30	Interarea	0.000003
1w2d 14:06:30	External	0.000002
1w2d 14:06:30	NSSA	0.000001
1w2d 14:06:30	Prefix SID	0.000237
1w2d 14:06:30	Adj SID	0.000087
1w2d 14:06:30	Cleanup	0.000430
1w2d 14:06:30	Total	0.001060
1w2d 14:06:25	SPF	0.000207
1w2d 14:06:25	Stub	0.000077
1w2d 14:06:25	Interarea	0.000002
1w2d 14:06:25	External	0.000002
1w2d 14:06:25	NSSA	0.000001
1w2d 14:06:25	Prefix SID	0.000250
1w2d 14:06:25	Adj SID	0.000245
1w2d 14:06:25	Cleanup	0.000399
1w2d 14:06:25	Total	0.001840
1w2d 14:05:56	FlexAlgo SPF	0.001781
1w2d 14:05:56	FlexAlgo SPF	0.000080
1w2d 14:05:55	SPF	0.000215
1w2d 14:05:55	Stub	0.000025
1w2d 14:05:55	Interarea	0.000002
1w2d 14:05:55	External	0.000001

```

1w2d 14:05:55  NSSA          0.000001
1w2d 14:05:55  Prefix SID    0.000240
1w2d 14:05:55  Adj SID      0.000073
1w2d 14:05:55  Cleanup      0.000422
1w2d 14:05:55  Total        0.001055
1w2d 14:05:50  SPF          0.000212
1w2d 14:05:50  Stub         0.000082
1w2d 14:05:50  Interarea    0.000003
1w2d 14:05:50  External     0.000001
1w2d 14:05:50  NSSA         0.000001
1w2d 14:05:50  Prefix SID   0.000264
1w2d 14:05:50  Adj SID      0.000239
1w2d 14:05:50  Cleanup      0.000458
1w2d 14:05:50  Total        0.002053
1w2d 13:59:23  FlexAlgo SPF 0.001603
1w2d 13:59:23  FlexAlgo SPF 0.000062
1w2d 13:59:23  SPF          0.000224
1w2d 13:59:23  Stub         0.000021
1w2d 13:59:23  Interarea    0.000002
1w2d 13:59:23  External     0.000001
1w2d 13:59:23  NSSA         0.000001
1w2d 13:59:23  Prefix SID   0.000222
1w2d 13:59:23  Adj SID      0.000087
1w2d 13:59:23  Cleanup      0.000413
1w2d 13:59:23  Total        0.001228
1w2d 13:59:18  SPF          0.000316
1w2d 13:59:18  Stub         0.000233
1w2d 13:59:18  Interarea    0.000002
1w2d 13:59:18  External     0.000004
1w2d 13:59:18  NSSA         0.000001
1w2d 13:59:18  Prefix SID   0.000324
1w2d 13:59:18  Adj SID      0.000318
1w2d 13:59:18  Cleanup      0.000551
1w2d 13:59:18  Total        0.002751

```

Meaning

The output displays the FlexAlgo keyword added for the SPF logs.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
22.2R1	Starting in Junos OS and Junos OS Evolved Release 22.2R1, you can advertise different te-attributes such as te-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link.
22.4R1	Starting in Junos OS Release 22.4R1, we've defined the Flexible Algorithm Prefix Metric (FAPM) to allow optimal end-to-end path for an interarea prefix.
change-completed	
25.4R1	Starting in Junos OS and Junos OS Evolved Release 25.4R1, you can control path selection by configuring the preference for OSPF Flexible Algorithm routes in <code>inetcolor.0</code> and <code>mpls.0</code> routing tables.
25.4R1	Starting in Junos OS and Junos OS Evolved Release 25.4R1, you can control path selection by configuring the preference for OSPF Flexible Algorithm routes in <code>inetcolor.0</code> and <code>mpls.0</code> routing tables.

Configuring Application-Specific Link Attribute on an OSPF Interface

Advertise different te-attributes such as te-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link. This is done using flexible algorithm specific application-specific link attribute as defined in RFC 8920.

To configure application-specific link attribute based flexible algorithm on an OSPF Interface:

1. Create an OSPF area.

```
[edit protocols]
user@host#set protocols ospf area area-id
```

For example:

```
[edit protocols]
user@host#set protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host#set interface interface-name
```

For example:

```
[edit protocols ospf area 0.0.0.0]
user@host#set interface ge-0/0/0.0
```

3. Configure application-specific link attribute on the OSPF interface of the device.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set application-specific
```

4. Specify the attribute group.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set attribute-group name
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set attribute-group asla
```

5. Configure flexible algorithm specific te-attributes such as te-metric, delay-metric, and admin-groups. Specify the te-metric for the attribute group. The te-metric indicates the metric type based on which OSPFv2 calculates the path.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group
asla]
user@host#set te-metric
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set 15
```

6. Specify the admin-group for the attribute group.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group
asla]
user@host#set admin-group
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set green
```

7. Specify delay-metric for the attribute group.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group asla]
user@host#set delay-metric
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set 123123
```

8. In case delay-metric is not configured, specify advertise-interface-delay to fetch the delay values from the interface configuration hierarchy, that is legacy delay values.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group asla]
user@host#set advertise-interface-delay
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific]
user@host#set 123125
```



NOTE: The following configuration can be committed only if all of the following criteria match:

- An application is associated with the attribute group.
- Delay-metric is not configured in the hierarchy.
- Interface level delay configurations are present.

9. Specify the application for the attribute group. In the current implementation, only flexible algorithm can be configured as an application. An attribute group can have more than one applications associated with it and it equates to a single application-specific link attribute with the

application bits set in the standard application identifier bit mask field of the application-specific link attribute sub.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group
asla]
user@host#set application application-name
```

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0 application-specific attribute-group
asla]
user@host#set application flex-algorithm
```

10. Enter `commit` from the configuration mode.
11. Specify `strict-asla-based-flex-algorithm` to mandate that flexible algorithm path computations use only the links which advertise relevant te-attributes through application-specific link attribute.

```
[edit protocols ospf source-packet-routing]
user@host#set strict-asla-based-flex-algorithm
```

12. Enter `commit` from the configuration mode.
To verify your configuration results, use the `show protocols operational` command.

```
ospf {
  area 0.0.0.0
    interface ge-0/0/0.0 {
      application-specific {
        attribute-group asla {
          te-metric 15;
          admin-group green;
          delay-metric 123123;
          advertise-interface-delay;
          application flex-algorithm;
        }
      }
    }
  source-packet-routing {
    strict-asla-based-flex-algorithm;
  }
}
```

```
}

```

The Junos OS and Junos OS Evolved implementation supports application-specific link attribute subTLV to comply with RFC 8920. The application-specific link attribute sub-TLV is a sub-TLV of the OSPFv2 extended Link TLV as defined in RFC 7684.

To verify the presence of application-specific link attribute sub-TLVs in the OSPF database use the `show ospf database extensive operational` command.

```
user@host> show ospf database advertising-router self extensive lsa-id 10.0.0.2

```

```
OSPF database, Area 0.0.0.0
  Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
OpaqArea*10.0.0.2      100.100.100.100 0x80000007    665  0x22 0x649d 104
Opaque LSA
  Extended Link (1), length 80:
    Link Type (1), length 1:
      1
    Link Id (2), length 4:
      10.1.1.1
    Link Data (3), length 4:
      10.21.1.1
  Adjacency Sid (2), length 7:
    Flags (1), length 1:
      0x60
    MT ID (2), length 1:
      0
    Weight (3), length 1:
      0
    Label (4), length 3:
      17
  Application Specific Link Attribute (10), length 52:
    SABM Length (1), length 1:
      4
    UDABM Length (2), length 1:
      0
    SABM (3), length 4:
      0x10
    UDABM (4), length 0:
      0x0

```

```

TEMetric (5), length 4:
  10
UnidirecLinkDelay (27), length 4:
  123
MinMaxUnidirecLinkDelay (28), length 8:
  Min DM: 123, Max DM: 123
UnidirecLinkDelayVar (29), length 4:
  0
Color (9), length 4:
  2
Gen timer 00:34:55
Aging timer 00:48:55
Installed 00:11:05 ago, expires in 00:48:55, sent 00:11:05 ago
Last changed 00:11:05 ago, Change count: 6, Ours, TE Link ID: 0

```

The output displays application-specific link attribute sub-TLV fields and attributes.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
22.2R1	Starting in Junos OS and Junos OS Evolved Release 22.2R1, you can advertise different te-attributes such as te-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link. This is done using flexible algorithm specific application-specific link attribute as defined in RFC 8920.

RELATED DOCUMENTATION

| *application-specific (Protocols OSPF)*

How to Enable Link Delay Measurement and Advertising in OSPF

IN THIS SECTION

- [Understanding Link Delay Measurement and Advertising in OSPF | 546](#)
- [Configuring OSPF Link Delay and Delay Normalization on an OSPF Interface | 548](#)

Understanding Link Delay Measurement and Advertising in OSPF

IN THIS SECTION

- [Benefits of link delay measurement and advertising in OSPF | 546](#)
- [Overview of link delay measurement and advertising in OSPF | 547](#)
- [Overview of Link Delay Normalization | 547](#)

Benefits of link delay measurement and advertising in OSPF

Link delay measurement and advertising in OSPF provides the following benefits:


- Highly beneficial in certain networks such as stock market data providers, where it is crucial to have access to market data in real-time to make trades faster than the competition. This is where network performance criteria or latency is becoming critical to data-path selection.
- Helps to make path-selection decisions based on performance data (such as latency) in a cost-effective and scalable way.
- Superior alternative to using metrics such as hop count or cost as routing metrics.

Overview of link delay measurement and advertising in OSPF

Network performance is measured by using TWAMP -Light. You can get the measurement of various performance metrics in IP networks, by using probe messages. OSPF Traffic Engineering Extensions helps to distribute network-performance information in a scalable fashion. This information can then be used to make path-selection decisions based on network performance.

Border Gateway Protocol Link-State (BGP-LS) allows BGP to carry link-state information acquired from IGP, which then allows internet service providers (ISP) to selectively expose the information with other ISPs, service providers, CDNs and so on, through normal BGP peering. New BGP-Link State (BGP-LS) TLVs are defined to carry the IGP Traffic Engineering Metric Extensions.

The following illustration depicts the minimum IGP metric and minimum delay metric in networks that consist a core, metro, and access network.



In this scenario, core network is cheaper but has longer delay. Access shortcut, with lowest latency is expensive. As core network is cheaper, majority of traffic typically go from 1>2>3>4>5> to 6 by using minimum IGP requirement by running OSPF with appropriate cost configured and default OSPF algorithm set to zero. In businesses where ultra-low latency is crucial, packets need to go from 1 to 6. As displayed in scenario b), you can achieve minimum delay metric by defining OSPF flex algorithm with minimum latency, which minimize the delay to the endpoint. This flex algorithm consists only node 1 and node 6.

Overview of Link Delay Normalization

The Delay Normalization feature computes a normalized delay value and uses the normalized value instead of measured delay value. This value is advertised and used as a metric during the flexible algorithm computation. The normalization is performed when the delay is received from the delay measurement component. When the next value is received, it is normalized and compared to the previously saved normalized value. If the values are different, then the LSP generation is triggered.

Delay normalization is disabled by default. To enable and configure delay normalization across IGP instances, use the delay-measurement command.

```
set protocols ospf area area interface interface-name delay-measurement normalize interval  
value offset value
```

SEE ALSO

[delay-measurement \(Protocols OSPF\)](#)

[delay-metric \(Protocols OSPF\)](#)

[Understand Two-Way Active Measurement Protocol](#)

Configuring OSPF Link Delay and Delay Normalization on an OSPF Interface

In IP networks, the bulk of traffic often goes through the core network, which reduces costs but might result in increased latency. Business traffic, however, often benefits from the ability to make path-selection decisions based on other performance metrics, such as path latency, rather than relying on the traditional path optimization based simply on IGP metrics. Optimizing a path to reduce latency can greatly benefit applications like real-time voice and video. It can also enable high performance access to financial market data where milliseconds can translate into significant gains or losses.

You can enable OSPF link delay in IP networks. You can achieve minimum IGP metric paths by configuring OSPF with the appropriate link cost using the default OSPF algorithm. Doing so optimizes paths to the endpoint that are based strictly on the sum of the link metrics. By using the OSPF delay flex algorithm you can optimize paths based on their end-to-end delay.

Link delay can be dynamically measured using Two-Way Active Measurement Probes (TWAMP). The routers then flood their link delay parameters. The routers in the area store these parameters in the shared Link State Database (LSDB). Ingress nodes run an SPF algorithm against the LSDB to compute paths that are optimized on various attributes, such as link colors, IGP metric, traffic-engineering (TE) metric.

To configure link delay measurement for an OSPF Interface:

1. Create an OSPF area.

```
[edit protocols]
user@host#set protocols ospf area area-id
```

For example:

```
[edit protocols]
user@host#set protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host#set interface interface-name
```

For example:

```
[edit protocols ospf area 0.0.0.0]
user@host#set interface ge-0/0/0.0
```

3. Configure dynamic OSPF link delay-measurement on the OSPF interface of the device.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement
```

4. Configure the delay-measurement advertisement on the OSPF interface of the device. You can either configure accelerated or periodic advertisement.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement advertisement (accelerated | periodic)
```



NOTE: Accelerated advertisement is disabled by default. To configure accelerated advertisement, configure the threshold percentage.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement advertisement accelerated threshold percentage
```

5. To configure periodic advertisement, you can either configure interval or the threshold percentage.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement advertisement periodic (interval interval seconds |
threshold threshold percentage)
```

For example: To configure periodic interval:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement advertisement periodic interval 35
```

For example: To configure periodic threshold:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement advertisement periodic threshold 100)
```

6. (Optional) Specify the probe count.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement probe-count seconds
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement probe-count 10
```

7. (Optional) Specify the probe interval.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement probe-interval seconds
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement probe-interval 100
```

8. (Optional) Specify the normalize interval and offset values.

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement normalize interval microseconds offset microseconds seconds
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-measurement normalize interval 50 offset 10
```

9. Enter `commit` from the configuration mode.

To configure delay metric for an OSPF Interface:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-metric microseconds
```

For example:

```
[edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0]
user@host#set delay-metric 20000
```

Enter `commit` from the configuration mode.

To verify your configuration results, use the `show protocols operational` command.

```
user@host# show protocols
```

```
ospf {
  area 0.0.0.0
    interface ge-0/0/0.0 {
      delay-measurement {
      advertisement {
        accelerated {
          threshold 100;
        }
      }
      periodic {
        interval 35;
        threshold 100;
      }
    }
  }
}
```

```

}
  probe-count 10;
  probe-interval 100;
}
      delay-metric {
        20000;
      )
}
  normalize interval 50 offset 10;

```

To verify that link delay parameters are present in the OSPF database use the `show ospf database extensive | match delay operational` command.

```
user@host> show ospf database extensive | match delay
```

```

Unidirectional link delay: 20000
Min unidirectional link delay: 20000
Max unidirectional link delay: 20000
Unidirectional delay variation: 20000

```

The output displays the delay of 20000 microseconds that is configured on the interface.

To verify the normalized delay value, use the `show ospf interface extensive operational` command.

```
user@host> show ospf interface extensive
```

```

root@R0# run show ospf interface ge-0/0/0.0 extensive
Interface      State  Area      DR ID      BDR ID      Nbrs
ge-0/0/0.0     PtToPt 0.0.0.0   0.0.0.0   0.0.0.0     1
Type: P2P, Address: 21.0.1.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None
Topology default (ID 0) -> Cost: 1

```

```

Unidirectional link delay: 1616
Min unidirectional link delay: 712
Max unidirectional link delay: 8125
Unidirectional delay variation: 2992
Delay Normalization: Enabled (Interval: 50, Offset: 10)
Normalized Delay: 860 (Last Normalized: Mon Jul 24 10:03:41)

```

The output displays that the unidirectional link delay 1616 microseconds has been normalized into a value (860) that OSPF uses in SPF computations for flexible algorithm. Last Normalized: Mon Jul 24 10:03:41 indicates the timestamp of the most recent normalization calculation.

SEE ALSO

delay-measurement (Protocols OSPF)

delay-metric (Protocols OSPF)

[Understand Two-Way Active Measurement Protocol](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
25.4R1	Starting in Junos OS and Junos OS Evolved Release 25.4R1, you can use delay normalization to compute and advertise a normalized delay metric for flexible algorithm, to improve path-selection consistency across all IGP instances.
21.4R1	Starting in Junos OS Release 21.4R1, you can get the measurement of various performance metrics in IP networks, by using probe messages.

How to Configure Microloop Avoidance in OSPFv2 Segment Routing Networks

SUMMARY

Microloops can consume the available bandwidth of the links, which impacts the efficient transmission of useful packets. Microloop avoidance can prevent forwarding of looping packets.

IN THIS SECTION

- [Understanding OSPF Microloop Avoidance | 554](#)
- [Configuring Segment Routing Microloop Avoidance in OSPFv2 Networks | 557](#)

Understanding OSPF Microloop Avoidance

IN THIS SECTION

- [Benefits of Avoiding Microloops in OSPFv2 Networks with Segment Routing | 554](#)
- [Microloop Avoidance in OSPFv2 Networks with Segment Routing | 555](#)
- [Supported and Unsupported Features | 556](#)


Benefits of Avoiding Microloops in OSPFv2 Networks with Segment Routing

- Micro loop-free path avoids delays and traffic loss.
- Microloop avoidance can prevent forwarding of looping packets and avoid wasteful bandwidth consumption.
- Microloop avoidance path is computed only for the impacted links in case of multiple link failures. If the second link failure does not impact the computed microloop avoidance path, OSPFv2 continues to use the same microloop avoidance path.

Junos OS enables a device to defer OSPFv2 route download when an OSPFv2 link fails in order to avoid micro loops. When local links go down, the OSPFv2 protocol floods an entire area with the database. If the node connected to the local interface that has failed converges faster than the neighboring node, then the connected node redirects traffic to the converged path. This redirection can result in micro

looping of traffic until the neighboring node converges. When the primary path of a protected node fails, the connected node does not need to converge quickly if the configured backup path is not impacted. In this case, traffic flow towards a converged path is deferred until the configured delay time. This time delay helps in avoiding microloops because all routers do not arrive at the post-convergence forwarding states simultaneously.

Figure 34: Microloop Avoidance in OSPFv2 Networks



In the [Figure 34 on page 555](#), the primary path from Source to Destination is S→R0→R1→R2→R3→D. When the link between R2 and R3 fails, traffic sent from S to D, is subject to transient forwarding loops while routers update their forwarding state for destination D.

- If R0 updates its forwarding state before R5, packets loop between R0 and R5
- If both R0 and R5, have updated their forwarding states, and R4 has not, packets loop between R4 and R5.
- R0 detects the link failure between R2 and R3, and temporarily steers traffic destined to Destination over SR path [NodeSID(R4), AdjSID(R4->R3), D].
- When the configured timeout elapses, R0 just uses the node-SID to D to reach the destination.

Microloop Avoidance in OSPFv2 Networks with Segment Routing

Enable a post convergence path calculation on a device to avoid microloops if a link or metric change occurs in an OSPFv2 segment routed network. To configure microloop avoidance in an OSPFv2 segment routing network for both local and remote network events including link down, link-up, and metric-change, include the `maximum-labelsdelay milliseconds` statement at the `[edit protocols ospf spf-options microloop avoidance post-convergence-path]` hierarchy level. For effective microloop avoidance, configure this feature on all the nodes in the network.



NOTE: Micro-loop avoidance is not a replacement for local repair mechanisms like TI-LFA which detects local failure very fast and activates a pre-computed loop-free-alternative path.

Routers that implement micro-loop avoidance compute the micro-loop avoiding path only after receiving the link state update for the event. So, micro-loop avoidance mechanism is not a replacement for local repair mechanisms like TI-LFA which detects local failure very fast and activates a pre-computed loop-free-alternative path at PFE level. In the above example, if local repair mechanism is not present for the R2↔R3 failure, there will be a lot of traffic loss before R0 can detect the failure (through global convergence) and program a micro-loop avoiding path. Micro-loop avoidance cannot avoid traffic loss due to delayed detection of the failure. Microloop avoidance avoids traffic loss due to micro-loops only. Both local-repair mechanisms like TI-LFA and micro-loop avoidance, have to be enabled on all the nodes in the network to ensure that traffic loss is in milli-seconds range.

To avoid micro-loops, the following process is used:

1. After computing the new path to D, for a predetermined time, R installs an entry for D that steers packets to D through a loop-free segment routed path. This time should be greater than worst case delay of any router in the network.
2. After the configured time delay, R installs the post-convergence route entry for D, which is without any SIDs.

Supported and Unsupported Features

Junos OS supports microloop avoidance in the following scenarios:

- Microloop avoidance is supported on all the Junos OS platforms that support OSPF routing protocol.
- Microloop avoidance is supported for IPv4 networks only.
- Microloop avoidance is supported for flexible algorithm topologies.

Junos OS does not support the following features in conjunction with microloop avoidance:

- Microloop avoidance path that needs more than 8 labels is not supported. The maximum number of labels installed for microloop avoidance path is 8. For the microloop avoidance ECMP path to be usable, the number of labels must be less than or equal to maximum labels.
- Cannot prevent traffic loss because of slow control plane convergence.
- OSPFv2 multi-topology is not supported with microloop avoidance.
- Adjacency SIDs are not supported with microloop avoidance.
- If shortcuts are available OSPFv2 does not provide a microloop avoidance path.

Configuring Segment Routing Microloop Avoidance in OSPFv2 Networks

IN THIS SECTION

- [Overview | 557](#)
- [Requirements | 557](#)
- [Topology | 557](#)
- [Configuration | 558](#)
- [Verification | 577](#)

Overview

Microloops are packet forwarding loops that occur in the network following network change events such as link down, link up, or metric change. When a network change event occurs, different routers update their forwarding states at different times. This can lead to packets getting looped between upstream and downstream routers for a transient period, resulting in packet loss, jitter, and out-of-order packets. Microloops can consume the available bandwidth of the links, which impacts the efficient transmission of useful packets.

Microloop avoidance can prevent forwarding of looping packets. The segment routing microloop avoidance detects if microloops are possible following a topology change. When a network change event is detected, the routes are programmed to take the post-convergence path, that uses a combination of node and adjacency SIDs. This ensures the routers that might not yet have converged do not loop the packets causing microloops. This behavior lasts for a configurable delay. Once the delay timer expires, routes are programmed normally by using node-SID of the destinations.

Requirements

This example uses the following hardware and software components:


- Eight MX Series routers.
- Junos OS Release 22.1R1 or later.

Topology

In [Figure 35 on page 558](#) device R0 and device R7 are the ingress and egress routers that support devices CE1 and CE2. The devices R1, R2, R3, R4, R5, and R6 comprise an IPv4 only provider core network. All the devices belong to the same autonomous system. OSPFv2 is the interior gateway protocol in the core configured to support microloop avoidance. In this example the device R2 is

configured as an IPv4 route reflector with IBGP peering sessions to both R0 and R7. No other routers speak BGP in this example. The Device R6 has the firewall filter configured to detect packets with microloops if any following a link down event.

Figure 35: Microloop Avoidance Topology



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 559](#)
- [Configuring Device R0 | 570](#)
- [Step-by-Step Procedure | 570](#)
- [Results | 574](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R0

```
set interfaces xe-0/0/0:0 description To_R1
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.1.1/30
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:3 description To_R4
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.4.1/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces xe-0/0/1:2 description to_CE1
set interfaces xe-0/0/1:2 unit 1 family inet address 172.16.10.2/30
set interfaces xe-0/0/1:2 unit 1 family mpls
set interfaces xe-0/0/1:2 unit 4 family inet address 172.16.11.2/30
set interfaces xe-0/0/1:2 unit 4 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.10/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.10/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1000
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.10
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols bgp group to-RR type internal
set protocols bgp group to-RR local-address 192.168.255.10
set protocols bgp group to-RR neighbor 192.168.255.12 family inet unicast
set protocols bgp group to-RR neighbor 192.168.255.12 family inet-vpn unicast per-prefix-label
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
```

```

set protocols ospf source-packet-routing node-segment ipv4-index 0
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection

```

Device R1

```

set interfaces xe-0/0/0:0 description To_R0
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.1.2/30
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:2 description To_R2
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.12.1/30
set interfaces xe-0/0/0:2 unit 0 family mpls
set interfaces xe-0/0/0:2 unit 1 family inet address 10.11.12.1/30
set interfaces xe-0/0/0:2 unit 1 family mpls
set interfaces xe-0/0/0:3 description to_R4
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.14.1/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces xe-0/0/1:0 description to_R5
set interfaces xe-0/0/1:0 unit 0 family inet address 10.10.15.1/30
set interfaces xe-0/0/1:0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.11/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.11/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1001
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.11
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable

```

```

set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 2
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 post-convergence-lfa
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.1 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.1 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.1 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.1 metric 10

```

Device R2

```

set interfaces xe-0/0/0:1 description To_R1
set interfaces xe-0/0/0:1 unit 0 family inet address 10.10.12.2/30
set interfaces xe-0/0/0:1 unit 0 family mpls
set interfaces xe-0/0/0:1 unit 1 family inet address 10.11.12.2/30
set interfaces xe-0/0/0:1 unit 1 family inet6
set interfaces xe-0/0/0:1 unit 1 family mpls
set interfaces xe-0/0/0:2 description To_R3
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.23.1/30
set interfaces xe-0/0/0:2 unit 0 family mpls
set interfaces xe-0/0/0:3 description To_R4
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.24.1/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces xe-0/0/1:0 description To_R5

```

```
set interfaces xe-0/0/1:0 unit 0 family inet address 10.10.25.1/30
set interfaces xe-0/0/1:0 unit 0 family mpls
set interfaces xe-0/0/1:1 description To_R6
set interfaces xe-0/0/1:1 unit 0 family inet address 10.10.26.1/30
set interfaces xe-0/0/1:1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.12/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.12/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1002
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.12
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols bgp group to-RR type internal
set protocols bgp group to-RR local-address 192.168.255.12
set protocols bgp group to-RR neighbor 192.168.255.17 family inet unicast
set protocols bgp cluster 192.168.255.12
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 4
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 10
```

```

set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.1 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.1 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.2 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.3 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.3 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.4 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.4 metric 10

```

Device R3

```

set interfaces xe-0/0/0:0 description To_R2
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.23.2/30
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:1 description To_R5
set interfaces xe-0/0/0:1 unit 0 family inet address 10.10.35.2/30
set interfaces xe-0/0/0:1 unit 0 family mpls
set interfaces xe-0/0/0:2 description To_R6
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.36.1/30
set interfaces xe-0/0/0:2 unit 0 family mpls
set interfaces xe-0/0/0:3 description To_R7
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.37.1/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.13/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.13/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1003
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.13
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable

```

```

set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 6
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection

```

Device R4

```

set interfaces xe-0/0/0:0 description To_R0
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.4.2/30
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:2 description To_R2
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.24.2/30
set interfaces xe-0/0/0:2 unit 0 family mpls
set interfaces xe-0/0/0:3 description To_R1
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.14.2/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces xe-0/0/1:0 description To_R5
set interfaces xe-0/0/1:0 unit 0 family inet address 10.10.45.1/30
set interfaces xe-0/0/1:0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.14/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.14/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1004

```

```

set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.14
set routing-options forwarding-table export pplb
set routing-options autonomous-system 65000
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 8
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection

```

Device R5

```

set interfaces xe-0/0/0:1 description To_R1
set interfaces xe-0/0/0:1 unit 0 family inet address 10.10.15.2/30
set interfaces xe-0/0/0:1 unit 0 family mpls
set interfaces xe-0/0/0:2 description To_R2
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.25.2/30
set interfaces xe-0/0/0:2 unit 0 family mpls
set interfaces xe-0/0/0:3 description To_R3
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.35.2/30

```

```
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces xe-0/0/1:0 description To_R4
set interfaces xe-0/0/1:0 unit 0 family inet address 10.10.45.2/30
set interfaces xe-0/0/1:0 unit 0 family mpls
set interfaces xe-0/0/1:1 description To_R6
set interfaces xe-0/0/1:1 unit 0 family inet address 10.10.56.1/30
set interfaces xe-0/0/1:1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.15/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.15/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1005
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set routing-options router-id 192.168.255.15
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing node-segment ipv4-index 10
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/1:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 interface-type p2p
```

```

set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/1:1.0 post-convergence-lfa node-protection

```

Device R6

```

set interfaces xe-0/0/0:0 description To_R2
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.26.2/30
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:1 description To_R3
set interfaces xe-0/0/0:1 unit 0 family inet address 10.10.36.2/30
set interfaces xe-0/0/0:1 unit 0 family mpls
set interfaces xe-0/0/0:2 description To_R5
set interfaces xe-0/0/0:2 unit 0 family inet filter output v4filter
set interfaces xe-0/0/0:2 unit 0 family inet address 10.10.56.2/30
set interfaces xe-0/0/0:2 unit 0 family mpls filter output mplsfilter
set interfaces xe-0/0/0:3 description To_R7
set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.67.1/30
set interfaces xe-0/0/0:3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.16/32
set interfaces lo0 unit 0 family inet address 192.168.255.61/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.16/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1006
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set policy-options policy-statement prefix-sid term 2 from route-filter 192.168.255.61/32 exact
set policy-options policy-statement prefix-sid term 2 then prefix-segment index 1106
set policy-options policy-statement prefix-sid term 2 then accept
set firewall family inet filter v4filter term t1 from destination-address 8.3.0.0/16
set firewall family inet filter v4filter term t1 then accept
set firewall family inet filter v4filter term t6 then accept
set firewall family mpls filter mplsfilter term t1 from ip-version ipv4 destination-address
10.8.0.1/16
set firewall family mpls filter mplsfilter term t1 then count v4sr-nsid-cnt
set firewall family mpls filter mplsfilter term t1 then accept
set firewall family mpls filter mplsfilter term t2 from ip-version ipv4 destination-address
10.9.0.1/16
set firewall family mpls filter mplsfilter term t2 then count v4sr-psid-cnt
set firewall family mpls filter mplsfilter term t2 then accept
set firewall family mpls filter mplsfilter term t3 then accept

```

```

set firewall family mpls filter mplsfilter term t4 then accept
set firewall family mpls filter mplsfilter term t6 then accept
set routing-options router-id 192.168.255.16
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf spf-options microloop-avoidance post-convergence-path delay 60000
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 12
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:2.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 metric 110
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 100
set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-protection

```

Device R7

```

set interfaces xe-0/0/0:0 description To_R3
set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.37.2/24
set interfaces xe-0/0/0:0 unit 0 family mpls
set interfaces xe-0/0/0:1 description To_R6
set interfaces xe-0/0/0:1 unit 0 family inet address 10.10.67.2/30
set interfaces xe-0/0/0:1 unit 0 family mpls
set interfaces xe-0/0/0:2 description to_CE2
set interfaces xe-0/0/0:2 unit 4 family inet address 172.16.20.1/30

```

```
set interfaces xe-0/0/0:2 unit 4 family mpls
set interfaces lo0 unit 0 family inet address 192.168.255.17/32
set interfaces lo0 unit 0 family inet address 192.168.255.71/32
set interfaces lo0 unit 0 family mpls
set policy-options policy-statement payload_9 term 1 from route-filter 10.7.0.1/16 orlonger
set policy-options policy-statement payload_9 term 1 then next-hop 192.168.255.17
set policy-options policy-statement payload_9 term 1 then accept
set policy-options policy-statement payload_9 term 2 from route-filter 10.8.0.1/16 orlonger
set policy-options policy-statement payload_9 term 2 then next-hop 192.168.255.17
set policy-options policy-statement payload_9 term 2 then accept
set policy-options policy-statement payload_9 term 3 from route-filter 8.2.0.0/16 orlonger
set policy-options policy-statement payload_9 term 3 then next-hop 192.168.255.71
set policy-options policy-statement payload_9 term 4 then reject
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement prefix-sid term 1 from route-filter 192.168.255.17/32 exact
set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1007
set policy-options policy-statement prefix-sid term 1 then prefix-segment node-segment
set policy-options policy-statement prefix-sid term 1 then accept
set policy-options policy-statement prefix-sid term 2 from route-filter 192.168.255.71/32 exact
set policy-options policy-statement prefix-sid term 2 then prefix-segment index 1107
set policy-options policy-statement prefix-sid term 2 then accept
set policy-options policy-statement v4stat term 1 from protocol static
set policy-options policy-statement v4stat term 1 from route-filter 100.100.100.1/32 orlonger
set policy-options policy-statement v4stat term 1 then accept
set policy-options policy-statement v4_prefixes term 1 from route-filter 8.3.0.0/16 orlonger
set policy-options policy-statement v4_prefixes term 1 then accept
set policy-options policy-statement v4_prefixes term 3 then reject
set routing-options rib inet.0 static route 100.100.100.1/32 receive
set routing-options router-id 192.168.255.17
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols bgp group to-RR type internal
set protocols bgp group to-RR local-address 192.168.255.17
set protocols bgp group to-RR neighbor 192.168.255.12 family inet unicast
set protocols bgp group to-RR neighbor 192.168.255.12 export payload_9
set protocols bgp group to-CE1 type external
set protocols bgp group to-CE1 local-address 172.16.20.1
set protocols bgp group to-CE1 neighbor 172.16.20.2 family inet unicast
set protocols bgp group to-CE1 neighbor 172.16.20.2 peer-as 700
set protocols bgp group to-CE1 neighbor 172.16.20.2 local-as 100
set protocols mpls traffic-engineering
set protocols mpls label-range static-label-range 60001 100000
set protocols mpls interface all
```

```

set protocols mpls interface fxp0.0 disable
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-paths 8
set protocols ospf backup-spf-options use-source-packet-routing
set protocols ospf source-packet-routing prefix-segment prefix-sid
set protocols ospf source-packet-routing node-segment ipv4-index 14
set protocols ospf source-packet-routing srgb start-label 800000
set protocols ospf source-packet-routing srgb index-range 80000
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-protection
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/0/0:1.0 post-convergence-lfa node-protection

```

Configuring Device R0

Step-by-Step Procedure

To configure segment routing microloop avoidance path in an OSPFv2 network, perform the following steps on the R0 device:

1. Configure the device interfaces to enable IP and MPLS transport.

```

[edit]
user@R0#set interfaces xe-0/0/0:0 description To_R1
user@R0#set interfaces xe-0/0/0:0 unit 0 family inet address 10.10.1.1/30
uesr@R0#set interfaces xe-0/0/0:0 unit 0 family mpls
user@R0#set interfaces xe-0/0/0:3 description To_R4
user@R0#set interfaces xe-0/0/0:3 unit 0 family inet address 10.10.4.1/30
uesr@R0#set interfaces xe-0/0/0:3 unit 0 family mpls
user@R0#set interfaces xe-0/0/1:2 description to_CE1
user@R0#set interfaces xe-0/0/1:2 unit 1 family inet address 172.16.10.2/30
user@R0#set interfaces xe-0/0/1:2 unit 1 family mpls

```

2. Configure the loopback interface (lo0) addresses that is used as router ID for OSPF sessions.

```
[edit]
user@R0#set interfaces lo0 unit 0 family inet address 192.168.255.10/32
user@R0#set interfaces lo0 unit 0 family inet address 192.168.255.18/32
```

3. Configure the router ID and autonomous system (AS) number to propagate routing information within a set of routing devices that belong to the same AS.

```
[edit]
user@R0#set routing-options router-id 192.168.255.10
user@R0#set routing-options autonomous-system 65000
```

4. Define a policy to load balance packets and apply the per-packet policy to enable load balancing of traffic.

```
[edit]
user@R0#set policy-options policy-statement pplb then load-balance per-packet
user@R0#set routing-options forwarding-table export pplb
```

5. Configure R0 to advertise the loopback address. The `prefix-segment index` option sets the base label for each router's loopback. In this example the base index is set to reflect the router number. As a result, R0 uses 1000.

```
[edit]
user@R0#set policy-options policy-statement prefix-sid term 1 from route-filter
192.168.255.10/32 exact
user@R0#set policy-options policy-statement prefix-sid term 1 then prefix-segment index 1000
user@R0#set policy-options policy-statement prefix-sid term 1 then prefix-segment node-
segment
user@R0#set policy-options policy-statement prefix-sid term 1 then accept
```

6. Configure MPLS on all interfaces excluding the management interface. Also enable traffic engineering.

```
[edit]
user@R0#set protocols mpls interface all
```

```
user@R0#set protocols mpls interface fxp0.0 disable
user@R0#set protocols mpls traffic-engineering
```

7. Configure the MPLS label range to assign static labels for the links.

```
[edit]
user@R0#set protocols mpls label-range static-label-range 60001 100000
```

8. Configure BGP peering between R0 and the route reflector R2. Configure the unicast network layer reachability information (NRLI) to allocate a unique label for each prefix on the devices.

```
[edit]
user@R0#set protocols bgp group to-RR type internal
user@R0#set protocols bgp group to-RR local-address 192.168.255.10
user@R0#set protocols bgp group to-RR neighbor 192.168.255.12 family inet unicast
user@R0#set protocols bgp group to-RR neighbor 192.168.255.12 family inet-vpn unicast per-
prefix-label
```

9. Configure TI-LFA to enable protection against link and node failures. SR using TI-LFA provides faster restoration of network connectivity by routing the traffic instantly to a backup or an alternate path if the primary path fails or becomes unavailable.

```
[edit]
user@host#set protocols ospf backup-spf-options use-source-packet-routing
```

10. Configure backup shortest path first (SPF) attributes such as maximum equal-cost multipath (ECMP) as 8 and maximum number of labels as 5 for TI-LFA for the OSPFv2 protocol.

```
[edit]
user@host#set protocols ospf backup-spf-options use-post-convergence-lfa maximum-labels 5
user@host#set protocols ospf backup-spf-options use-post-convergence-lfa maximum-backup-
paths 8
```

11. Configure prefix segment attributes, the start label and the index range for segment routing global blocks (SRGBs) in SPRING for the OSPFv2 protocol.

```
[edit]
user@host#set protocols ospf source-packet-routing prefix-segment prefix-sid
```

```

user@host#set protocols ospf source-packet-routing node-segment ipv4-index 0
user@host#set protocols ospf source-packet-routing srgb start-label 800000
user@host#set protocols ospf source-packet-routing srgb index-range 80000

```

12. Configure the loopback interface as passive to ensure the protocols do not run over the loopback interface and that the loopback interface is advertised correctly throughout the network.

```

[edit]
user@host#set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

13. Configure OSPF area 0 on the point-to-point interface of the device R0.

```

[edit]
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 interface-type p2p
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 metric 10
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 interface-type p2p
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 metric 10

```

14. Configure the computation and installation of a backup path that follows the post-convergence path on the given area and interface for the OSPFv2 protocol. Also enable node-link protection on the these interfaces that follow post-convergence path.

```

[edit]
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:0.0 post-convergence-lfa node-
protection
user@host#set protocols ospf area 0.0.0.0 interface xe-0/0/0:3.0 post-convergence-lfa node-
protection

```

15. Configure microloop avoidance that temporarily installs a post-convergence path for routes potentially affected by microloops and specify a delay time period of 60000 milliseconds for the OSPFv2 protocol. The temporary path reverts to the node SIDs of the destination after the delay timer expires.

```

[edit]
user@host#set protocols ospf spf-options microloop-avoidance post-convergence-path delay
60000

```

Results

Check the results of the configuration:

```
interfaces {
  xe-0/0/0:0 {
    description To_R1;
    unit 0 {
      family inet {
        address 10.10.1.1/30;
      }
      family mpls;
    }
  }
  xe-0/0/0:3 {
    description To_R4;
    unit 0 {
      family inet {
        address 10.10.4.1/30;
      }
      family mpls;
    }
  }
  xe-0/0/1:2 {
    description to_CE1;
    unit 1 {
      family inet {
        address 172.16.10.2/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.255.10/32;
        address 192.168.255.18/32;
      }
      family mpls;
    }
  }
}
```

```
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
  policy-statement prefix-sid {
    term 1 {
      from {
        route-filter 192.168.255.10/32 exact;
      }
      then {
        prefix-segment {
          index 1000;
          node-segment;
        }
        accept;
      }
    }
    term 2 {
      from {
        route-filter 192.168.255.18/32 exact;
      }
      then {
        prefix-segment {
          index 1100;
        }
        accept;
      }
    }
  }
}
routing-options {
  router-id 192.168.255.10;
  autonomous-system 100;
  forwarding-table {
    export pplb;
  }
}
protocols {
  bgp {
    group to-RR {
```

```
    type internal;
    local-address 192.168.255.10;
    neighbor 192.168.255.12 {
        family inet {
            unicast;
        }
        family inet-vpn {
            unicast {
                per-prefix-label;
            }
        }
    }
}
mpls {
    traffic-engineering;
    label-range {
        static-label-range 60001 100000;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
ospf {
    spf-options {
        microloop-avoidance {
            post-convergence-path {
                delay 60000;
            }
        }
    }
    backup-spf-options {
        use-post-convergence-lfa {
            maximum-labels 5;
            maximum-backup-paths 8;
        }
        use-source-packet-routing;
    }
    source-packet-routing {
        prefix-segment prefix-sid;
        node-segment ipv4-index 0;
        srgb start-label 800000 index-range 80000;
```

```
}
area 0.0.0.0 {
  interface lo0.0 {
    passive;
  }
  interface xe-0/0/0:0.0 {
    interface-type p2p;
    metric 10;
    post-convergence-lfa;
  }
  interface xe-0/0/0:3.0 {
    interface-type p2p;
    metric 10;
    post-convergence-lfa;
  }
}
}
```

Verification

IN THIS SECTION

- [Verify Connectivity Between R0 and R7 Before the Link is Disabled Between R0 and R1 | 578](#)
- [Verify Disabling the Link Between R0 and R1 | 578](#)
- [Verify Microloop-avoidance Path Installed for the Destination After the Link is Disabled | 579](#)
- [Verify Packets With Microloops | 581](#)
- [Verify Microloop-avoidance Path Changes to Post-convergence- path After the Delay Timer Expires | 581](#)
- [Verify Connectivity Between R0 and R7 | 583](#)
- [Verify the Path Changes to Microloop-avoidance Path After the Link is Enabled | 583](#)

Confirm that the configuration is working properly.

The following section explains microloop avoidance for a link down event.

Verify Connectivity Between R0 and R7 Before the Link is Disabled Between R0 and R1

Purpose

Verify that the Device R0 can reach the destinations on Device R7.

Action

From operational mode, run the **ping** command on the device R0.

```
user@R0>ping 192.168.255.17
PING 192.168.255.17 (192.168.255.17): 56 data bytes
64 bytes from 192.168.255.17: icmp_seq=0 ttl=61 time=41.493 ms
64 bytes from 192.168.255.17: icmp_seq=1 ttl=61 time=57.242 ms
64 bytes from 192.168.255.17: icmp_seq=2 ttl=61 time=44.977 ms
64 bytes from 192.168.255.17: icmp_seq=3 ttl=61 time=202.092 ms
64 bytes from 192.168.255.17: icmp_seq=4 ttl=61 time=60.495 ms
64 bytes from 192.168.255.17: icmp_seq=5 ttl=61 time=39.396 ms
64 bytes from 192.168.255.17: icmp_seq=6 ttl=61 time=79.993 ms
64 bytes from 192.168.255.17: icmp_seq=7 ttl=61 time=78.741 ms
8 packets transmitted, 8 received, 0% packet loss, time 7007ms
rtt min/avg/max/mdev = 38.194/47.998/60.879/8.727 ms
```

Meaning

These results confirm that the device R0 can reach device R7 in the OSPFv2 network.

Verify Disabling the Link Between R0 and R1

Purpose

To verify disabling the link between R0 and R1 on the device R0

Action

From configuration mode, run the **disable interface** command on the device R0

```
user@R0#disble interface xe-0/0/0:0
```

To verify the link is disabled, from operational mode, run the **show interfaces** command on the device R0

```

user@R0>show interfaces xe-0/0/0:0
Physical interface: xe-0/0/0:0, Administratively down, Physical link is Down
  Interface index: 149, SNMP ifIndex: 527
  Description: To_R1_1
  Link-level type: Ethernet, MTU: 1518, MRU: 1526, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering:
Disabled, Flow control: Enabled, Speed Configuration: Auto
  Pad to minimum frame size: Disabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Down SNMP-Traps Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 2c:6b:f5:42:fe:00, Hardware address: 2c:6b:f5:42:fe:00
  Last flapped   : 2022-02-15 09:53:51 PST (00:00:10 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics
    Bit errors           Seconds
    Bit errors           0
    Errored blocks       0
  Link Degrade :
    Link Monitoring      : Disable
  Interface transmit statistics: Disabled

```

Meaning

The output indicates the physical link between R0 and R1 is disabled and is administratively down.

Verify Microloop-avoidance Path Installed for the Destination After the Link is Disabled

Purpose

Verify microloop-avoidance path installed for the destination routes R7 from R0 when the link is disabled between R0 and R1 by verifying routes in the inet.3 table and route label details in the mpls.0 table.

Action

From operational mode, run the **show route table inet.3** command on the device R0.

```
user@R0>show route table inet.3 192.168.255.17/32
inet.3: 25 destinations, 26 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.255.17/32          *[L-OSPF/10/5] 00:00:31, metric 130
                        > to 192.168.255.14 via xe-0/0/0:3, Push 16, Push 801006(top)
```

From operational mode, run the **show route label *label value* protocol ospf extensive** command on the device R0.

```
user@R0>show route label 801007 protocol ospf extensive
mpls.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
801007 (1 entry, 1 announced)
TSI:
KRT in-kernel 801007 /52 -> {Swap 16, Push 801006 (top)}
  *L-OSPF Preference: 10/5
    Next hop type: Router, Next hop index: 649
    Address: 0x7a1ed58
    Next-hop reference count: 4, key opaque handle: 0x0
    Next hop: 10.10.4.2 via xe-0/0/0:3.0 weight 0x1, selected
    Label operation: Swap 16, Push 801006(top)
    Load balance label: Label 16: None; Label 801006: None
    Label element ptr: 0x8fd6ed0
    Label parent element ptr: 0x0
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 321
    State: <Active Int>
    Local AS: 100
    Age: 2:55:13 Metric: 130
    Validation State: unverified
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (1): 1-KRT
```

```
AS path: I
Thread: junos-main
```

Meaning

The output indicates that when the link between R0 and R1 goes down, the microloop-avoidance path is installed for R7 from R0 through R4 until the delay timer expires.

Verify Packets With Microloops

Purpose

Verify packets with microloops by using firewall counter information

Action

From operational mode, run the **show firewall** command on the device R6.

```
user@R6>show firewall
Filter: mplsfilter
Counters:
Name                               Bytes      Packets
v4sr-nsid-cnt                       0           0
v4sr-psid-cnt                       0           0
```

Meaning

The output displays the mplsfilter configured on the device R6 to display microloops if there are any. The value 0 indicates there are no packets with microloops.

Verify Microloop-avoidance Path Changes to Post-convergence- path After the Delay Timer Expires

Purpose

Verify microloop-avoidance path installed for the destination routes R7 from R0 changes to post-convergence-path after the delay timer 60000 ms expires.

Action

From operational mode, run the **show route table inet.3** command on the device R0.

```
user@R0>show route table inet.3 192.168.255.17/32
inet.3: 25 destinations, 26 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.255.17/32          *[L-OSPF/10/5] 00:00:31, metric 130
                        > to 192.168.255.14 via xe-0/0/0:3, Push 801007
```

From operational mode, run the **show route label *label value* protocol ospf extensive** command on the device R0.

```
user@R0>show route label 801007 protocol ospf extensive
mpls.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
801007 (1 entry, 1 announced)
TSI:
KRT in-kernel 801007 /52 -> {Swap 801007}
  *L-OSPF Preference: 10/5
    Next hop type: Router, Next hop index: 615
    Address: 0x7a1c400
    Next-hop reference count: 4, key opaque handle: 0x0
    Next hop: 10.10.4.2 via xe-0/0/0:3.0 weight 0x1, selected
Label operation: Swap 801007
    Load balance label: Label 801007: None;
    Label element ptr: 0x8fd6458
    Label parent element ptr: 0x0
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 321
    State: <Active Int>
    Local AS: 100
    Age: 2:55:13 Metric: 130
    Validation State: unverified
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (1): 1-KRT
    AS path: I
    Thread: junos-main
```

Meaning

The output indicates that the microloop-avoidance path is changed to post-convergence-path after the delay timer expires.

Verify Connectivity Between R0 and R7

Purpose

Verify that the Device R0 can reach the destinations on Device R7.

Action

From operational mode, run the **ping** command on the device R0.

```
user@R0>ping 192.168.255.17
PING 192.168.255.17 (192.168.255.17): 56 data bytes
64 bytes from 192.168.255.17: icmp_seq=0 ttl=61 time=41.493 ms
64 bytes from 192.168.255.17: icmp_seq=1 ttl=61 time=57.242 ms
64 bytes from 192.168.255.17: icmp_seq=2 ttl=61 time=44.977 ms
64 bytes from 192.168.255.17: icmp_seq=3 ttl=61 time=202.092 ms
64 bytes from 192.168.255.17: icmp_seq=4 ttl=61 time=60.495 ms
64 bytes from 192.168.255.17: icmp_seq=5 ttl=61 time=39.396 ms
64 bytes from 192.168.255.17: icmp_seq=6 ttl=61 time=79.993 ms
64 bytes from 192.168.255.17: icmp_seq=7 ttl=61 time=78.741 ms
8 packets transmitted, 8 received, 0% packet loss, time 7007ms
rtt min/avg/max/mdev = 38.194/47.998/60.879/8.727 ms
```

Meaning

These results confirm that the device R0 can reach device R7 in the OSPFv2 network and that the traffic flows with 0% packet loss in case of link down because of the microloop-avoidance path configured.

Verify the Path Changes to Microloop-avoidance Path After the Link is Enabled

Purpose

Verify the path changes to microloop-avoidance path for the destination when the link is enabled between R0 and R1.

Action

From operational mode, run the **show route table inet.3** command on the device R0.

```
user@R0>show route table inet.3 192.168.255.17/32
inet.3: 26 destinations, 27 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.255.17/32      *[L-OSPF/10/5] 00:02:05, metric 40
                    > to 192.168.255.11 via xe-0/0/0:0, Push 801007
                    to 192.168.255.14 via xe-0/0/0:3, Push 16, Push 801006(top)
```

From operational mode, run the **show route label *label value* protocol ospf extensive** command on the device R0.

```
user@R0>show route label 801007 protocol ospf extensive
mpls.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
801007 (1 entry, 1 announced)
TSI:
KRT in-kernel 801007 /52 -> {list:Swap 801007, Swap 16, Push 801006(top)}
  *L-OSPF Preference: 10/5
    Next hop type: Router, Next hop index: 615
    Address: 0x79329ac
    Next-hop reference count: 3, key opaque handle: 0x0
    Next hop: 10.10.4.2 via xe-0/0/0:3.0 weight 0x1, selected
    Label operation: Push 801007
    Load balance label: Label 801007: None;
    Label element ptr: 0x8fd6458
    Label parent element ptr: 0x0
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 0
    Next hop: 10.10.1.2 via xe-0/0/0:0.0 weight 0xf000, selected
    Label operation: Swap 16, Push 801006(top)
    Load balance label: Label 16: None; Label 801006: None;
    Label element ptr: 0x8fd8e60
    Label parent element ptr: 0x0
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
```

```
Session Id: 0
State: <Active Int>
Local AS: 100
Age: 2:55:13 Metric: 40
Validation State: unverified
Area: 0.0.0.0
Task: OSPF
Announcement bits (1): 1-KRT
AS path: I
Thread: junos-main
```

Meaning

The output displays the routes to the destination R7 from R0 which includes microloop-avoidance path and the post-convergence path after the link is enabled between R0 and R7.

13

CHAPTER

Configure OSPF Database Protection

IN THIS CHAPTER

- [Configuring OSPF Database Protection | 587](#)
-

Configuring OSPF Database Protection

IN THIS SECTION

- [OSPF Database Protection Overview | 587](#)
- [Configuring OSPF Database Protection | 588](#)

OSPF Database Protection Overview

OSPF database protection allows you to limit the number of link-state advertisements (LSAs) not generated by the local router in a given OSPF routing instance, helping to protect the link-state database from being flooded with excessive LSAs. This feature is particularly useful if VPN routing and forwarding is configured on your provider edge and customer edge routers using OSPF as the routing protocol. An overrun link-state database on the customer edge router can exhaust resources on the provider edge router and impact the rest of the service provider network.

When you enable OSPF database protection, the maximum number of LSAs you specify includes all LSAs whose advertising router ID is not equal to the local router ID (nonself-generated LSAs). These might include external LSAs as well as LSAs with any scope such as the link, area, and autonomous system (AS).

Once the specified maximum LSA count is exceeded, the database typically enters into the ignore state. In this state, all neighbors are brought down, and nonself-generated LSAs are destroyed. In addition, the database sends out hellos but ignores all received packets. As a result, the database does not form any full neighbors, and therefore does not learn about new LSAs. However, if you have configured the **warning-only** option, only a warning is issued and the database does not enter the ignore state but continues to operate as before.

You can also configure one or more of the following options:

- A warning threshold for issuing a warning message before the LSA limit is reached.
- An ignore state time during which the database must remain in the ignore state and after which normal operations can be resumed.
- An ignore state count that limits the number of times the database can enter the ignore state, after which it must enter the isolate state. The isolate state is very similar to the ignore state, but has one important difference: once the database enters the isolate state, it must remain there until you issue a command to clear database protection before it can return to normal operations.

- A reset time during which the database must stay out of the ignore or isolate state before it is returned to a normal operating state.

SEE ALSO

| *database-protection*

Configuring OSPF Database Protection

By configuring OSPF database protection, you can help prevent your OSPF link-state database from being overrun with excessive LSAs that are not generated by the local router. You specify the maximum number of LSAs whose advertising router ID is not the same as the local router ID in an OSPF instance. This feature is particularly useful if your provider edge and customer edge routers are configured with VPN routing and forwarding using OSPF.

OSPF database protection is supported on:

- Logical systems
- All routing instances supported by OSPFv2 and OSPFv3
- OSPFv2 and OSPFv3 topologies
- OSPFv3 realms

To configure OSPF database protection:

1. Include the `database-protection` statement at one of the following hierarchy levels:

- [edit protocols ospf | ospf3]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf |ospf3)]
- [edit routing-instances *routing-instance-name* protocols (ospf |ospf3)]
- [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast)]

2. Include the `maximum-lsa` *number* statement.



NOTE: The `maximum-lsa` statement is mandatory, and there is no default value for it. If you omit this statement, you cannot configure OSPF database protection.

3. (Optional) Include the following statements:

- **ignore-count *number***—Specify the number of times the database can enter the ignore state before it goes into the isolate state.
 - **ignore-time *seconds***—Specify the time limit the database must remain in the ignore state before it resumes regular operations.
 - **reset-time *seconds***—Specify the time during which the database must operate without being in either the ignore or isolate state before it is reset to a normal operating state.
 - **warning-threshold *percent***—Specify the percent of the maximum LSA number that must be exceeded before a warning message is issued.
4. (Optional) Include the `warning-only` statement to prevent the database from entering the ignore state or isolate state when the maximum LSA count is exceeded.



NOTE: If you include the `warning-only` statement, values for the other optional statements at the same hierarchy level are not used when the maximum LSA number is exceeded.

5. Verify your configuration by checking the database protection fields in the output of the `show ospf overview` command.

RELATED DOCUMENTATION

| [*database-protection*](#)

14

CHAPTER

Configure OSPF Routing Policy

IN THIS CHAPTER

- [Configuring OSPF Routing Policy | 591](#)
-

Configuring OSPF Routing Policy

IN THIS SECTION

- [Understanding Routing Policies | 591](#)
- [Understanding OSPF Routing Policy | 595](#)
- [Understanding Backup Selection Policy for OSPF Protocol | 597](#)
- [Configuring Backup Selection Policy for the OSPF Protocol | 599](#)
- [Topology-Independent Loop-Free Alternate with Segment Routing for OSPF | 607](#)
- [Example: Configuring Backup Selection Policy for the OSPF or OSPF3 Protocol | 610](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table | 644](#)
- [Example: Redistributing Static Routes into OSPF | 649](#)
- [Example: Configuring an OSPF Import Policy | 654](#)
- [Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF | 660](#)
- [Import and Export Policies for Network Summaries Overview | 667](#)
- [Example: Configuring an OSPF Export Policy for Network Summaries | 667](#)
- [Example: Configuring an OSPF Import Policy for Network Summaries | 680](#)
- [Example: Redistributing OSPF Routes into IS-IS | 693](#)

Understanding Routing Policies

IN THIS SECTION

- [Importing and Exporting Routes | 592](#)
- [Active and Inactive Routes | 594](#)
- [Explicitly Configured Routes | 594](#)
- [Dynamic Database | 594](#)

For some routing platform vendors, the flow of routes occurs between various protocols. If, for example, you want to configure redistribution from RIP to OSPF, the RIP process tells the OSPF process that it has routes that might be included for redistribution. In Junos OS, there is not much direct interaction between the routing protocols. Instead, there are central gathering points where all protocols install their routing information. These are the main unicast routing tables `inet.0` and `inet6.0`.

From these tables, the routing protocols calculate the best route to each destination and place these routes in a forwarding table. These routes are then used to forward routing protocol traffic toward a destination, and they can be advertised to neighbors.

Importing and Exporting Routes

Two terms—*import* and *export*—explain how routes move between the routing protocols and the routing table.


- When the Routing Engine places the routes of a routing protocol into the routing table, it is *importing* routes into the routing table.
- When the Routing Engine uses active routes from the routing table to send a protocol advertisement, it is *exporting* routes from the routing table.



NOTE: The process of moving routes between a routing protocol and the routing table is described always *from the point of view of the routing table*. That is, routes are *imported into* a routing table from a routing protocol and they are *exported from* a routing table to a routing protocol. Remember this distinction when working with routing policies.

As shown in [Figure 36 on page 593](#), you use import routing policies to control which routes are placed in the routing table, and export routing policies to control which routes are advertised from the routing table to neighbors.

Figure 36: Importing and Exporting Routes



6001706

In general, the routing protocols place all their routes in the routing table and advertise a limited set of routes from the routing table. The general rules for handling the routing information between the routing protocols and the routing table are known as the *routing policy framework*.

The routing policy framework is composed of default rules for each routing protocol that determine which routes the protocol places in the routing table and advertises from the routing table. The default rules for each routing protocol are known as *default routing policies*.

You can create routing policies to preempt the default policies, which are always present. A *routing policy* allows you to modify the routing policy framework to suit your needs. You can create and implement your own routing policies to do the following:

- Control which routes a routing protocol places in the routing table.
- Control which active routes a routing protocol advertises from the routing table. An *active route* is a route that is chosen from all routes in the routing table to reach a destination.
- Manipulate the route characteristics as a routing protocol places the route in the routing table or advertises the route from the routing table.

You can manipulate the route characteristics to control which route is selected as the active route to reach a destination. The active route is placed in the forwarding table and is used to forward traffic toward the route's destination. In general, the active route is also advertised to a router's neighbors.

Active and Inactive Routes

When multiple routes for a destination exist in the routing table, the protocol selects an active route and that route is placed in the appropriate routing table. For equal-cost routes, the Junos OS places multiple next hops in the appropriate routing table.

When a protocol is exporting routes from the routing table, it exports active routes only. This applies to actions specified by both default and user-defined export policies.

When evaluating routes for export, the Routing Engine uses only active routes from the routing table. For example, if a routing table contains multiple routes to the same destination and one route has a preferable metric, only that route is evaluated. In other words, an export policy does not evaluate all routes; it evaluates only those routes that a routing protocol is allowed to advertise to a neighbor.



NOTE: By default, BGP advertises active routes. However, you can configure BGP to advertise *inactive routes*, which go to the same destination as other routes but have less preferable metrics.

Explicitly Configured Routes

An *explicitly configured route* is a route that you have configured. *Direct routes* are not explicitly configured. They are created as a result of IP addresses being configured on an interface. Explicitly configured routes include aggregate, generated, local, and static routes. (An *aggregate route* is a route that distills groups of routes with common addresses into one route. A *generated route* is a route used when the routing table has no information about how to reach a particular destination. A *local route* is an IP address assigned to a router interface. A *static route* is an unchanging route to a destination.)

The policy framework software treats direct and explicitly configured routes as if they are learned through routing protocols; therefore, they can be imported into the routing table. Routes cannot be exported from the routing table to the pseudoprotocol, because this protocol is not a real routing protocol. However, aggregate, direct, generated, and static routes can be exported from the routing table to routing protocols, whereas local routes cannot.

Dynamic Database

In Junos OS Release 9.5 and later, you can configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required by the standard configuration database. As a result, you can quickly commit these routing policies and policy objects, which can be referenced and applied in the standard configuration as needed. BGP is the only protocol to which you can apply routing policies that reference policies configured in the dynamic database. After a routing policy based on the dynamic database is configured and committed in the standard configuration, you can quickly make changes to existing routing policies by modifying policy objects in the dynamic

database. Because Junos OS does not validate configuration changes to the dynamic database, when you use this feature, you should test and verify all configuration changes before committing them.

SEE ALSO

| [Example: Configuring Dynamic Routing Policies](#)

Understanding OSPF Routing Policy

IN THIS SECTION

- [Routing Policy Terms | 596](#)
- [Routing Policy Match Conditions | 596](#)
- [Routing Policy Actions | 597](#)

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration. Once a policy is created and named, it must be applied before it is active.

In the `import` statement, you list the name of the routing policy used to filter OSPF external routes from being installed into the routing tables of OSPF neighbors. You can filter the routes, but not link-state address (LSA) flooding. An external route is a route that is outside the OSPF Autonomous System (AS). The `import` policy does not impact the OSPF database. This means that the `import` policy has no impact on the link-state advertisements. The default `import` policy for OSPF is to accept all learned routes and import them into the routing table.

In the `export` statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF. The default `export` policy for OSPF is to reject everything, except LSA type1 and type 2. OSPF does not actually export its internally learned routes (the directly connected routes on interfaces that are running the protocol). OSPF uses link-state advertisement (LSA) flooding to advertise both local routes and learned routes, and LSA flooding is not affected by the `export` policy.

By default, if a routing device has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

This topic describes the following information:

Routing Policy Terms

Routing policies are made up of one or more terms. A term is a named structure in which match conditions and actions are defined. You can define one or more terms. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Routing Policy Match Conditions

A match condition defines the criteria that a route must match for an action to take place. You can define one or more match conditions for each term. If a route matches all of the match conditions for a particular term, the actions defined for that term are processed.

Each term can include two statements, `from` and `to`, that define the match conditions:

- In the `from` statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The `from` statement is optional. If you omit the `from` and the `to` statements, all routes are considered to match.



NOTE: In export policies, omitting the `from` statement from a routing policy term might lead to unexpected results.

- In the `to` statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The order of the match conditions in a term is not important because a route must match all match conditions in a term for an action to be taken.

For a complete list of match conditions, see [Configuring Match Conditions in Routing Policy Terms](#).

Routing Policy Actions

An action defines what the routing device does with the route when the route matches all the match conditions in the `from` and `to` statements for a particular term. If a term does not have `from` and `to` statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the `then` statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.

The `then` statement is optional. If you omit it, one of the following occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the `accept` or `reject` action specified by the default policy is executed.

For a complete list of routing policy actions, see [Configuring Actions in Routing Policy Terms](#).

Understanding Backup Selection Policy for OSPF Protocol

Support for OSPF loop-free alternate (LFA) routes essentially adds IP fast-reroute capability for OSPF. Junos OS precomputes multiple loop-free backup routes for all OSPF routes. These backup routes are pre-installed in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. The selection of LFA is done randomly by selecting any matching LFA to progress to the given destination. This does not ensure best backup coverage available for the network. In order to choose the best LFA, Junos OS allows you to configure network-wide backup selection policies for each destination (IPv4 and IPv6) and a primary next-hop interface. These policies are evaluated based on `admin-group`, `srlg`, `bandwidth`, `protection-type`, `metric`, and `node` information.

During backup shortest-path-first (SPF) computation, each node and link attribute of the backup path is accumulated by IGP and is associated with every node (router) in the topology. The next hop in the best backup path is selected as the backup next hop in the routing table. In general, backup evaluation policy rules are categorized into the following types:

- Pruning – Rules configured to select the eligible backup path.
- Ordering – Rules configured to select the best among the eligible backup paths.

The backup selection policies can be configured with both pruning and ordering rules. While evaluating the backup policies, each backup path is assigned a score, an integer value that signifies the total weight of the evaluated criteria. The backup path with the highest score is selected.

To enforce LFA selection, configure various rules for the following attributes:

- admin-group– Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. These configured administrative groups are defined under protocol MPLS. You can use administrative groups to implement a variety of backup selection policies using exclude, include-all, include-any, or preference.
- srlg– A shared risk link group (SRLG) is a set of links sharing a common resource, which affects all links in the set if the common resource fails. These links share the same risk of failure and are therefore considered to belong to the same SRLG. For example, links sharing a common fiber are said to be in the same SRLG because a fault with the fiber might cause all links in the group to fail. An SRLG is represented by a 32-bit number unique within an IGP (OSPF) domain. A link might belong to multiple SRLGs. You can define the backup selection to either allow or reject the common SRLGs between the primary and the backup path. This rejection of common SRLGs are based on the non-existence of link having common SRLGs in the primary next-hop and the backup SPF.



NOTE: Administrative groups and SRLGs can be created only for default topologies.

- bandwidth–The bandwidth specifies the bandwidth constraints between the primary and the backup path. The backup next-hop link can be used only if the bandwidth of the backup next-hop interface is greater than or equal to the bandwidth of the primary next hop.
- protection-type– The protection-type protects the destination from node failure of the primary node or link failure of the primary link. You can configure node, link, or node-link to protect the destination. If link-node is configured, then the node-protecting LFA is preferred over link-protection LFA.
- node- The node is per-node policy information. Here, node can be a directly connected router, remote router like RSVP backup LSP tail-end, or any other router in the backup SPF path. The nodes are identified through the route-id advertised by a node in the LSP. You can list the nodes to either prefer or exclude them in the backup path.

- **metric**— Metric decides how the LFAs should be preferred. In backup selection path, root metric and dest-metric are the two types of metrics. root-metric indicates the metric to the one-hop neighbor or a remote router such as an RSVP backup LSP tail-end router. The dest-metric indicates the metric from a one-hop neighbor or remote router such as an RSVP backup LSP tail-end router to the final destination. The metric evaluation is done either in ascending or descending order. By default, the first preference is given to backup paths with lowest destination evaluation and then to backup paths with lowest root metrics.

The evaluation-order allows you to control the order and criteria of evaluating these attributes in the backup path. You can explicitly configure the evaluation order. Only the configured attributes influence the backup path selection. The default order of evaluation of these attributes for the LFA is [admin-group srlg bandwidth protection-type node metric] .



NOTE: TE attributes are not supported in OSPFv3 and cannot be used for backup selection policy evaluation for IPv6 prefixes.

SEE ALSO

| [backup-selection \(Protocols IS-IS\)](#)

Configuring Backup Selection Policy for the OSPF Protocol

Support for OSPF loop-free alternate (LFA) routes essentially adds IP fast-reroute capability for OSPF. Junos OS precomputes multiple loop-free backup routes for all OSPF routes. These backup routes are pre-installed in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. The selection of LFA is done randomly by selecting any matching LFA to progress to the given destination. This does not ensure best backup coverage available for the network. In order to choose the best LFA, Junos OS allows you to configure network-wide backup selection policies for each destination (IPv4 and IPv6) and a primary next-hop interface. These policies are evaluated based on admin-group, srlg, bandwidth, protection-type, metric, and node information.

Before you begin to configure the backup selection policy for the OSPF protocol:

- Configure the router interfaces. See the *Junos OS Network Management Administration Guide for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.

To configure the backup selection policy for the OSPF protocol:

1. Configure per-packet load balancing.

```
[edit policy-options]
user@host# set policy-statement ecmp term 1 then load-balance per-packet
```

2. Enable RSVP on all the interfaces.

```
[edit protocols]
user@host# set rsvp interface all
```

3. Configure administrative groups.

```
[edit protocols mpls]
user@host# set admin-groups group-name
```

4. Configure srlg values.

```
[edit routing-options]
user@host# set srlg srlg-name srlg-value srlg-value
```

5. Enable MPLS on all the interfaces.

```
[edit protocols mpls]
user@host# set interface all
```

6. Apply MPLS to an interface configured with an administrative group.

```
[edit protocols mpls]
user@host# set interface interface-name admin-group group-name
```

7. Configure the ID of the router.

```
[edit routing-options]
user@host# set router-id router-id
```

8. Apply the routing policy to all equal cost multipaths exported from the routing table to the forwarding table.

```
[edit routing-options]
user@host# set forwarding-table export ecmp
```

9. Enable link protection and configure metric values on all the interfaces for an area.

```
[edit protocols ospf]
user@host# set area area-id interface interface-name link-protection
user@host# set area area-id interface interface-name metric metric
```

10. Configure the administrative group of the backup selection policy for an IP address.
You can choose to exclude, include all, include any, or prefer the administrative groups from the backup path.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name admin-group
```

- Specify the administrative group to be excluded.

```
[edit routing-options backup-selection destination ip-address interface interface-name
admin-group]
user@host# set exclude group-name
```

The backup path is not selected as the loop-free alternate (LFA) or backup nexthop if any of the links in the path have any one of the listed administrative groups.

For example, to exclude the group c1 from the administrative group:

```
[edit routing-options backup-selection destination 0.0.0.0/0 interface all admin-group]
user@host# set exclude c1
```

- Configure all the administrative groups if each link in the backup path requires all the listed administrative groups in order to accept the path.

```
[edit routing-options backup-selection destination ip-address interface interface-name
admin-group]
user@host# set include-all group-name
```

For example, to set all the administrative groups if each link requires all the listed administrative groups in order to accept the path:

```
[edit routing-options backup-selection destination 0.0.0.0/0 interface all admin-group]
user@host# set include-all c2
```

- Configure any administrative group if each link in the backup path requires at least one of the listed administrative groups in order to select the path.

```
[edit routing-options backup-selection destination ip-address interface interface-name
admin-group]
user@host# set include-any group-name
```

For example, to set any administrative group if each link in the backup path requires at least one of the listed administrative groups in order to select the path:

```
[edit routing-options backup-selection destination 0.0.0.0/0 interface all admin-group]
user@host# set include-any c3
```

- Define an ordered set of an administrative group that specifies the preference of the backup path.

The leftmost element in the set is given the highest preference.

```
[edit routing-options backup-selection destination ip-address interface interface-name
admin-group]
user@host# set preference group-name
```

For example, to set an ordered set of an administrative group that specifies the preference of the backup path:

```
[edit routing-options backup-selection destination 0.0.0.0/0 interface all admin-group]
user@host# set preference c4
```

11. Configure the backup path to allow the selection of the backup next hop only if the bandwidth is greater than or equal to the bandwidth of the primary next hop.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name bandwidth-
greater-equal-primary
```

12. Configure the backup path to specify the metric from the one-hop neighbor or from the remote router such as an RSVP backup label-switched-path (LSP) tail-end router to the final destination. The destination metric can be either highest or lowest.
 - Configure the backup path that has the highest destination metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name dest-
metric highest
```

- Configure the backup path that has the lowest destination metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name dest-
metric lowest
```

13. Configure the backup path that is a downstream path to the destination.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name downstream-
paths-only
```

14. Set the order of preference of the root and the destination metric during backup path selection. The preference order can be :

- [root dest] – Backup path selection or preference is first based on the root-metric criteria. If the criteria of all the root-metric is the same, then the selection or preference is based on the dest-metric.
- [dest root] – Backup path selection or preference is first based on the dest-metric criteria. If the criteria of all the dest-metric is the same, then the selection is based on the root-metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name metric-
order dest
user@host# set backup-selection destination ip-address interface interface-name metric-
order root
```

15. Configure the backup path to define a list of loop-back IP addresses of the adjacent neighbors to either exclude or prefer in the backup path selection.

The neighbor can be a local (adjacent router) neighbor, remote neighbor, or any other router in the backup path.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name node
```

- Configure the list of neighbors to be excluded.

```
[edit routing-options backup-selection destination ip-address interface interface-name
node]
user@host# set exclude node-address
```

The backup path that has a router from the list is not selected as the loop-free alternative or backup next hop.

- Configure an ordered set of neighbors to be preferred.

```
[edit routing-options backup-selection destination ip-address interface interface-name
node]
user@host# set preference node-address
```

The backup path having the leftmost neighbor is selected.

16. Configure the backup path to specify the required protection type of the backup path to be link, node, or node-link.

- Select the backup path that provides link protection.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name
protection-type link
```

- Select the backup path that provides node protection.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name
protection-type node
```

- Select the backup path that allows either node or link protection LFA where node-protection LFA is preferred over link-protection LFA.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface interface-name
protection-type node-link
```

17. Specify the metric to the one-hop neighbor or to the remote router such as an RSVP backup label-switched-path (LSP) tail-end router.

- Select the path with highest root metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface all root-metric highest
```

- Select the path with lowest root metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface all root-metric lowest
```

18. Configure the backup selection path to either allow or reject the common shared risk link groups (SRLGs) between the primary link and each link in the backup path.

- Configure the backup path to allow common srlgs between the primary link and each link in the backup path.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface all srlg loose
```

A backup path with a fewer number of srlg collisions is preferred.

- Configure the backup path to reject the backup path that has common srlgs between the primary next-hop link and each link in the backup path.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface all srlg strict
```

- 19.** Configure the backup path to control the order and the criteria of evaluating the backup path based on the administrative group, srlg, bandwidth, protection type, node, and metric.

The default order of evaluation is admin-group, srlg, bandwidth, protection-type, node, and metric.

```
[edit routing-options]
user@host# set backup-selection destination ip-address interface all evaluation-order admin-
group
user@host# set backup-selection destination ip-address interface all evaluation-order srlg
user@host# set backup-selection destination ip-address interface all evaluation-order
bandwidth
```

SEE ALSO

| [backup-selection \(Protocols IS-IS\)](#)

Topology-Independent Loop-Free Alternate with Segment Routing for OSPF

IN THIS SECTION

- [Topology-Independent Loop-Free Alternate with Segment Routing for OSPF Overview | 607](#)
- [Configuring Topology-Independent Loop-Free Alternate with Segment Routing for OSPF | 609](#)

Topology-Independent Loop-Free Alternate with Segment Routing for OSPF Overview

IN THIS SECTION

- [Benefits of Using Topology-Independent Loop-Free Alternate with Segment Routing | 608](#)

This section describes the TI-LFA feature for OSPF.

When used with OSPF, TI-LFA provides protection against link failure, node failure, fate-sharing failures, and shared risk link group failures. In link failure mode, the destination is protected if the link fails. In node protection mode, the destination is protected if the neighbor connected to the primary link fails. To determine the node-protecting post-convergence path, the cost of all the links leaving the neighbor is assumed to increase by a configurable amount.

Configure fate-sharing protection in TI-LFA networks for segment routing to choose a fast reroute path that does not include fate-sharing groups in the topology-independent loop-free alternate (TI-LFA) backup paths to avoid fate-sharing failures. With fate-sharing protection, a list of fate-sharing groups are configured on each PLR with the links in each fate-sharing group identified by their respective IP addresses. The PLR associates a cost with each fate-sharing group. The fate-sharing-aware post-convergence path is computed by assuming that the cost of each link in the same fate-sharing group as the failed link has increased the cost associated with that group.

Configure Shared Risk Link Group (SRLG) protection in TI-LFA networks for segment routing to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. SRLGs share a common fibre and they also share the risks of a broken link. When one link in an SRLG fails, other links in the group might also fail. Therefore, you need to avoid links that share the same risk as the protected link in the backup path. Configuring SRLG protection prevents TI-

LFA from selecting backup paths that include a shared risk link. If you have configured SRLG protection then OSPFv2 computes the fast reroute path that is aligned with the post convergence path and excludes the links that belong to the SRLG of the protected link. All local and remote links that are from the same SRLG as the protected link are excluded from the TI-LFA backup path. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface. Currently you cannot enable SRLG protection in IPv6 networks and in networks with mult topology.

In order to construct a backup path that follows the post-convergence path, TI-LFA can use several labels in the label stack that define the backup path. If the number of labels required to construct a particular post-convergence backup path exceeds a certain amount, it is useful in some circumstances to not install that backup path. You can configure the maximum number of labels that a backup path can have in order to be installed. The default value is 3, with a range of 2 through 5.

It is often the case that the post-convergence path for a given failure is actually a set of equal-cost paths. TI-LFA attempts to construct the backup paths to a given destination using multiple equal-cost paths in the post-failure topology. Depending on the topology, TI-LFA might need to use different label stacks to accurately construct those equal-cost backup paths. By default, TI-LFA only installs one backup path for a given destination. However, you can configure the value in the range from 1 through 8.

Benefits of Using Topology-Independent Loop-Free Alternate with Segment Routing

- Loop-free alternate (LFA) and remote LFA (RLFA) have been used to provide fast-reroute protection for several years. With LFA, a point of local repair (PLR) determines whether or not a packet sent to one of its direct neighbors reaches its destination without looping back through the PLR. In a typical network topology, approximately 40 to 60 percent of the destinations can be protected by LFA. Remote LFA expands on the concept of LFA by allowing the PLR to impose a single label to tunnel the packet to a repair tunnel endpoint from which the packet can reach its destination without looping back through the PLR. Using remote LFA, more destinations can be protected by the PLR compared to LFA. However, depending on the network topology, the percentage of destinations protected by remote LFA is usually less than 100 percent.
- Topology-independent LFA (TI-LFA) extends the concept of LFA and remote LFA by allowing the PLR to use deeper label stacks to construct backup paths. In addition, TI-LFA imposes the constraint that the backup path used by the PLR be the same path that a packet takes once the interior gateway protocol (IGP) has converged for a given failure scenario. This path is referred to as the post-convergence path.
- Using the post-convergence path as the backup path has some desirable characteristics. For some topologies, a network operator only needs to make sure that the network has enough capacity to carry the traffic along the post-convergence path after a failure. In these cases, a network operator does not need to allocate additional capacity to deal with the traffic pattern immediately after the failure while the backup path is active, because the backup path follows the post-convergence path.

- When used with OSPF, TI-LFA provides protection against link failure and node failure.

Configuring Topology-Independent Loop-Free Alternate with Segment Routing for OSPF

Before you configure TI-LFA for OSPF, be sure you configure SPRING or segment routing.

Junos supports creation of OSPF topology-independent TI-LFA backup paths where the prefix SID is learned from a segment routing mapping server advertisement when the PLR and mapping server are both in the same OSPF area.

To configure TI-LFA using SPRING for OSPF, you must do the following:

1. Enable TI-LFA for OSPF protocol.

```
[edit protocols ospf backup-spf-options]
user@R1# set use-post-convergence-lfa
```

2. (Optional) Configure backup shortest path first (SPF) attributes such as maximum equal-cost multipath (ECMP) backup paths and maximum labels for TI-LFA for the OSPF protocol.

```
[edit protocols ospf backup-spf-options use-post-convergence-lfa]
user@R1# set maximum-backup-paths maximum-backup-paths
user@R1# set maximum-labels maximum-labels
```

3. Configure the computation and installation of a backup path that follows the post-convergence path on the given area and interface for the OSPF protocol.

```
[edit protocols ospf area area-id interface interface-name]
user@R1# set post-convergence-lfa
```

4. (Optional) Enable node protection for a given area and interface.

```
[edit protocols ospf area area-id interface interface-name post-convergence-lfa]
user@R1# set node-protection
```

5. (Optional) Enable fate-sharing protection for a given area and interface.

```
[edit protocols ospf area area-id interface interface-name post-convergence-lfa]
user@R1# set fate-sharing-protection
```

6. (Optional) Enable SRLG protection for a given area and interface.

```
[edit protocols ospf area area-id interface interface-name post-convergence-lfa]
user@R1# set srlg-protection
```

RELATED DOCUMENTATION

[source-packet-routing](#)

[use-post-convergence-lfa](#)

[post-convergence-lfa](#)

Example: Configuring Backup Selection Policy for the OSPF or OSPF3 Protocol

IN THIS SECTION

- [Requirements | 610](#)
- [Overview | 611](#)
- [Configuration | 612](#)
- [Verification | 637](#)

This example shows how to configure the backup selection policy for the OSPF or OSPF3 protocol, which enables you to select a loop-free alternate (LFA) in the network.

When you enable backup selection policies, Junos OS allows selection of LFA based on the policy rules and attributes of the links and nodes in the network. These attributes are admin-group, srlg, bandwidth, protection-type, metric, and node.

Requirements

This example uses the following hardware and software components:

- Combination of eight routers
- Supported Junos OS Release running on all devices

Before you begin:

1. Configure the device interfaces.
2. Configure OSPF.

Overview

IN THIS SECTION


- [Topology | 611](#)

In Junos OS, the default loop-free alternative (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are configured for each destination (IPv4 and IPv6) and a primary next-hop interface. These backup policies enforce LFA selection based on admin-group, srlg, bandwidth, protection-type, metric, and node attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup next-hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured for each destination (IPv4 and IPv6) and a primary next-hop interface. The first or the best backup path is selected and installed as the backup next hop in the routing table. To configure the backup selection policy, include the backup-selection configuration statement at the [edit routing-options] hierarchy level. The show backup-selection command displays the configured policies for a given interface and destination. The display can be filtered against a particular destination, prefix, interface, or logical systems.

Topology

In this topology shown in [Figure 37 on page 612](#), the backup selection policy is configured on Device R3.

Figure 37: Example Backup Selection Policy for OSPF or OPSF3



8043139

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 613](#)
- [Configuring Device R3 | 627](#)
- [Results | 631](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

RO

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/30
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:db8:10:1:1::1/64
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/1/0 unit 0 family inet address 172.16.15.1/30
set interfaces ge-0/1/0 unit 0 family inet6 address 2001:db8:15:1:1::1/64
set interfaces ge-0/1/0 unit 0 family mpls
set interfaces xe-0/2/0 unit 0 family inet address 172.16.20.1/30
set interfaces xe-0/2/0 unit 0 family inet6 address 2001:db8:20:1:1::1/64
set interfaces xe-0/2/0 unit 0 family mpls
set interfaces ge-1/0/5 unit 0 family inet address 172.16.150.1/24
set interfaces ge-1/0/5 unit 0 family inet6 address 2001:db8:150:1:1::1/64
set interfaces ge-1/0/5 unit 0 family mpls
set interfaces ge-1/1/1 unit 0 family inet address 172.16.30.1/30
set interfaces ge-1/1/1 unit 0 family inet6 address 2001:db8:30:1:1::1/64
set interfaces ge-1/1/1 unit 0 family mpls
set interfaces xe-1/3/0 unit 0 family inet address 172.16.25.1/30
set interfaces xe-1/3/0 unit 0 family inet6 address 2001:db8:25:1:1::1/64
set interfaces xe-1/3/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.10.10.10/32 primary
set interfaces lo0 unit 0 family inet6 address 2001:db8::10:10:10:10/128 primary
set interfaces lo0 unit 0 family mpls
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 10.10.10.10
```

```
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/1/0.0 metric 18
set protocols ospf area 0.0.0.0 interface xe-0/2/0.0 metric 51
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 metric 23
set protocols ospf area 0.0.0.0 interface xe-1/3/0.0 metric 52
set protocols ospf area 0.0.0.0 interface ge-1/0/5.0
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/1/0.0 metric 18
set protocols ospf3 area 0.0.0.0 interface xe-0/2/0.0 metric 51
```

```

set protocols ospf3 area 0.0.0.0 interface ge-1/1/1.0 metric 23
set protocols ospf3 area 0.0.0.0 interface xe-1/3/0.0 metric 52
set protocols ospf3 area 0.0.0.0 interface ge-1/0/5.0

```

R1

```

set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.2/30
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:db8:10:1:1::2/64
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/5 unit 0 family inet address 172.16.35.1/30
set interfaces ge-0/0/5 unit 0 family inet6 address 2001:db8:35:1:1::1/64
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces xe-0/2/0 unit 0 family inet address 172.16.40.1/30
set interfaces xe-0/2/0 unit 0 family inet6 address 2001:db8:40:1:1::1/64
set interfaces xe-0/2/0 unit 0 family mpls
set interfaces xe-0/3/0 unit 0 family inet address 172.16.45.1/30
set interfaces xe-0/3/0 unit 0 family inet6 address 2001:db8:45:1:1::1/64
set interfaces xe-0/3/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 172.16.1.1/32 primary
set interfaces lo0 unit 0 family inet6 address 2001:db8::1:1:1:1/128 primary
set interfaces lo0 unit 0 family mpls
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.1.1
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6

```

```
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols mpls interface ge-0/0/0.0 srlg srlg9
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/2/0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/3/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/2/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/3/0.0 metric 10
```

R2

```
set interfaces ge-0/0/2 unit 0 family inet address 172.16.35.2/30
set interfaces ge-0/0/2 unit 0 family inet6 address 2001:db8:35:1:1::2/64
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/1/0 unit 0 family inet address 172.16.50.1/30
```

```
set interfaces ge-0/1/0 unit 0 family inet6 address 2001:db8:50:1:1::1/64
set interfaces ge-0/1/0 unit 0 family mpls
set interfaces xe-0/2/1 unit 0 family inet address 172.16.55.1/30
set interfaces xe-0/2/1 unit 0 family inet6 address 2001:db8:55:1:1::1/64
set interfaces xe-0/2/1 unit 0 family mpls
set interfaces ge-1/0/2 unit 0 family inet address 172.16.60.1/30
set interfaces ge-1/0/2 unit 0 family inet6 address 2001:db8:60:1:1::1/64
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-1/0/9 unit 0 family inet address 172.16.65.1/30
set interfaces ge-1/0/9 unit 0 family inet6 address 2001:db8:65:1:1::1/64
set interfaces ge-1/0/9 unit 0 family mpls
set interfaces ge-1/1/5 unit 0 family inet address 172.16.70.1/30
set interfaces ge-1/1/5 unit 0 family inet6 address 2001:db8:70:1:1::1/64
set interfaces ge-1/1/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 172.16.2.2/32 primary
set interfaces lo0 unit 0 family inet6 address 2001:db8::2:2:2:2/128 primary
set interfaces lo0 unit 0 family mpls
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.2.2
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
```

```

set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols mpls interface ge-0/1/0.0 srlg srlg1
set protocols mpls interface ge-1/0/9.0 srlg srlg1
set protocols mpls interface ge-1/1/5.0 srlg srlg7
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/1/0.0 link-protection
set protocols ospf area 0.0.0.0 interface xe-0/2/1.0 metric 12
set protocols ospf area 0.0.0.0 interface ge-1/0/2.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/0/9.0 metric 12
set protocols ospf area 0.0.0.0 interface ge-1/1/5.0 metric 13
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/1/0.0 link-protection
set protocols ospf3 area 0.0.0.0 interface xe-0/2/1.0 metric 12
set protocols ospf3 area 0.0.0.0 interface ge-1/0/2.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-1/0/9.0 metric 12
set protocols ospf3 area 0.0.0.0 interface ge-1/1/5.0 metric 13

```

R3

```

set interfaces ge-0/0/5 unit 0 family inet address 172.16.50.2/30
set interfaces ge-0/0/5 unit 0 family inet6 address 2001:db8:50:1:1::2/64
set interfaces ge-0/0/5 unit 0 family mpls

```

```
set interfaces xe-0/3/1 unit 0 family inet address 172.16.75.1/30
set interfaces xe-0/3/1 unit 0 family inet6 address 2001:db8:75:1:1::1/64
set interfaces xe-0/3/1 unit 0 family mpls
set interfaces ge-1/0/0 unit 0 family inet address 172.16.80.1/30
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8:80:1:1::1/64
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/5 unit 0 family inet address 172.16.200.1/24
set interfaces ge-1/0/5 unit 0 family inet6 address 2001:db8:200:1:1::1/64
set interfaces ge-1/0/6 unit 0 family inet address 172.16.85.1/30
set interfaces ge-1/0/6 unit 0 family inet6 address 2001:db8:85:1:1::1/64
set interfaces ge-1/0/6 unit 0 family mpls
set interfaces xe-1/3/0 unit 0 family inet address 172.16.90.1/30
set interfaces xe-1/3/0 unit 0 family inet6 address 2001:db8:90:1:1::1/64
set interfaces xe-1/3/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 172.16.3.3/32 primary
set interfaces lo0 unit 0 family inet6 address 2001:db8::3:3:3:3/128 primary
set interfaces lo0 unit 0 family mpls
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.3.3
set routing-options forwarding-table export ecmp
set routing-options backup-selection destination 10.1.1.0/30 interface xe-1/3/0.0 admin-group
include-all c2
set routing-options backup-selection destination 10.1.1.0/30 interface all admin-group exclude c3
set routing-options backup-selection destination 10.1.1.0/30 interface all srlg strict
set routing-options backup-selection destination 10.1.1.0/30 interface all protection-type node
set routing-options backup-selection destination 10.1.1.0/30 interface all bandwidth-greater-
equal-primary
set routing-options backup-selection destination 10.1.1.0/30 interface all neighbor preference
172.16.7.7
set routing-options backup-selection destination 10.1.1.0/30 interface all root-metric lowest
set routing-options backup-selection destination 10.1.1.0/30 interface all metric-order root
set routing-options backup-selection destination 172.16.30.0/30 interface all admin-group
```

```
exclude c5
set routing-options backup-selection destination 172.16.30.0/30 interface all srlg strict
set routing-options backup-selection destination 172.16.30.0/30 interface all protection-type
node
set routing-options backup-selection destination 172.16.30.0/30 interface all bandwidth-greater-
equal-primary
set routing-options backup-selection destination 172.16.30.0/30 interface all neighbor
preference 172.16.7.7
set routing-options backup-selection destination 172.16.30.0/30 interface all root-metric lowest
set routing-options backup-selection destination 172.16.30.0/30 interface all metric-order root
set routing-options backup-selection destination 172.16.45.0/30 interface all admin-group
exclude c5
set routing-options backup-selection destination 172.16.45.0/30 interface all srlg strict
set routing-options backup-selection destination 172.16.45.0/30 interface all protection-type
node
set routing-options backup-selection destination 172.16.45.0/30 interface all bandwidth-greater-
equal-primary
set routing-options backup-selection destination 172.16.45.0/30 interface all neighbor
preference 172.16.7.7
set routing-options backup-selection destination 172.16.45.0/30 interface all root-metric lowest
set routing-options backup-selection destination 172.16.45.1/30 interface all metric-order root
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
```

```
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols mpls interface ge-0/0/5.0 admin-group c0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/3/1.0 metric 21
set protocols ospf area 0.0.0.0 interface ge-1/0/0.0 metric 13
set protocols ospf area 0.0.0.0 interface ge-1/0/6.0 metric 15
set protocols ospf area 0.0.0.0 interface xe-1/3/0.0 link-protection
set protocols ospf area 0.0.0.0 interface xe-1/3/0.0 metric 22
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 link-protection
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/3/1.0 metric 21
set protocols ospf3 area 0.0.0.0 interface ge-1/0/0.0 metric 13
set protocols ospf3 area 0.0.0.0 interface ge-1/0/6.0 metric 15
set protocols ospf3 area 0.0.0.0 interface xe-1/3/0.0 link-protection
set protocols ospf3 area 0.0.0.0 interface xe-1/3/0.0 metric 22
set policy-options policy-statement ecmp term 1 then load-balance per-packet
```

R4

```
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
```

```
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.4.4
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface ge-0/1/0.0 metric 18
set protocols ospf area 0.0.0.0 interface xe-0/2/0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-1/3/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/1/0.0 metric 10
set protocols ospf area 0.0.0.0 interface xe-0/3/1.0 metric 21
set protocols ospf3 area 0.0.0.0 interface ge-0/1/0.0 metric 18
```

```
set protocols ospf3 area 0.0.0.0 interface xe-0/2/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-1/3/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-1/1/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/3/1.0 metric 21
```

R5

```
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.5.5
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
```

```
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface xe-0/2/0.0 metric 51
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0 metric 13
set protocols ospf area 0.0.0.0 interface ge-0/1/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/2/0.0 metric 51
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 metric 13
set protocols ospf3 area 0.0.0.0 interface ge-0/1/0.0 metric 10
```

R6

```
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.6.6
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
```

```
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 25
set protocols mpls admin-groups c26 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface xe-0/3/0.0 metric 52
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 12
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 metric 15
set protocols ospf area 0.0.0.0 interface xe-0/2/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface xe-0/3/0.0 metric 52
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 metric 12
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 metric 15
set protocols ospf3 area 0.0.0.0 interface xe-0/2/0.0 metric 10
```

R7

```
set routing-options srlg srlg1 srlg-value 1001
set routing-options srlg srlg2 srlg-value 1002
```

```
set routing-options srlg srlg3 srlg-value 1003
set routing-options srlg srlg4 srlg-value 1004
set routing-options srlg srlg5 srlg-value 1005
set routing-options srlg srlg6 srlg-value 1006
set routing-options srlg srlg7 srlg-value 1007
set routing-options srlg srlg8 srlg-value 1008
set routing-options srlg srlg9 srlg-value 1009
set routing-options srlg srlg10 srlg-value 10010
set routing-options srlg srlg11 srlg-value 10011
set routing-options srlg srlg12 srlg-value 10012
set routing-options router-id 172.16.7.7
set protocols rsvp interface all
set protocols mpls admin-groups c0 0
set protocols mpls admin-groups c1 1
set protocols mpls admin-groups c2 2
set protocols mpls admin-groups c3 3
set protocols mpls admin-groups c4 4
set protocols mpls admin-groups c5 5
set protocols mpls admin-groups c6 6
set protocols mpls admin-groups c7 7
set protocols mpls admin-groups c8 8
set protocols mpls admin-groups c9 9
set protocols mpls admin-groups c10 10
set protocols mpls admin-groups c11 11
set protocols mpls admin-groups c12 12
set protocols mpls admin-groups c13 13
set protocols mpls admin-groups c14 14
set protocols mpls admin-groups c15 15
set protocols mpls admin-groups c16 16
set protocols mpls admin-groups c17 17
set protocols mpls admin-groups c18 18
set protocols mpls admin-groups c19 19
set protocols mpls admin-groups c20 20
set protocols mpls admin-groups c21 21
set protocols mpls admin-groups c22 22
set protocols mpls admin-groups c23 23
set protocols mpls admin-groups c24 24
set protocols mpls admin-groups c25 26
set protocols mpls admin-groups c27 27
set protocols mpls admin-groups c28 28
set protocols mpls admin-groups c29 29
set protocols mpls admin-groups c30 30
set protocols mpls admin-groups c31 31
```

```

set protocols mpls interface all
set protocols mpls interface xe-0/3/0.0 srlg srlg8
set protocols ospf area 0.0.0.0 interface ge-0/1/5.0 metric 23
set protocols ospf area 0.0.0.0 interface xe-0/3/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/0/0.0 metric 13
set protocols ospf area 0.0.0.0 interface xe-1/3/0.0 metric 22
set protocols ospf area 0.0.0.0 interface xe-1/2/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-0/1/5.0 metric 23
set protocols ospf3 area 0.0.0.0 interface xe-0/3/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-1/0/0.0 metric 13
set protocols ospf3 area 0.0.0.0 interface xe-1/3/0.0 metric 22
set protocols ospf3 area 0.0.0.0 interface xe-1/2/0.0 metric 10

```

Configuring Device R3

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure Device R3:

1. Configure the interfaces.

```

[edit interfaces]
user@R3# set ge-0/0/5 unit 0 family inet address 172.16.50.2/30
user@R3# set ge-0/0/5 unit 0 family inet6 address 2001:db8:50:1:1::2/64
user@R3# set ge-0/0/5 unit 0 family mpls
user@R3# set xe-0/3/1 unit 0 family inet address 172.16.75.1/30
user@R3# set xe-0/3/1 unit 0 family inet6 address 2001:db8:75:1:1::1/64
user@R3# set xe-0/3/1 unit 0 family mpls
user@R3# set ge-1/0/0 unit 0 family inet address 172.16.80.1/30
user@R3# set ge-1/0/0 unit 0 family inet6 address 2001:db8:80:1:1::1/64
user@R3# set ge-1/0/0 unit 0 family mpls
user@R3# set ge-1/0/5 unit 0 family inet address 172.16.200.1/24
user@R3# set ge-1/0/5 unit 0 family inet6 address 2001:db8:200:1:1::1/64
user@R3# set ge-1/0/6 unit 0 family inet address 172.16.85.1/30
user@R3# set ge-1/0/6 unit 0 family inet6 address 2001:db8:85:1:1::1/64
user@R3# set ge-1/0/6 unit 0 family mpls
user@R3# set xe-1/3/0 unit 0 family inet address 172.16.90.1/30
user@R3# set xe-1/3/0 unit 0 family inet6 address 2001:db8:90:1:1::1/64

```

```

user@R3# set xe-1/3/0 unit 0 family mpls
user@R3# set lo0 unit 0 family inet address 172.16.3.3/32 primary
user@R3# set lo0 unit 0 family inet6 address 2001:db8::3:3:3:3/128 primary
user@R3# set lo0 unit 0 family mpls

```

2. Configure srlg values.

```

[edit routing-options]
user@R3# set srlg srlg1 srlg-value 1001
user@R3# set srlg srlg2 srlg-value 1002
user@R3# set srlg srlg3 srlg-value 1003
user@R3# set srlg srlg4 srlg-value 1004
user@R3# set srlg srlg5 srlg-value 1005
user@R3# set srlg srlg6 srlg-value 1006
user@R3# set srlg srlg7 srlg-value 1007
user@R3# set srlg srlg8 srlg-value 1008
user@R3# set srlg srlg9 srlg-value 1009
user@R3# set srlg srlg10 srlg-value 10010
user@R3# set srlg srlg11 srlg-value 10011
user@R3# set srlg srlg12 srlg-value 10012

```

3. Configure the ID of the router.

```

[edit routing-options]
user@R3# set router-id 172.16.3.3

```

4. Apply the routing policy to all equal-cost multipaths exported from the routing table to the forwarding table.

```

[edit routing-options]
user@R3# set forwarding-table export ecmp

```

5. Configure attributes of the backup selection policy.

```

[edit routing-options backup-selection]
user@R3# set destination 10.1.1.0/30 interface xe-1/3/0.0 admin-group include-all c2
user@R3# set destination 10.1.1.0/30 interface all admin-group exclude c3
user@R3# set destination 10.1.1.0/30 interface all srlg strict
user@R3# set destination 10.1.1.0/30 interface all protection-type node

```

```

user@R3# set destination 10.1.1.0/30 interface all bandwidth-greater-equal-primary
user@R3# set destination 10.1.1.0/30 interface all neighbor preference 172.16.7.7
user@R3# set destination 10.1.1.0/30 interface all root-metric lowest
user@R3# set destination 10.1.1.0/30 interface all metric-order root
user@R3# set destination 172.16.30.0/30 interface all admin-group exclude c5
user@R3# set destination 172.16.30.0/30 interface all srlg strict
user@R3# set destination 172.16.30.0/30 interface all protection-type node
user@R3# set destination 172.16.30.0/30 interface all bandwidth-greater-equal-primary
user@R3# set destination 172.16.30.0/30 interface all neighbor preference 172.16.7.7
user@R3# set destination 172.16.30.0/30 interface all root-metric lowest
user@R3# set destination 172.16.30.0/30 interface all metric-order root
user@R3# set destination 192.168.45.0/30 interface all admin-group exclude c5
user@R3# set destination 192.168.45.0/30 interface all srlg strict
user@R3# set destination 192.168.45.0/30 interface all protection-type node
user@R3# set destination 192.168.45.0/30 interface all bandwidth-greater-equal-primary
user@R3# set destination 192.168.45.0/30 interface all neighbor preference 172.16.7.7
user@R3# set destination 192.168.45.0/30 interface all root-metric lowest
user@R3# set destination 192.168.45.0/30 interface all metric-order root

```

6. Enable RSVP on all the interfaces.

```

[edit protocols]
user@R3# set rsvp interface all

```

7. Configure administrative groups.

```

[edit protocols mpls]
user@R3# set admin-groups c0 0
user@R3# set admin-groups c1 1
user@R3# set admin-groups c2 2
user@R3# set admin-groups c3 3
user@R3# set admin-groups c4 4
user@R3# set admin-groups c5 5
user@R3# set admin-groups c6 6
user@R3# set admin-groups c7 7
user@R3# set admin-groups c8 8
user@R3# set admin-groups c9 9
user@R3# set admin-groups c10 10
user@R3# set admin-groups c11 11
user@R3# set admin-groups c12 12
user@R3# set admin-groups c13 13

```

```

user@R3# set admin-groups c14 14
user@R3# set admin-groups c15 15
user@R3# set admin-groups c16 16
user@R3# set admin-groups c17 17
user@R3# set admin-groups c18 18
user@R3# set admin-groups c19 19
user@R3# set admin-groups c20 20
user@R3# set admin-groups c21 21
user@R3# set admin-groups c22 22
user@R3# set admin-groups c23 23
user@R3# set admin-groups c24 24
user@R3# set admin-groups c25 25
user@R3# set admin-groups c26 26
user@R3# set admin-groups c27 27
user@R3# set admin-groups c28 28
user@R3# set admin-groups c29 29
user@R3# set admin-groups c30 30
user@R3# set admin-groups c31 31

```

8. Enable MPLS on all the interfaces and configure administrative group for an interface.

```

[edit protocols mpls]
user@R3# set interface all
user@R3# set interface ge-0/0/5.0 admin-group c0

```

9. Enable link protection and configure metric values on all the interfaces for an OSPF area.

```

[edit protocols ospf]
user@R3# set area 0.0.0.0 interface ge-0/0/5.0 link-protection
user@R3# set area 0.0.0.0 interface ge-0/0/5.0 metric 10
user@R3# set area 0.0.0.0 interface xe-0/3/1.0 metric 21
user@R3# set area 0.0.0.0 interface ge-1/0/0.0 metric 13
user@R3# set area 0.0.0.0 interface ge-1/0/6.0 metric 15
user@R3# set area 0.0.0.0 interface xe-1/3/0.0 link-protection
user@R3# set area 0.0.0.0 interface xe-1/3/0.0 metric 22

```

10. Enable link protection and configure metric values on all the interfaces for an OSPF3 area.

```

[edit protocols ospf3]
user@R3# set area 0.0.0.0 interface ge-0/0/5.0 link-protection

```

```

user@R3# set area 0.0.0.0 interface ge-0/0/5.0 metric 10
user@R3# set area 0.0.0.0 interface xe-0/3/1.0 metric 21
user@R3# set area 0.0.0.0 interface ge-1/0/0.0 metric 13
user@R3# set area 0.0.0.0 interface ge-1/0/6.0 metric 15
user@R3# set area 0.0.0.0 interface xe-1/3/0.0 link-protection
user@R3# set area 0.0.0.0 interface xe-1/3/0.0 metric 22

```

11. Configure the routing policy.

```

[edit policy-options]
user@R3# set policy-statement ecmp term 1 then load-balance per-packet

```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R3# show interfaces
ge-0/0/5 {
  unit 0 {
    family inet {
      address 192.168.50.2/30;
    }
    family inet6 {
      address 2001:db8:50:1:1::2/64;
    }
    family mpls;
  }
}
xe-0/3/1 {
  unit 0 {
    family inet {
      address 192.168.75.1/30;
    }
    family inet6 {
      address 2001:db8:75:1:1::1/64;
    }
    family mpls;
  }
}

```

```
}
ge-1/0/0 {
  unit 0 {
    family inet {
      address 192.168.80.1/30;
    }
    family inet6 {
      address 2001:db8:80:1:1::1/64;
    }
    family mpls;
  }
}
ge-1/0/5 {
  unit 0 {
    family inet {
      address 172.16.200.1/24;
    }
    family inet6 {
      address 2001:db8:200:1:1::1/64;
    }
  }
}
ge-1/0/6 {
  unit 0 {
    family inet {
      address 192.168.85.1/30;
    }
    family inet6 {
      address 2001:db8:85:1:1::1/64;
    }
    family mpls;
  }
}
xe-1/3/0 {
  unit 0 {
    family inet {
      address 192.168.90.1/30;
    }
    family inet6 {
      address 2001:db8:90:1:1::1/64;
    }
    family mpls;
  }
}
```

```
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.3.3/32 {
        primary;
      }
    }
    family inet6 {
      address 2001:db8:3:3:3:3/128 {
        primary;
      }
    }
    family mpls;
  }
}
```

```
user@R3# show protocols
```

```
  rsvp {
    interface all;
  }
  mpls {
    admin-groups {
      c0 0;
      c1 1;
      c2 2;
      c3 3;
      c4 4;
      c5 5;
      c6 6;
      c7 7;
      c8 8;
      c9 9;
      c10 10;
      c11 11;
      c12 12;
      c13 13;
      c14 14;
      c15 15;
      c16 16;
      c17 17;
    }
  }
}
```

```
c18 18;
c19 19;
c20 20;
c21 21;
c22 22;
c23 23;
c24 24;
c25 25;
c26 26;
c27 27;
c28 28;
c29 29;
c30 30;
c31 31;
}
interface all;
interface ge-0/0/5.0 {
    admin-group c0;
}
}
ospf {
    area 0.0.0.0 {
        interface ge-0/0/5.0 {
            link-protection;
            metric 10;
        }
        interface xe-0/3/1.0 {
            metric 21;
        }
        interface ge-1/0/0.0 {
            metric 13;
        }
        interface ge-1/0/6.0 {
            metric 15;
        }
        interface xe-1/3/0.0 {
            link-protection;
            metric 22;
        }
    }
}
ospf3 {
    area 0.0.0.0 {
```

```
    interface ge-0/0/5.0 {
        link-protection;
        metric 10;
    }
    interface xe-0/3/1.0 {
        metric 21;
    }
    interface ge-1/0/0.0 {
        metric 13;
    }
    interface ge-1/0/6.0 {
        metric 15;
    }
    interface xe-1/3/0.0 {
        link-protection;
        metric 22;
    }
}
}
```

```
user@R3# show routing-options
srlg {
    srlg1 srlg-value 1001;
    srlg2 srlg-value 1002;
    srlg3 srlg-value 1003;
    srlg4 srlg-value 1004;
    srlg5 srlg-value 1005;
    srlg6 srlg-value 1006;
    srlg7 srlg-value 1007;
    srlg8 srlg-value 1008;
    srlg9 srlg-value 1009;
    srlg10 srlg-value 10010;
    srlg11 srlg-value 10011;
    srlg12 srlg-value 10012;
}
router-id 172.16.3.3;
forwarding-table {
    export ecmp;
}
backup-selection {
    destination 10.1.1.0/30 {
```

```
interface xe-1/3/0.0 {
  admin-group {
    include-all c2;
  }
}
interface all {
  admin-group {
    exclude c3;
  }
  srlg strict;
  protection-type node;
  bandwidth-greater-equal-primary;
  node {
    preference 172.16.7.7;
  }
  root-metric lowest;
  metric-order root;
}
}
destination 172.16.30.0/30 {
  interface all {
    admin-group {
      exclude c5;
    }
    srlg strict;
    protection-type node;
    bandwidth-greater-equal-primary;
    node {
      preference 172.16.7.7;
    }
    root-metric lowest;
    metric-order root;
  }
}
}
destination 192.168.45.0/30 {
  interface all {
    admin-group {
      exclude c5;
    }
    srlg strict;
    protection-type node;
    bandwidth-greater-equal-primary;
    node {
```



```
    > to 10.92.31.254 via fxp0.0
10.5.0.0/16    *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.6.128.0/17 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.9.0.0/16   *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.10.0.0/16  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.13.4.0/23  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.13.10.0/23 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.82.0.0/15  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.84.0.0/16  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.85.12.0/22 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.92.0.0/16  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.92.16.0/20 *[Direct/0] 02:22:57
    > via fxp0.0
10.92.24.195/32 *[Local/0] 02:22:57
    Local via fxp0.0
10.94.0.0/16  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.99.0.0/16  *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.102.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.150.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.155.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.157.64.0/19 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.160.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.204.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
10.205.0.0/16 *[Static/5] 02:22:57
    > to 10.92.31.254 via fxp0.0
```

```
10.206.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.207.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.209.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.212.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.213.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.214.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.215.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.216.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.218.13.0/24     *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.218.14.0/24     *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.218.16.0/20     *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.218.32.0/20     *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
10.227.0.0/16      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
172.16.50.0/30     *[Direct/0] 02:19:55
                  > via ge-0/0/5.0
172.16.50.2/32     *[Local/0] 02:19:58
                  Local via ge-0/0/5.0
172.16.75.0/30     *[Direct/0] 02:19:55
                  > via xe-0/3/1.0
172.16.75.1/32     *[Local/0] 02:19:57
                  Local via xe-0/3/1.0
172.16.24.195/32   *[Direct/0] 02:22:57
                  > via lo0.0
172.16.0.0/12      *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
192.168.0.0/16     *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
192.168.102.0/23   *[Static/5] 02:22:57
                  > to 10.92.31.254 via fxp0.0
192.168.136.0/24   *[Static/5] 02:22:57
```

```

                > to 10.92.31.254 via fxp0.0
192.168.136.192/32 *[Static/5] 02:22:57
                > to 10.92.31.254 via fxp0.0
192.168.137.0/24  *[Static/5] 02:22:57
                > to 10.92.31.254 via fxp0.0
192.168.233.5/32  *[OSPF/10] 00:16:55, metric 1
                MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1280.9202.4195/152
                *[Direct/0] 02:22:57
                > via lo0.0

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:16:55, metric 1
                Receive
1                *[MPLS/0] 00:16:55, metric 1
                Receive
2                *[MPLS/0] 00:16:55, metric 1
                Receive
13               *[MPLS/0] 00:16:55, metric 1
                Receive

inet6.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:50:1:1::/64  *[Direct/0] 02:19:44
                > via ge-0/0/5.0
2001:db8:50:1:1::2/128 *[Local/0] 02:19:58
                Local via ge-0/0/5.0
2001:db8:75:1:1::/64  *[Direct/0] 02:19:44
                > via xe-0/3/1.0
2001:db8:75:1:1::1/128 *[Local/0] 02:19:57
                Local via xe-0/3/1.0
2001:db8::3:3:3:3/128 *[Direct/0] 02:22:27
                > via lo0.0
2001:db8::128:92:24:195/128
                *[Direct/0] 02:22:57
                > via lo0.0

```

```

fe80::/64          *[Direct/0] 02:19:44
                  > via ge-0/0/5.0
                  [Direct/0] 02:19:43
                  > via xe-0/3/1.0
fe80::205:86ff:fe00:ed05/128
                  *[Local/0] 02:19:58
                  Local via ge-0/0/5.0
fe80::205:86ff:fe00:ed3d/128
                  *[Local/0] 02:19:57
                  Local via xe-0/3/1.0
fe80::5668:a50f:fcc1:3ca2/128
                  *[Direct/0] 02:22:57
                  > via lo0.0

```

Meaning

The output shows all Device R3 routes.

Verifying the OSPF Route

Purpose

Verify the routing table of OSPF.

Action

From operational mode, run the `show ospf route detail` command for Device R3.

```

user@R3> show ospf route detail
Topology default Route Table:

Prefix          Path Route   NH      Metric NextHop      Nexthop
                Type Type     Type                    Interface  Address/LSP
172.16.50.0/30  Intra Network IP        10 ge-0/0/5.0
  area 0.0.0.0, origin 172.16.3.3, priority low
172.16.75.0/30  Intra Network IP        21 xe-0/3/1.0
  area 0.0.0.0, origin 172.16.3.3, priority low

```

Meaning

The output displays the routing table of OSPF routers.

Verifying the OSPF3 Route

Purpose

Verify the routing table of OSPF3.

Action

From operational mode, run the `show ospf3 route detail` command for Device R3.

```
user@R3> show ospf3 route detail
```

Prefix	Path Type	Route Type	NH Type	Metric
2001:db8:50:1:1::/64		Intra Network	IP	10
NH-interface ge-0/0/5.0				
Area 0.0.0.0, Origin 172.16.3.3, Priority low				
2001:db8:75:1:1::/64		Intra Network	IP	21
NH-interface xe-0/3/1.0				
Area 0.0.0.0, Origin 172.16.3.3, Priority low				

Meaning

The output displays the routing table of OSPF3 routers.

Verifying the Backup Selection Policy for Device R3

Purpose

Verify the backup selection policy for Device R3.

Action

From operational mode, run the `show backup-selection` command for Device R3.

```

user@R3> show backup-selection
Prefix: 10.1.1.0/30
  Interface: all
    Admin-group exclude: c3
    Neighbor preference: 172.16.7.7
    Protection Type: Node, Downstream Paths Only: Disabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection, node, Metric
  Interface: xe-1/3/0.0
    Admin-group include-all: c2
    Protection Type: Link, Downstream Paths Only: Disabled, SRLG: Loose, B/w >= Primary:
Disabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Dest-metric, Root-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection, node, Metric Prefix:
172.16.30.0/30
  Interface: all
    Admin-group exclude: c5
    Neighbor preference: 172.16.7.7
    Protection Type: Node, Downstream Paths Only: Disabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection, node, Metric
Prefix: 172.16.45.0/30
  Interface: all
    Admin-group exclude: c5
    Neighbor preference: 172.16.7.7
    Protection Type: Node, Downstream Paths Only: Disabled, SRLG: Strict, B/w >= Primary:
Enabled, Root-metric: lowest, Dest-metric: lowest
    Metric Evaluation Order: Root-metric, Dest-metric
    Policy Evaluation Order: Admin-group, SRLG, Bandwidth, Protection, node, Metric

```

Meaning

The output displays the configured policies per prefix per primary next-hop interface.

SEE ALSO

| [backup-selection \(Protocols IS-IS\)](#)

Example: Injecting OSPF Routes into the BGP Routing Table

IN THIS SECTION

- [Requirements | 644](#)
- [Overview | 644](#)
- [Configuration | 645](#)
- [Verification | 648](#)
- [Troubleshooting | 649](#)

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

Requirements

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [Example: Configuring External BGP Point-to-Point Peer Sessions](#).
- Configure interior gateway protocol (IGP) sessions between peers.

Overview

IN THIS SECTION

- [Topology | 645](#)

In this example, you create a routing policy called `injectpolicy1` and a routing term called `injectterm1`. The policy injects OSPF routes into the BGP routing table.

Topology

Configuration

IN THIS SECTION

- [Configuring the Routing Policy | 645](#)
- [Configuring Tracing for the Routing Policy | 647](#)

Configuring the Routing Policy

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```

2. Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

3. Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```

4. Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

5. Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

Results

Confirm your configuration by entering the `show policy-options` and `show protocols bgp` commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    from {
      protocol ospf;
      area 0.0.0.1;
    }
  }
  then accept;
```

```

    }
}

```

```

user@host# show protocols bgp
export injectpolicy1;

```

If you are done configuring the device, enter `commit` from configuration mode.

Configuring Tracing for the Routing Policy

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```

set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

1. Include a trace action in the policy.

```

[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# then trace

```

2. Configure the tracing file for the output.

```

[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m

```

```
user@host# set file files 5
user@host# set flag policy
```

Results

Confirm your configuration by entering the `show policy-options` and `show routing-options` commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    then {
      trace;
    }
  }
}
```

```
user@host# show routing-options
traceoptions {
  file ospf-bgp-policy-log size 5m files 5;
  flag policy;
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That the Expected BGP Routes Are Present | 649](#)

Confirm that the configuration is working properly.

Verifying That the Expected BGP Routes Are Present

Purpose

Verify the effect of the export policy.

Action

From operational mode, enter the `show route` command.

Troubleshooting

IN THIS SECTION

- [Using the show log Command to Examine the Actions of the Routing Policy | 649](#)

Using the show log Command to Examine the Actions of the Routing Policy

Problem

The routing table contains unexpected routes, or routes are missing from the routing table.

Solution

If you configure policy tracing as shown in this example, you can run the `show log ospf-bgp-policy-log` command to diagnose problems with the routing policy. The `show log ospf-bgp-policy-log` command displays information about the routes that the `injectpolicy1` policy term analyzes and acts upon.

Example: Redistributing Static Routes into OSPF

IN THIS SECTION

- [Requirements | 650](#)
- [Overview | 650](#)

- Configuration | 651
- Verification | 653

This example shows how to create a policy that redistributes static routes into OSPF.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).

Overview

IN THIS SECTION

- Topology | 650

In this example, you create a routing policy called `exportstatic1` and a routing term called `exportstatic1`. The policy injects static routes into OSPF. This example includes the following settings:

- `policy-statement`—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens (-) and be up to 255 characters long.
- `term`—Defines the match condition and applicable actions for the routing policy. The term name can contain letters, numbers, and hyphens (-) and be up to 255 characters long. You specify the name of the term and define the criteria that an incoming route must match by including the `from` statement and the action to take if the route matches the conditions by including the `then` statement. In this example you specify the static protocol match condition and the `accept` action.
- `export`—Applies the export policy you created to be evaluated when routes are being exported from the routing table into OSPF.

Topology

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 651](#)
- [Procedure | 651](#)

CLI Quick Configuration

To quickly create a policy that injects static routes into OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set policy-options policy-statement exportstatic1 term exportstatic1 from protocol static
set policy-options policy-statement exportstatic1 term exportstatic1 then accept
set protocols ospf export exportstatic1
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To inject static routes into OSPF:

1. Create the routing policy.

```
[edit]
user@host# edit policy-options policy-statement exportstatic1
```

2. Create the policy term.

```
[edit policy-options policy-statement exportstatic1]
user@host# set term exportstatic1
```

3. Specify static as a match condition.

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set from protocol static
```

4. Specify that the route is to be accepted if the previous condition is matched.

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set then accept
```

5. Apply the routing policy to OSPF.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf export exportstatic1
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show policy-options` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement exportstatic1 {
  term exportstatic1 {
    from protocol static;
    then accept;
```

```
}  
}
```

```
user@host# show protocols ospf  
export exportstatic1;
```

To confirm your OSPFv3 configuration, enter the `show policy-options` and the `show protocols ospf3` commands.

Verification

IN THIS SECTION

- [Verifying That the Expected Static Routes Are Present | 653](#)
- [Verifying That AS External LSAs Are Added to the Routing Table | 653](#)

Confirm that the configuration is working properly.

Verifying That the Expected Static Routes Are Present

Purpose

Verify the effect of the export policy.

Action

From operational mode, enter the `show route` command.

Verifying That AS External LSAs Are Added to the Routing Table

Purpose

On the routing device where you configured the export policy, verify that the routing device originates an AS external LSA for the static routes that are added to the routing table.

Action

From operational mode, enter the `show ospf database` command for OSPFv2, and enter the `show ospf3 database` command for OSPFv3.

Example: Configuring an OSPF Import Policy

IN THIS SECTION

- [Requirements | 654](#)
- [Overview | 654](#)
- [Configuration | 655](#)
- [Verification | 659](#)

This example shows how to create an OSPF import policy. OSPF import policies apply to external routes only. An external route is a route that is outside the OSPF autonomous system (AS).

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57.
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62 .

Overview

External routes are learned by AS boundary routers. External routes can be advertised throughout the OSPF domain if you configure the AS boundary router to redistribute the route into OSPF. An external route might be learned by the AS boundary router from a routing protocol other than OSPF, or the external route might be a static route that you configure on the AS boundary router.

For OSPFv3, the link-state advertisement (LSA) is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An area border router (ABR) originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements. The filtering is done only on external routes in OSPF. The intra-area and interarea routes are not considered for filtering. The default action is to accept the route when the route does not match the policy.

This example includes the following OSPF policy settings:

- `policy-statement`—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens (-) and be up to 255 characters long.
- `export`—Applies the export policy you created to be evaluated when network summary LSAs are flooded into an area. In this example, the export policy is named `export_static`.
- `import`—Applies the import policy you created to prevent external routes from being added to the routing table. In this example, the import policy is named `filter_routes`.

The devices you configure in this example represent the following functions:

- **R1**—Device R1 is in area 0.0.0.0 and has a direct connection to device R2. R1 has an OSPF export policy configured. The export policy redistributes static routes from R1's routing table into R1's OSPF database. Because the static route is in R1's OSPF database, the route is advertised in an LSA to R1's OSPF neighbor. R1's OSPF neighbor is device R2.
- **R2**—Device R2 is in area 0.0.0.0 and has a direct connection to device R1. R2 has an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network and prevents the static route from being installed in R2's routing table. R2's OSPF neighbor is device R1.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 656](#)
- [Procedure | 656](#)

CLI Quick Configuration

To quickly configure an OSPF import policy, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
set protocols ospf export export_static
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement export_static from protocol static
set policy-options policy-statement export_static then accept
```

Configuration on Device R2:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
set protocols ospf import filter_routes
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement filter_routes from route-filter 10.0.16.0/30 exact
set policy-options policy-statement filter_routes then reject
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To configure an OSPF import policy:

1. Configure the interfaces.

```
[edit]
user@R1# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R2# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

3. On R1, redistribute the static route into OSPF.

```
[edit]
user@R1# set protocols ospf export export_static
user@R1# set policy-options policy-statement export_static from protocol static
user@R1# set policy-options policy-statement export_static then accept
```

4. On R2, configure the OSPF import policy.

```
[edit]
user@R2# set protocols ospf import filter_routes
user@R2# set policy-options policy-statement filter_routes from route-filter 10.0.16.0/30
exact
user@R2# set policy-options policy-statement filter_routes then reject
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
```

```
user@R1# show policy-options
policy-statement export_static {
  from protocol static;
  then accept;
}
```

```
user@R1# show protocols ospf
export export_static;
area 0.0.0.0 {
  interface so-0/2/0.0;
}
```

Output for R2:

```
user@R2# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.2.2/30;
    }
  }
}
```

```
user@R2# show policy-options
policy-statement filter_routes {
  from {
    route-filter 10.0.16.0/30 exact;
  }
  then reject;
}
```

```
user@R2# show protocols ospf
import filter_routes;
area 0.0.0.0 {
  interface so-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces`, `show policy-options`, `show routing-options`, and `show protocols ospf3` commands on the appropriate device.

Verification

IN THIS SECTION

- [Verifying the OSPF Database | 660](#)
- [Verifying the Routing Table | 660](#)

Confirm that the configuration is working properly.

Verifying the OSPF Database

Purpose

Verify that OSPF is advertising the static route in the OSPF database.

Action

From operational mode, enter the `show ospf database` for OSPFv2, and enter the `show ospf3 database` command for OSPFv3.

Verifying the Routing Table

Purpose

Verify the entries in the routing table.

Action

From operational mode, enter the `show route` command.

Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF

IN THIS SECTION

- [Requirements | 661](#)
- [Overview | 661](#)
- [Configuration | 662](#)
- [Verification | 666](#)

This example shows how to create an OSPF import policy that prioritizes specific prefixes learned through OSPF.

Requirements

Before you begin:

- Configure the device interfaces. See the [Interfaces User Guide for Security Devices](#).
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election See "[Example: Controlling OSPF Designated Router Election](#)" on page 57
- Configure a single-area OSPF network. See "[Example: Configuring a Single-Area OSPF Network](#)" on page 62 .
- Configure a multiarea OSPF network. See "[Example: Configuring a Multiarea OSPF Network](#)" on page 65.

Overview

IN THIS SECTION

- [Topology | 662](#)

In a network with a large number of OSPF routes, it can be useful to control the order in which routes are updated in response to a network topology change. In Junos OS Release 9.3 and later, you can specify a priority of high, medium, or low for prefixes included in an OSPF import policy. In the event of an OSPF topology change, high priority prefixes are updated in the routing table first, followed by medium and then low priority prefixes.

OSPF import policy can only be used to set priority or to filter OSPF external routes. If an OSPF import policy is applied that results in a `reject` terminating action for a nonexternal route, then the `reject` action is ignored and the route is accepted anyway. By default, such a route is now installed in the routing table with a priority of low. This behavior prevents traffic black holes, that is, silently discarded traffic, by ensuring consistent routing within the OSPF domain.

In general, OSPF routes that are not explicitly assigned a priority are treated as priority medium, except for the following:

- Summary discard routes have a default priority of low.
- Local routes that are not added to the routing table are assigned a priority of low.

- External routes that are rejected by import policy and thus not added to the routing table are assigned a priority of low.

Any available match criteria applicable to OSPF routes can be used to determine the priority. Two of the most commonly used match criteria for OSPF are the `route-filter` and `tag` statements.

In this example, the routing device is in area 0.0.0.0, with interfaces `fe-0/1/0` and `fe-1/1/0` connecting to neighboring devices. You configure an import routing policy named `ospf-import` to specify a priority for prefixes learned through OSPF. Routes associated with these prefixes are installed in the routing table in the order of the prefixes' specified priority. Routes matching `192.0.2.0/24 orlonger` are installed first because they have a priority of high. Routes matching `198.51.100.0/24 orlonger` are installed next because they have a priority of medium. Routes matching `203.0.113.0/24 orlonger` are installed last because they have a priority of low. You then apply the import policy to OSPF.



NOTE: The priority value takes effect when a new route is installed, or when there is a change to an existing route.

Topology

Configuration

IN THIS SECTION

● [CLI Quick Configuration | 662](#)

● [Procedure | 663](#)

CLI Quick Configuration

To quickly configure an OSPF import policy that prioritizes specific prefixes learned through OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.5/30
set policy-options policy-statement ospf-import term t1 from route-filter 203.0.113.0/24 orlonger
```

```

set policy-options policy-statement ospf-import term t1 then priority low
set policy-options policy-statement ospf-import term t1 then accept
set policy-options policy-statement ospf-import term t2 from route-filter 198.51.100.0/24
orlonger
set policy-options policy-statement ospf-import term t2 then priority medium
set policy-options policy-statement ospf-import term t2 then accept
set policy-options policy-statement ospf-import term t3 from route-filter 192.0.2.0/24 orlonger
set policy-options policy-statement ospf-import term t3 then priority high
set policy-options policy-statement ospf-import term t3 then accept
set protocols ospf import ospf-import
set protocols ospf area 0.0.0.0 interface fe-0/1/0
set protocols ospf area 0.0.0.0 interface fe-1/1/0

```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To configure an OSPF import policy that prioritizes specific prefixes:

1. Configure the interfaces.

```

[edit]
user@host# set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
user@host# set interfaces fe-0/2/0 unit 0 family inet address 192.168.8.5/30

```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```

[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-0/1/0
user@host# set protocols ospf area 0.0.0.0 interface fe-0/2/0

```

3. Configure the policy to specify the priority for prefixes learned through OSPF.

```
[edit ]
user@host# set policy-options policy-statement ospf-import term t1 from route-filter
203.0.113.0/24 orlonger
user@host# set policy-options policy-statement ospf-import term t1 then priority low
user@host# set policy-options policy-statement ospf-import term t1 then accept
user@host# set policy-options policy-statement ospf-import term t2 from route-filter
198.51.100.0/24 orlonger
user@host# set policy-options policy-statement ospf-import term t2 then priority medium
user@host# set policy-options policy-statement ospf-import term t2 then accept
user@host# set policy-options policy-statement ospf-import term t3 from route-filter
192.0.2.0/24 orlonger
user@host# set policy-options policy-statement ospf-import term t3 then priority high
user@host# set policy-options policy-statement ospf-import term t3 then accept
```

4. Apply the policy to OSPF.

```
[edit]
user@host# set protocols ospf import ospf-import
```

5. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show policy-options`, and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 192.168.8.4/30;
    }
  }
}
```

```
}
  fe-0/2/0 {
    unit 0 {
      family inet {
        address 192.168.8.5/30;
      }
    }
  }
}
```

```
user@host# show protocols ospf
import ospf-import;
area 0.0.0.0 {
  interface fe-0/1/0.0;
  interface fe-0/2/0.0;
}
```

```
user@host# show policy-options
policy-statement ospf-import {
  term t1 {
    from {
      route-filter 203.0.113.0/24 orlonger;
    }
    then {
      priority low;
      accept;
    }
  }
  term t2 {
    from {
      route-filter 198.51.100.0/24 orlonger;
    }
    then {
      priority medium;
      accept;
    }
  }
  term t3 {
    from {
      route-filter 192.0.2.0/24 orlonger;
    }
  }
}
```

```
        then {  
            priority high;  
            accept;  
        }  
    }  
}
```

```
user@host# show protocols ospf  
import ospf-import;  
area 0.0.0.0 {  
    interface fe-0/1/0.0;  
    interface fe-0/2/0.0;  
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces`, `show policy-options`, and `show protocols ospf3` commands.

Verification

IN THIS SECTION

- [Verifying the Prefix Priority in the OSPF Routing Table | 666](#)

Confirm that the configuration is working properly.

Verifying the Prefix Priority in the OSPF Routing Table

Purpose

Verify the priority assigned to the prefix in the OSPF routing table.

Action

From operational mode, enter the `show ospf route detail` for OSPFv2, and enter the `show ospf3 route detail` command for OSPFv3.

Import and Export Policies for Network Summaries Overview

By default, OSPF uses network-summary link-state advertisements (LSAs) to transmit route information across area boundaries. Each area border router (ABR) floods network-summary LSAs to other routing devices in the same area. The ABR also controls which routes from the area are used to generate network-summary LSAs into other areas. Each ABR maintains a separate topological database for each area to which they are connected. In Junos OS Release 9.1 and later, you can configure export and import policies for OSPFv2 and OSPFv3 that enable you to control how network-summary LSAs, which contain information about interarea OSPF prefixes, are distributed and generated. For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

The export policy enables you to specify which summary LSAs are flooded into an area. The import policy enables you to control which routes learned from an area are used to generate summary LSAs into other areas. You define a routing policy at the [edit policy-options policy-statement *policy-name*] hierarchy level. As with all OSPF export policies, the default for network-summary LSA export policies is to reject everything. Similarly, as with all OSPF import policies, the default for network-summary LSA import policies is to accept all OSPF routes.

Example: Configuring an OSPF Export Policy for Network Summaries

IN THIS SECTION

- [Requirements | 667](#)
- [Overview | 668](#)
- [Configuration | 670](#)
- [Verification | 679](#)

This example shows how to create an OSPF export policy to control the network-summary (Type 3) LSAs that the ABR floods into an OSPF area.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See ["Example: Configuring an OSPF Router Identifier" on page 54](#).
- Control OSPF designated router election. See ["Example: Controlling OSPF Designated Router Election" on page 57](#)

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.




NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to export network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

[Figure 38 on page 669](#) shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 38: Sample Topology Used for an OSPF Export Network Summary Policy



In this example, you configure R4 with an export network summary policy named `export-policy` that only allows routes that match the `10.0.4.4` prefix from area 3 into area 4. The export policy controls the network-summary LSAs that R4 floods into area 4. This results in only the allowed interarea route to enter area 4, and all other interarea routes to be purged from the OSPF database and the routing table of the devices in area 4. You first define the policy and then apply it to the ABR by including the `network-summary-export` statement for OSPFv2 or the `inter-area-prefix-export` statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface `fe-0/1/0` has an IP address of `10.0.4.13/30` and connects to R3. Interface `fe-0/0/1` has an IP address of `10.0.4.5/30` and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface `fe-0/0/1` has an IP address of `10.0.4.6/30` and connects to R1. Interface `fe-1/0/0` has an IP address of `10.0.4.1` and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface `fe-1/0/0` has an IP address of `10.0.4.2/30` and connects to R2. Interface `fe-1/1/0` has an IP address of `10.0.4.14/30` and connects to R1. Interface `fe-0/0/1` has an IP address of `10.0.2.1/30` and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface `fe-0/0/1` has an IP address of `10.0.2.4/30` and connects to R3. Interface `fe-1/1/0` has an IP address of `10.0.8.6/30` and connects to R5. Interface `fe-1/0/0` has an IP address of `10.0.8.9/30` and connects to R6.

- R5—Device R5 is an internal router in area 4. Interface fe-1/1/0 has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface fe-1/0/0 has an IP address of 10.0.8.10/30 and connects to R4.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 670](#)
- [Procedure | 672](#)

CLI Quick Configuration

To quickly configure an OSPF export policy for network summaries, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set policy-options policy-statement export-policy term term1 from route-filter 10.0.4.4/30
prefix-length-range /30-/30
set policy-options policy-statement export-policy term term1 then accept
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0
set protocols ospf area 0.0.0.4 network-summary-export export-policy
```

Configuration on Device R5:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface fe-0/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To configure an configure an OSPF export policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
```

```
[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
```

```
[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3#set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
```

```
[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
```

```
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

```
[edit]
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

3. On R4, configure the export network summary policy.

```
[edit ]
user@R4# set policy-options policy-statement export-policy term term1 from route-filter
10.0.4.4/30 prefix-length-range /30-/30
user@R4# set policy-options policy-statement export-policy term term1 then accept
```

4. On R4, apply the export network summary policy to OSPF.



NOTE: For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```
[edit]
user@R4# set protocols ospf area 0.0.0.4 network-summary-export export-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}
```

```
user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}
```

Output for R2:

```
user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.3/30;
    }
  }
}
```

```
    }  
  }  
}
```

```
user@R2# show protocols ospf  
area 0.0.0.3 {  
    interface fe-0/1/0.0;  
    interface fe-1/0/0.0;  
}
```

Output for R3:

```
user@R3# show interfaces  
fe-0/0/1 {  
    unit 0 {  
        family inet {  
            address 10.0.2.3/30;  
        }  
    }  
}  
fe-1/0/0 {  
    unit 0 {  
        family inet {  
            address 10.0.4.2/30;  
        }  
    }  
}  
fe-1/1/0 {  
    unit 0 {  
        family inet {  
            address 10.0.4.14/30;  
        }  
    }  
}
```

```
user@R3# show protocols ospf  
area 0.0.0.0 {  
    interface fe-0/0/1.0;  
}  
area 0.0.0.3 {
```

```
interface fe-1/0/0.0;  
interface fe-1/1/0.0;  
}
```

Output for R4:

```
user@R4# show interfaces  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.2.4/30;  
    }  
  }  
}  
fe-1/0/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.6/30;  
    }  
  }  
}  
fe-1/1/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.3/30;  
    }  
  }  
}
```

```
user@R4# show protocols ospf  
area 0.0.0.0 {  
  interface fe-0/0/1.0;  
}  
area 0.0.0.4 {  
  network-summary-export export-policy;  
  interface fe-1/0/0.0;
```

```
interface fe-1/1/0.0;  
}
```

```
user@R4# show policy-options  
policy-statement export-policy {  
  term term1 {  
    from {  
      route-filter 10.0.4.4/30 prefix-length-range /30-/30;  
    }  
    then accept;  
  }  
}
```

Output for R5:

```
user@R5# show interfaces  
fe-1/1/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.5/30;  
    }  
  }  
}
```

```
user@R5# show protocols ospf  
area 0.0.0.4 {  
  interface fe-1/1/0.0;  
}
```

Output for R6:

```
user@R6# show interfaces  
fe-1/0/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.7/30;  
    }  
  }  
}
```

```
}  
}
```

```
user@R6# show protocols ospf  
area 0.0.0.4 {  
    interface fe-1/0/0.0;  
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces`, `show policy-options`, and `show protocols ospf3` commands on the appropriate device.

Verification

IN THIS SECTION

- [Verifying the OSPF Database | 679](#)
- [Verifying the Routing Table | 680](#)

Confirm that the configuration is working properly.

Verifying the OSPF Database

Purpose

Verify that the OSPF database for the devices in area 4 includes the interarea route that we permitted on the ABR R4. The other interarea routes that are not specified should age out or no longer be present in the OSPF database.

Action

From operational mode, enter the `show ospf database netsummary area 0.0.0.4` command for OSPFv2, and enter the `show ospf3 database inter-area-prefix area 0.0.0.4` command for OSPFv3.

Verifying the Routing Table

Purpose

Verify that the routes corresponding to the rejected network summaries are no longer present in R4's, R5's, or R6's routing table.

Action

From operational mode, enter the `show route protocol ospf` command for both OSPFv2 and OSPFv3.

Example: Configuring an OSPF Import Policy for Network Summaries

IN THIS SECTION

- [Requirements | 680](#)
- [Overview | 680](#)
- [Configuration | 682](#)
- [Verification | 692](#)

This example shows how to create an OSPF import policy to control the network-summary (Type 3) LSAs that the ABR advertises out of an OSPF area.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 54.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 57.

Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to

advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.


NOTE: For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to import network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the `route-filter` policy match condition for these types of policies.
- We recommend that you use the `accept` and `reject` routing policy terms for these types of policies.

Figure 39 on page 681 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 39: Sample Topology Used for an OSPF Import Network Summary Policy



In this example, you configure R3 with an import network summary policy named `import-policy` so R3 only generates network summaries for the route `10.0.4.12/30`. The import policy controls the routes and therefore the network summaries that R3 advertises out of area 3, so applying this policy means that R3 only advertises route `10.0.4.12/30` out of area 3. This results in existing network summaries from other interarea routes getting purged from the OSPF database in area 0 and area 4, as well as the routing tables of the devices in areas 0 and area 4. You first define the policy and then apply it to the ABR by including the `network-summary-import` statement for OSPFv2 or the `inter-area-prefix-import` statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface `ge-0/1/0` has an IP address of `10.0.4.13/30` and connects to R3. Interface `ge-0/0/1` has an IP address of `10.0.4.5/30` and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface `ge-0/0/1` has an IP address of `10.0.4.6/30` and connects to R1. Interface `ge-1/0/0` has an IP address of `10.0.4.1/30` and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface `ge-1/0/0` has an IP address of `10.0.4.2/30` and connects to R2. Interface `ge-1/1/0` has an IP address of `10.0.4.14/30` and connects to R1. Interface `ge-0/0/1` has an IP address of `10.0.2.1/30` and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface `ge-0/0/1` has an IP address of `10.0.2.1/30` and connects to R3. Interface `ge-1/1/0` has an IP address of `10.0.8.6/30` and connects to R5. Interface `ge-1/0/0` has an IP address of `10.0.8.9/30` and connects to R6.
- R5—Device R5 is an internal router in area 4. Interface `ge-1/1/0` has an IP address of `10.0.8.5/30` and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface `ge-1/0/0` has an IP address of `10.0.8.10/30` and connects to R4.

Configuration

IN THIS SECTION

- [Procedure | 683](#)

Procedure

CLI Quick Configuration

To quickly configure an OSPF import policy for network summaries, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

Configuration on Device R1:

```
[edit]
set interfaces ge-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface ge-0/1/0
set protocols ospf area 0.0.0.3 interface ge-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces ge-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces ge-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface ge-0/1/0
set protocols ospf area 0.0.0.3 interface ge-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces ge-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces ge-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.1/30
set policy-options policy-statement import-policy term term1 from route-filter 10.0.4.12/30
prefix-length-range /30-/30
set policy-options policy-statement import-policy term term1 then accept
set policy-options policy-statement import-policy term term2 then reject
set protocols ospf area 0.0.0.3 interface ge-1/0/0
set protocols ospf area 0.0.0.3 interface ge-1/1/0
set protocols ospf area 0.0.0.0 interface ge-0/0/1
set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

Configuration on Device R4:

```
[edit]
set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces ge-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces ge-1/0/0 unit 0 family inet address 10.0.8.9/30
set protocols ospf area 0.0.0.0 interface ge-0/0/1
set protocols ospf area 0.0.0.4 interface ge-1/1/0
set protocols ospf area 0.0.0.4 interface ge-1/0/0
```

Configuration on Device R5:

```
[edit]
set interfaces ge-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface ge-1/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces ge-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface ge-1/0/0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To configure an OSPF import policy for network summaries:

1. Configure the interfaces.



NOTE: For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces ge-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.5/30
```

```
[edit]
user@R2# set interfaces ge-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces ge-1/0/0 unit 0 family inet address 10.0.4.1/30
```

```
[edit]
user@R3# set interfaces ge-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces ge-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3#set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R4# set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces ge-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces ge-1/0/0 unit 0 family inet address 10.0.8.9/30
```

```
[edit]
user@R5# set interfaces ge-1/1/0 unit 0 family inet address 10.0.8.5/30
```

```
[edit]
user@R6# set interfaces ge-1/0/0 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



NOTE: For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface ge-0/1/0
user@R1# set protocols ospf area 0.0.0.3 interface ge-0/0/1
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.3 interface ge-0/1/0
user@R2# set protocols ospf area 0.0.0.3 interface ge-1/0/0
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 interface ge-1/0/0
user@R3# set protocols ospf area 0.0.0.3 interface ge-1/1/0
user@R3# set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface ge-0/0/1
user@R4# set protocols ospf area 0.0.0.4 interface ge-1/1/0
user@R4# set protocols ospf area 0.0.0.4 interface ge-1/0/0
```

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 interface ge-1/1/0
```

```
[edit]
user@R6# set protocols ospf area 0.0.0.4 interface ge-1/0/0
```

3. On R3, configure the import network summary policy.

```
[edit ]
user@R3# set policy-options policy-statement import-policy term term1 from route-filter
10.0.4.12/30 prefix-length-range /30-/30
user@R3# set policy-options policy-statement import-policy term term1 then accept
```

4. On R3, apply the import network summary policy to OSPF.



NOTE: For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfacesge-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
ge-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}
```

```
    }  
}
```

```
user@R1# show protocols ospf  
area 0.0.0.3 {  
    interface ge-0/1/0.0;  
    interface ge-0/0/1.0;  
}
```

Output for R2:

```
user@R2# show interfaces  
ge-0/1/0 {  
    unit 0 {  
        family inet {  
            address 10.0.4.6/30;  
        }  
    }  
}  
ge-1/0/0 {  
    unit 0 {  
        family inet {  
            address 10.0.4.1/30;  
        }  
    }  
}
```

```
user@R2# show protocols ospf  
area 0.0.0.3 {  
    interface ge-0/1/0.0;  
    interface ge-1/0/0.0;  
}
```

Output for R3:

```
user@R3# show interfaces  
ge-0/0/1 {  
    unit 0 {  
        family inet {
```

```

        address 10.0.2.1/30;
    }
}
ge-1/0/0 {
    unit 0 {
        family inet {
            address 10.0.4.2/30;
        }
    }
}
ge-1/1/0 {
    unit 0 {
        family inet {
            address 10.0.4.14/30;
        }
    }
}
}

```

```

user@R3# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/1.0;
}
area 0.0.0.3 {
    network-summary-import import policy;
    interface ge-1/0/0.0;
    interface ge-1/1/0.0;
}

```

```

user@R3# show policy-options
policy-statement import-policy {
    term term1 {
        from {
            route-filter 10.0.4.12/30 prefix-length-range /30-/30;
        }
        then accept;
    }
    term term2 {
        then reject;
    }
}

```

```
}  
}
```

Output for R4:

```
user@R4# show interfaces  
ge-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.2.1/30;  
    }  
  }  
}  
ge-1/0/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.9/30;  
    }  
  }  
}  
ge-1/1/0 {  
  unit 0 {  
    family inet {  
      address 10.0.8.6/30;  
    }  
  }  
}
```

```
user@R4# show protocols ospf  
area 0.0.0.0 {  
  interface ge-0/0/1.0;  
}  
area 0.0.0.4 {  
  interface ge-0/1/0.0;  
  interface ge-1/0/0.0;  
}
```

Output for R5:

```
user@R5# show interfaces
ge-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}
```

```
user@R5# show protocols ospf
area 0.0.0.4 {
  interface ge-1/1/0.0;
}
```

Output for R6:

```
user@R6# show interfaces
ge-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.10/30;
    }
  }
}
```

```
user@R6# show protocols ospf
area 0.0.0.4 {
  interface ge-1/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the `show interfaces`, `show policy-options`, and `show protocols ospf3` commands on the appropriate device.

Verification

IN THIS SECTION

- [Verifying the OSPF Database | 692](#)
- [Verifying the Routing Table | 692](#)

Confirm that the configuration is working properly.

Verifying the OSPF Database

Purpose

Verify that the OSPF database for the devices in area 4 includes the interarea route that we are advertising from R3. Any other routes from area 3 should not be advertised into area 4, so those entries should age out or no longer be present in the OSPF database.

Action

From operational mode, enter the `show ospf database netsummary area 0.0.0.4` command for OSPFv2, and enter the `show ospf3 database inter-area-prefix area 0.0.0.4` command for OSPFv3.

Verifying the Routing Table

Purpose

Verify that the specified route is included in R4's, R5's, or R6's routing table. Any other routes from area 3 should not be advertised into area 4.

Action

From operational mode, enter the `show route protocol ospf` command for both OSPFv2 and OSPFv3.

Example: Redistributing OSPF Routes into IS-IS

IN THIS SECTION

- Requirements | 693
- Overview | 693
- Configuration | 694
- Verification | 702

This example shows how to redistribute OSPF routes into an IS-IS network.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

IN THIS SECTION

- Topology | 694

Export policy can be applied to IS-IS to facilitate route redistribution.


Junos OS does not support the application of import policy for link-state routing protocols like IS-IS because such policies can lead to inconsistent link-state database (LSDB) entries, which in turn can result in routing inconsistencies.

In this example, OSPF routes 192.168.0/24 through 192.168.3/24 are redistributed into IS-IS area 49.0002 from Device R2.

In addition, policies are configured to ensure that Device R1 can reach destinations on the 10.0.0.44/30 network, and that Device R3 can reach destinations on the 10.0.0.36/30 network. This enables end-to-end reachability.

[Figure 40 on page 694](#) shows the topology used in this example.

Figure 40: IS-IS Route Redistribution Topology



"CLI Quick Configuration" on page 695 shows the configuration for all of the devices in Figure 40 on page 694. The section "No Link Title" on page 696 describes the steps on Device R2. "No Link Title" on page 698 describes the steps on Device R3.

Topology

Configuration

IN THIS SECTION

- Procedure | 695

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 0 description to-R7
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 172.16.3.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0305.00
set protocols isis interface fe-1/2/0.0
set protocols isis interface lo0.0
```

Device R2

```
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/0 unit 0 description to-OSPF-network
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.45/30
set interfaces lo0 unit 0 family inet address 172.16.9.7/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
set protocols isis export ospf-isis
set protocols isis export send-direct-to-isis-neighbors
set protocols isis interface fe-1/2/1.0
set protocols isis interface lo0.0
set protocols ospf export send-direct-to-ospf-neighbors
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf-isis term 1 from protocol ospf
set policy-options policy-statement ospf-isis term 1 from route-filter 192.168.0.0/22 longer
set policy-options policy-statement ospf-isis term 1 then accept
set policy-options policy-statement send-direct-to-isis-neighbors from protocol direct
set policy-options policy-statement send-direct-to-isis-neighbors from route-filter 10.0.0.44/30
exact
set policy-options policy-statement send-direct-to-isis-neighbors then accept
set policy-options policy-statement send-direct-to-ospf-neighbors from protocol direct
```

```
set policy-options policy-statement send-direct-to-ospf-neighbors from route-filter 10.0.0.36/30
exact
set policy-options policy-statement send-direct-to-ospf-neighbors then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.46/30
set interfaces lo0 unit 0 family inet address 192.168.1.1/32
set interfaces lo0 unit 0 family inet address 192.168.2.1/32
set interfaces lo0 unit 0 family inet address 192.168.3.1/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf export ospf
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf term 1 from protocol static
set policy-options policy-statement ospf term 1 then accept
set routing-options static route 192.168.0.0/24 discard
set routing-options static route 192.168.1.0/24 discard
set routing-options static route 192.168.3.0/24 discard
set routing-options static route 192.168.2.0/24 discard
```

Step-by-Step Procedure

To configure Device R2:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/1 unit 0 description to-R5
user@R2# set fe-1/2/1 unit 0 family inet address 10.0.0.37/30
user@R2# set fe-1/2/1 unit 0 family iso
user@R2# set fe-1/2/0 unit 0 description to-OSPF-network
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.45/30
user@R2# set lo0 unit 0 family inet address 172.16.9.7/32
user@R2# set lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
```

2. Configure IS-IS on the interface facing Device R1 and the loopback interface.

```
[edit protocols isis]
user@R2# set interface fe-1/2/1.0
user@R2# set interface lo0.0
```

3. Configure the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit policy-options policy-statement send-direct-to-isis-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.44/30 exact
user@R2# set then accept
```

4. Apply the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit protocols isis]
user@R2# set export send-direct-to-isis-neighbors
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R2# set area 0.0.0.1 interface fe-1/2/0.0
user@R2# set area 0.0.0.1 interface lo0.0 passive
```

6. Configure the OSPF route redistribution policy.

```
[edit policy-options policy-statement ospf-isis term 1]
user@R2# set from protocol ospf
user@R2# set from route-filter 192.168.0.0/22 longer
user@R2# set then accept
```

7. Apply the OSPF route redistribution policy to the IS-IS instance.

```
[edit protocols isis]
user@R2# set export ospf-isis
```

- Configure the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit policy-options policy-statement send-direct-to-ospf-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.36/30 exact
user@R2# set then accept
```

- Apply the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit protocols ospf]
user@R2# set export send-direct-to-ospf-neighbors
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure multi-level IS-IS:

- Configure the network interfaces.

Multiple addresses are configured on the loopback interface to simulate multiple route destinations.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.46/30
user@R3# set lo0 unit 0 family inet address 192.168.1.1/32
user@R3# set lo0 unit 0 family inet address 192.168.2.1/32
user@R3# set lo0 unit 0 family inet address 192.168.3.1/32
user@R3# set lo0 unit 0 family inet address 192.168.0.1/32
```

- Configure static routes to the loopback interface addresses.

These are the routes that are redistributed into IS-IS.

```
[edit routing-options static]
user@R3# set route 192.168.0.0/24 discard
user@R3# set route 192.168.1.0/24 discard
```

```

user@R3# set route 192.168.3.0/24 discard
user@R3# set route 192.168.2.0/24 discard

```

3. Configure OSPF on the interfaces.

```

[edit protocols ospf area 0.0.0.1]
user@R3# set interface fe-1/2/0.0
user@R3# set interface lo0.0 passive

```

4. Configure the OSPF policy to export the static routes.

```

[edit policy-options policy-statement ospf term 1]
user@R3# set from protocol static
user@R3# set then accept

```

5. Apply the OSPF export policy.

```

[edit protocols ospf]
user@R3# set export ospf

```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device R2

```

user@R2# show interfaces
fe-1/2/1 {
  unit 0 {
    description to-R5;
    family inet {
      address 10.0.0.37/30;
    }
    family iso;
  }
}
fe-1/2/0 {

```

```

unit 0 {
    description to-OSPF-network;
    family inet {
        address 10.0.0.45/30;
    }
}
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.9.7/32;
        }
        family iso {
            address 49.0002.0172.0016.0907.00;
        }
    }
}
}

```

```

user@R2# show protocols
isis {
    export [ ospf-isis send-direct-to-isis-neighbors ];
    interface fe-1/2/1.0;
    interface lo0.0;
}
ospf {
    export send-direct-to-ospf-neighbors;
    area 0.0.0.1 {
        interface fe-1/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
}

```

```

user@R2# show policy-options
policy-statement ospf-isis {
    term 1 {
        from {
            protocol ospf;
            route-filter 192.168.0.0/22 longer;

```

```
    }
    then accept;
  }
}
policy-statement send-direct-to-isis-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.44/30 exact;
  }
  then accept;
}
policy-statement send-direct-to-ospf-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.36/30 exact;
  }
  then accept;
}
```

Device R3

```
user@R3# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.46/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.1.1/32;
      address 192.168.2.1/32;
      address 192.168.3.1/32;
      address 192.168.0.1/32;
    }
  }
}
```

```
}  
}
```

```
user@R3# show protocols  
ospf {  
  export ospf;  
  area 0.0.0.1 {  
    interface fe-1/2/0.0;  
    interface lo0.0 {  
      passive;  
    }  
  }  
}
```

```
user@R3# show policy-options  
policy-statement ospf {  
  term 1 {  
    from protocol static;  
    then accept;  
  }  
}
```

```
user@R3# show routing-options  
static {  
  route 192.168.0.0/24 discard;  
  route 192.168.1.0/24 discard;  
  route 192.168.3.0/24 discard;  
  route 192.168.2.0/24 discard;  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

 [Verifying OSPF Route Advertisement | 703](#)

- [Verifying Route Redistribution | 704](#)
- [Verifying Connectivity | 705](#)

Confirm that the configuration is working properly.

Verifying OSPF Route Advertisement

Purpose

Make sure that the expected routes are advertised by OSPF.

Action

From operational mode on Device R2, enter the `show route protocol ospf` command.

```
user@R2> show route protocol ospf

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/24    *[OSPF/150] 03:54:21, metric 0, tag 0
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.0.1/32  *[OSPF/10] 03:54:21, metric 1
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.1.0/24  *[OSPF/150] 03:54:21, metric 0, tag 0
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.1.1/32  *[OSPF/10] 03:54:21, metric 1
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.2.0/24  *[OSPF/150] 03:54:21, metric 0, tag 0
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.2.1/32  *[OSPF/10] 03:54:21, metric 1
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.3.0/24  *[OSPF/150] 03:54:21, metric 0, tag 0
                 > to 10.0.0.46 via fe-1/2/0.0
192.168.3.1/32  *[OSPF/10] 03:54:21, metric 1
                 > to 10.0.0.46 via fe-1/2/0.0
224.0.0.5/32    *[OSPF/10] 03:56:03, metric 1
```

MultiRecv

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning

The 192.168/16 routes are advertised by OSPF.

Verifying Route Redistribution**Purpose**

Make sure that the expected routes are redistributed from OSPF into IS-IS.

Action

From operational mode on Device R1, enter the `show route protocol isis` command.

```
user@R1> show route protocol isis

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.44/30      *[IS-IS/160] 03:45:24, metric 20
                  > to 10.0.0.37 via fe-1/2/0.0
172.16.9.7/32    *[IS-IS/15] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.0.0/24   *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.0.1/32   *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.1.0/24   *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.1.1/32   *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.2.0/24   *[IS-IS/160] 03:49:46, metric 10
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.2.1/32   *[IS-IS/160] 03:49:46, metric 11, tag2 1
                  > to 10.0.0.37 via fe-1/2/0.0
192.168.3.0/24   *[IS-IS/160] 03:49:46, metric 10
```

```
                > to 10.0.0.37 via fe-1/2/0.0
192.168.3.1/32   *[IS-IS/160] 03:49:46, metric 11, tag2 1
                > to 10.0.0.37 via fe-1/2/0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning

The 192.168/16 routes are redistributed into IS-IS.

Verifying Connectivity

Purpose

Check that Device R1 can reach the destinations on Device R3.

Action

From operational mode, enter the ping command.

```
user@R1> ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=63 time=2.089 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=63 time=1.270 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=63 time=2.135 ms
```

Meaning

These results confirm that Device R1 can reach the destinations in the OSPF network.

RELATED DOCUMENTATION

[OSPF Routing Policy Overview](#)

[Understanding Route Filters for Use in Routing Policy Match Conditions](#)

15

CHAPTER

Configure OSPFv2 Sham Links

IN THIS CHAPTER

- [Configuring OSPFv2 Sham Links | 707](#)
-

Configuring OSPFv2 Sham Links

IN THIS SECTION

- [OSPFv2 Sham Links Overview | 707](#)
- [Example: Configuring OSPFv2 Sham Links | 709](#)

OSPFv2 Sham Links Overview


You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is because the backup link is considered an intra-area link, while the VPN backbone is always considered an interarea link. Intra-area links are always preferred over interarea links.

The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.

The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.

Each sham link is identified by the combination of a local endpoint address and a remote endpoint address. [Figure 41 on page 708](#) shows an OSPFv2 sham link. Router CE1 and Router CE2 are located in the same OSPFv2 area. These customer edge (CE) routing devices are linked together by a Layer 3 VPN over Router PE1 and Router PE2. In addition, Router CE1 and Router CE2 are connected by an intra-area link used as a backup.

Figure 41: OSPFv2 Sham Link



OSPFv2 treats the link through the Layer 3 VPN as an interarea link. By default, OSPFv2 prefers intra-area links to interarea links, so OSPFv2 selects the backup intra-area link as the active path. This is not acceptable in a configuration where the intra-area link is not the expected primary path for traffic between the CE routing devices. You can configure the metric for the sham link to ensure that the path over the Layer 3 VPN is preferred to a backup path over an intra-area link connecting the CE routing devices.

For the remote endpoint, you can configure the OSPFv2 interface as a demand circuit, configure IPsec authentication (you configure the actual IPsec authentication separately), and define the metric value.

You should configure an OSPFv2 sham link under the following circumstances:

- Two CE routing devices are linked together by a Layer 3 VPN.
- These CE routing devices are in the same OSPFv2 area.
- An intra-area link is configured between the two CE routing devices.

If there is no intra-area link between the CE routing devices, you do not need to configure an OSPFv2 sham link.



NOTE: In Junos OS Release 9.6 and later, an OSPFv2 sham link is installed in the routing table as a hidden route. Additionally, a BGP route is not exported to OSPFv2 if a corresponding OSPF sham link is available.



NOTE: In Junos OS Release 16.1 and later, OSPF sham-links are supported on default instances. The cost of the sham-link is dynamically set to the aigp-metric of the BGP route if no metric is configured on the sham-link by the user. If the aigp-metric is not present in the BGP route then the sham-link cost defaults to 1.

Example: Configuring OSPFv2 Sham Links

IN THIS SECTION

- [Requirements | 709](#)
- [Overview | 709](#)
- [Configuration | 710](#)
- [Verification | 718](#)

This example shows how to enable OSPFv2 sham links on a PE routing device.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

IN THIS SECTION

- [Topology | 710](#)

The sham link is an unnumbered point-to-point intra-area link and is advertised by means of a type 1 link-state advertisement (LSA). Sham links are valid only for routing instances and OSPFv2.

Each sham link is identified by a combination of the local endpoint address and a remote endpoint address and the OSPFv2 area to which it belongs. You manually configure the sham link between two PE devices, both of which are within the same VPN routing and forwarding (VRF) routing instance, and you specify the address for the local end point of the sham link. This address is used as the source for the sham link packets and is also used by the remote PE routing device as the sham link remote end point. You can also include the optional `metric` option to set a metric value for the remote end point. The metric value specifies the cost of using the link. Routes with lower total path metrics are preferred over those with higher path metrics.

To enable OSPFv2 sham links on a PE routing device:


- Configure an extra loopback interface on the PE routing device.

- Configure the VRF routing instance that supports Layer 3 VPNs on the PE routing device, and associate the sham link with an existing OSPF area. The OSPFv2 sham link configuration is also included in the routing instance. You configure the sham link's local endpoint address, which is the loopback address of the local VPN, and the remote endpoint address, which is the loopback address of the remote VPN. In this example, the VRF routing instance is named red.

Figure 42 on page 710 shows an OSPFv2 sham link.

Topology

Figure 42: OSPFv2 Sham Link Example



The devices in the figure represent the following functions:

- CE1 and CE2 are the customer edge devices.
- PE1 and PE2 are the provider edge devices.
- P is the provider device.

"CLI Quick Configuration" on page 711 shows the configuration for all of the devices in Figure 42 on page 710. The section "Step-by-Step Procedure" on page 713 describes the steps on Device PE1.

Configuration

IN THIS SECTION

- Procedure | 711

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

CE1

```
set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces lo0 unit 0 family inet address 192.0.2.1/24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0 metric 100
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
set routing-options router-id 192.0.2.1
set routing-options autonomous-system 1
```

PE1

```
set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.5/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set interfaces lo0 unit 1 family inet address 198.51.100.2/24
set protocols mpls interface fe-1/2/1.0
set protocols bgp group toR4 type internal
set protocols bgp group toR4 local-address 192.0.2.2
set protocols bgp group toR4 family inet-vpn unicast
set protocols bgp group toR4 neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
set policy-options policy-statement bgp-to-ospf term 2 then reject
```

```

set routing-instances red instance-type vrf
set routing-instances red interface fe-1/2/0.0
set routing-instances red interface lo0.1
set routing-instances red route-distinguisher 2:1
set routing-instances red vrf-target target:2:1
set routing-instances red protocols ospf export bgp-to-ospf
set routing-instances red protocols ospf sham-link local 198.51.100.2
set routing-instances red protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.4 metric 10
set routing-instances red protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set routing-instances red protocols ospf area 0.0.0.0 interface lo0.1
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 2

```

P

```

set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.6/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.9/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 3 family inet address 192.0.2.3/24
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface all
set routing-options router-id 192.0.2.3

```

PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.10/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.13/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.4/32
set interfaces lo0 unit 1 family inet address 198.51.100.4/32
set protocols mpls interface fe-1/2/0.0
set protocols bgp group toR2 type internal
set protocols bgp group toR2 local-address 192.0.2.4
set protocols bgp group toR2 family inet-vpn unicast
set protocols bgp group toR2 neighbor 192.0.2.2
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0

```

```

set protocols ldp interface fe-1/2/0.0
set protocols ldp interface lo0.0
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
set policy-options policy-statement bgp-to-ospf term 2 then reject
set routing-instances red instance-type vrf
set routing-instances red interface fe-1/2/1.0
set routing-instances red interface lo0.1
set routing-instances red route-distinguisher 2:1
set routing-instances red vrf-target target:2:1
set routing-instances red protocols ospf export bgp-to-ospf
set routing-instances red protocols ospf sham-link local 198.51.100.4
set routing-instances red protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.2 metric 10
set routing-instances red protocols ospf area 0.0.0.0 interface fe-1/2/1.0
set routing-instances red protocols ospf area 0.0.0.0 interface lo0.1
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 2

```

CE2

```

set interfaces fe-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces fe-1/2/0 unit 14 family mpls
set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30
set interfaces lo0 unit 5 family inet address 192.0.2.5/24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.14
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 3

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in [CLI User Guide](#).

To configure OSPFv2 sham links on each PE device:

1. Configure the interfaces, including two loopback interfaces.

```
[edit interfaces]
user@PE1# set fe-1/2/0 unit 0 family inet address 10.1.1.2/30
user@PE1# set fe-1/2/0 unit 0 family mpls
user@PE1# set fe-1/2/1 unit 0 family inet address 10.1.1.5/30
user@PE1# set fe-1/2/1 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 192.0.2.2/24
user@PE1# set lo0 unit 1 family inet address 198.51.100.2/24
```

2. Configure MPLS on the core-facing interface.

```
[edit protocols mpls]
user@PE1# set interface fe-1/2/1.0
```

3. Configure internal BGP (IBGP).

```
[edit ]
user@PE1# set protocols bgp group toR4 type internal
user@PE1# set protocols bgp group toR4 local-address 192.0.2.2
user@PE1# set protocols bgp group toR4 family inet-vpn unicast
user@PE1# set protocols bgp group toR4 neighbor 192.0.2.4
```

4. Configure OSPF on the core-facing interface and on the loopback interface that is being used in the main instance.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface fe-1/2/1.0
user@PE1# set interface lo0.0 passive
```

5. Configure LDP or RSVP on the core-facing interface and on the loopback interface that is being used in the main instance.

```
[edit protocols ldp]
user@PE1# set interface fe-1/2/1.0
user@PE1# set interface lo0.0
```

6. Configure a routing policy for use in the routing instance.

```
[edit policy-options policy-statement bgp-to-ospf]
user@PE1# set term 1 from protocol bgp
user@PE1# set term 1 then accept
user@PE1# set term 2 then reject
```

7. Configure the routing instance.

```
[edit routing-instances red]
user@PE1# set instance-type vrf
user@PE1# set interface fe-1/2/0.0
user@PE1# set route-distinguisher 2:1
user@PE1# set vrf-target target:2:1
user@PE1# set protocols ospf export bgp-to-ospf
user@PE1# set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
```

8. Configure the OSPFv2 sham link.

Include the extra loopback interface in the routing instance and also in the OSPF configuration.

Notice that the metric on the sham-link interface is set to 10. On Device CE1's backup OSPF link, the metric is set to 100. This causes the sham link to be the preferred link.

```
[edit routing-instances red]
user@PE1# set interface lo0.1
user@PE1# set protocols ospf sham-link local 198.51.100.2
user@PE1# set protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.4 metric 10
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.1
```

9. Configure the autonomous system (AS) number and the router ID.

```
[edit routing-options]
user@PE1# set router-id 192.0.2.2
user@PE1# set autonomous-system 2
```

10. If you are done configuring the device, commit the configuration.

```
[edit]  
user@R1# commit
```

Results

Confirm your configuration by entering the `show interfaces` and the `show routing-instances` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for PE1:

```
user@PE1# show interfaces  
fe-1/2/0 {  
  unit 0 {  
    family inet {  
      address 10.1.1.2/30;  
    }  
    family mpls;  
  }  
}  
fe-1/2/1 {  
  unit 0 {  
    family inet {  
      address 10.1.1.5/30;  
    }  
    family mpls;  
  }  
}  
lo0 {  
  unit 0 {  
    family inet {  
      address 192.0.2.2/24;  
    }  
  }  
  unit 1 {  
    family inet {  
      address 198.51.100.2/24;  
    }  
  }  
}
```

```
}  
}
```

```
user@PE1# show protocols  
mpls {  
    interface fe-1/2/1.0;  
}  
bgp {  
    group toR4 {  
        type internal;  
        local-address 192.0.2.2;  
        family inet-vpn {  
            unicast;  
        }  
        neighbor 192.0.2.4;  
    }  
}  
ospf {  
    area 0.0.0.0 {  
        interface fe-1/2/1.0;  
        interface lo0.0 {  
            passive;  
        }  
    }  
}  
ldp {  
    interface fe-1/2/1.0;  
    interface lo0.0;  
}
```

```
user@PE1# show policy-options  
policy-statement bgp-to-ospf {  
    term 1 {  
        from protocol bgp;  
        then accept;  
    }  
    term 2 {  
        then reject;  
    }  
}
```

```
}  
}
```

```
user@PE1# show routing-instances  
red {  
  instance-type vrf;  
  interface fe-1/2/0.0;  
  interface lo0.1;  
  route-distinguisher 2:1;  
  vrf-target target:2:1;  
  protocols {  
    ospf {  
      export bgp-to-ospf;  
      sham-link local 198.51.100.2;  
      area 0.0.0.0 {  
        sham-link-remote 198.51.100.4 metric 10;  
        interface fe-1/2/0.0;  
        interface lo0.1;  
      }  
    }  
  }  
}
```

```
user@PE1# show routing-options  
router-id 192.0.2.2;  
autonomous-system 2;
```

Verification

IN THIS SECTION

- [Verifying the Sham Link Interfaces | 719](#)
- [Verifying the Local and Remote End Points of the Sham Link | 719](#)
- [Verifying the Sham Link Adjacencies | 720](#)
- [Verifying the Link-State Advertisement | 720](#)
- [Verifying the Path Selection | 721](#)

Confirm that the configuration is working properly.

Verifying the Sham Link Interfaces

Purpose

Verify the sham link interface. The sham link is treated as an interface in OSPFv2, with the named displayed as shamlink.<unique identifier>, where the unique identifier is a number. For example, shamlink.0. The sham link appears as a point-to-point interface.

Action

From operational mode, enter the show ospf interface instance *instance-name* command.

```
user@PE1> show ospf interface instance red
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.1          DR     0.0.0.0   198.51.100.2  0.0.0.0      0
fe-1/2/0.0    PtToPt 0.0.0.0   0.0.0.0     0.0.0.0      1
shamlink.0    PtToPt 0.0.0.0   0.0.0.0     0.0.0.0      1
```

Verifying the Local and Remote End Points of the Sham Link

Purpose

Verify the local and remote end points of the sham link. The MTU for the sham link interface is always zero.

Action

From operational mode, enter the show ospf interface instance *instance-name* detail command.

```
user@PE1> show ospf interface shamlink.0 instance red
Interface      State   Area      DR ID      BDR ID      Nbrs
shamlink.0    PtToPt 0.0.0.0   0.0.0.0     0.0.0.0      1
  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 0, Cost: 10
  Local: 198.51.100.2, Remote: 198.51.100.4
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
```

```
Protection type: None, No eligible backup
Topology default (ID 0) -> Cost: 10
```

Verifying the Sham Link Adjacencies

Purpose

Verify the adjacencies between the configured sham links.

Action

From operational mode, enter the `show ospf neighbor instance instance-name` command.

```
user@PE1> show ospf neighbor instance red
Address          Interface          State   ID              Pri  Dead
10.1.1.1         fe-1/2/0.0        Full   192.0.2.1       128  35
198.51.100.4    shamlink.0         Full   198.51.100.4   0    31
```

Verifying the Link-State Advertisement

Purpose

Verify that the router LSA originated by the instance carries the sham link adjacency as an unnumbered point-to-point link. The link data for sham links is a number ranging from 0x80010000 through 0x8001ffff.

Action

From operational mode, enter the `show ospf database instance instance-name` command.

```
user@PE1> show ospf database instance red

OSPF database, Area 0.0.0.0
Type   ID              Adv Rtr          Seq      Age Opt  Cksum Len
Router 192.0.2.1       192.0.2.1       0x80000009  1803 0x22 0x6ec7 72
Router 192.0.2.5       192.0.2.5       0x80000007   70 0x22 0x2746 72
Router *198.51.100.2 198.51.100.2   0x80000006   55 0x22 0xda6b 60
Router 198.51.100.4   198.51.100.4   0x80000005   63 0x22 0xb19 60
Network 10.0.0.18      192.0.2.5       0x80000002   70 0x22 0x9a71 32
```

OSPF AS SCOPE link state database

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Extern	198.51.100.2	198.51.100.4	0x80000002	72	0xa2	0x343	36
Extern	*198.51.100.4	198.51.100.2	0x80000002	71	0xa2	0xe263	36

Verifying the Path Selection

Purpose

Verify that the Layer 3 VPN path is used instead of the backup path.

Action

From operational mode, enter the traceroute command from Device CE1 to Device CE2.

```
user@CE1> traceroute 192.0.2.5

traceroute to 192.0.2.5 (192.0.2.5), 30 hops max, 40 byte packets
 1 10.1.1.2 (10.1.1.2) 1.930 ms 1.664 ms 1.643 ms
 2 * * *
 3 10.1.1.10 (10.1.1.10) 2.485 ms 1.435 ms 1.422 ms
    MPLS Label=299808 CoS=0 TTL=1 S=1
 4 192.0.2.5 (192.0.2.5) 1.347 ms 1.362 ms 1.329 ms
```

Meaning

The traceroute operation shows that the Layer 3 VPN is the preferred path. If you were to remove the sham link or if you were to modify the OSPF metric to prefer that backup path, the traceroute would show that the backup path is preferred.

RELATED DOCUMENTATION

| [Day One: Advanced OSPF in the Enterprise](#)

16

CHAPTER

Configure OSPF on Logical Systems

IN THIS CHAPTER

- [Configuring OSPF on Logical Systems | 723](#)
-

Configuring OSPF on Logical Systems

IN THIS SECTION

- [OSPF Support for Logical Systems | 723](#)
- [Example: Configuring OSPF on Logical Systems Within the Same Router | 724](#)

OSPF Support for Logical Systems

IN THIS SECTION

- [Introduction to Logical Systems | 723](#)
- [OSPF and Logical Systems | 723](#)

This topic describes the following information:

Introduction to Logical Systems

With Junos OS, you can partition a single physical router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the main router, logical systems offer an effective way to maximize the use of a single routing or switching platform. Logical systems have their own unique routing tables, interfaces, policies, and routing instances.

OSPF and Logical Systems

You can configure both OSPF Version 2 (OSPFv2) and OSPF Version 3 (OSPFv3) for logical systems. In the case of OSPFv3, you can also configure OSPFv3 realms for logical systems, which allows OSPFv3 to advertise address families other than unicast IPv6.

You configure OSPF for logical systems at the following hierarchy levels:

- `[edit logical-systems logical-system-name protocols (ospf | ospf3)]`

- [edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3)]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Example: Configuring OSPF on Logical Systems Within the Same Router

IN THIS SECTION

- Requirements | 724
- Overview | 724
- Configuration | 726
- Verification | 731

This example shows how to configure an OSPF network using multiple logical systems that are running on a single physical router. The logical systems are connected by logical tunnel interfaces.

Requirements

You must connect the logical systems by using logical tunnel (**lt**) interfaces. See [Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches](#).

Overview

IN THIS SECTION

- Topology | 725


This example shows the configuration of a single OSPF area with three logical systems running on one physical router. Each logical system has its own routing table. The configuration enables the protocol on

all logical system interfaces that participate in the OSPF domain and specifies the area that the interfaces are in.

Topology

Figure 43 on page 725 shows the sample network.

Figure 43: OSPF on Logical Systems



g040567

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 726](#)
- [Procedure | 727](#)
- [Results | 729](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```

set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS3
set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 5
set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 10.0.1.2/30
set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address 10.0.0.1/30
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.0
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30

```

```

set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3

```

Procedure

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure OSPF on logical systems:

1. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```

[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address 10.0.0.1/30

```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS3.

```

[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS3
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 5
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 10.0.1.2/30

```

3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.

```

[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet

```

```

user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 10.0.0.2/30

```

4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.

```

[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address 10.0.2.2/30

```

5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.

```

[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30

```

6. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS1.

```

[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30

```

7. Configure OSPF on all the interfaces.

```

[edit]
user@host# set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.0
user@host# set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2
user@host# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
user@host# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
user@host# set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.5
user@host# set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3

```

8. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

Results

Confirm your configuration by issuing the `show logical-systems` command.

```
show logical-systems  
LS1 {  
  interfaces {  
    lt-1/2/0 {  
      unit 0 {  
        description LS1->LS3;  
        encapsulation ethernet;  
        peer-unit 5;  
        family inet {  
          address 10.0.1.2/30;  
        }  
      }  
      unit 2 {  
        description LS1->LS2;  
        encapsulation ethernet;  
        peer-unit 1;  
        family inet {  
          address 10.0.0.1/30;  
        }  
      }  
    }  
  }  
  protocols {  
    ospf {  
      area 0.0.0.0 {  
        interface lt-1/2/0.0;  
        interface lt-1/2/0.2;  
      }  
    }  
  }  
}  
LS2 {
```

```
interfaces {
  lt-1/2/0 {
    unit 1 {
      description LS2->LS1;
      encapsulation ethernet;
      peer-unit 2;
      family inet {
        address 10.0.0.2/30;
      }
    }
    unit 4 {
      description LS2->LS3;
      encapsulation ethernet;
      peer-unit 3;
      family inet {
        address 10.0.2.2/30;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface lt-1/2/0.1;
      interface lt-1/2/0.4;
    }
  }
}
LS3 {
  interfaces {
    lt-1/2/0 {
      unit 3 {
        description LS3->LS2;
        encapsulation ethernet;
        peer-unit 4;
        family inet {
          address 10.0.2.1/30;
        }
      }
      unit 5 {
        description LS3->LS1;
        encapsulation ethernet;
```



```

lt-1/2/0.0      up   up   inet   10.0.1.2/30
lt-1/2/0.1      up   up   inet   10.0.0.2/30
lt-1/2/0.2      up   up   inet   10.0.0.1/30
lt-1/2/0.3      up   up   inet   10.0.2.1/30
lt-1/2/0.4      up   up   inet   10.0.2.2/30
lt-1/2/0.5      up   up   inet   10.0.1.1/30
...

```

Verifying Connectivity Between the Logical Systems

Purpose

Make sure that the OSPF adjacencies are established by checking the OSPF neighbor tables, checking the routing tables, and pinging the logical systems.

Action

```

user@host> show ospf neighbor logical-system LS1
Address      Interface      State  ID           Pri  Dead
10.0.1.1     lt-1/2/0.0    Full  10.0.1.1    128  37
10.0.0.2     lt-1/2/0.2    Full  10.0.0.2    128  33

```

```

user@host> show ospf neighbor logical-system LS2
Address      Interface      State  ID           Pri  Dead
10.0.0.1     lt-1/2/0.1    Full  10.0.0.1    128  32
10.0.2.1     lt-1/2/0.4    Full  10.0.1.1    128  36

```

```

user@host> show ospf neighbor logical-system LS3
Address      Interface      State  ID           Pri  Dead
10.0.2.2     lt-1/2/0.3    Full  10.0.0.2    128  36
10.0.1.2     lt-1/2/0.5    Full  10.0.0.1    128  37

```

```

user@host> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30      *[Direct/0] 00:28:00

```

```

> via lt-1/2/0.2
10.0.0.1/32      *[Local/0] 00:28:00
                 Local via lt-1/2/0.2
10.0.1.0/30     *[Direct/0] 00:28:00
                 > via lt-1/2/0.0
10.0.1.2/32     *[Local/0] 00:28:00
                 Local via lt-1/2/0.0
10.0.2.0/30     *[OSPF/10] 00:27:05, metric 2
                 > to 10.0.1.1 via lt-1/2/0.0
                 to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32   *[OSPF/10] 00:28:03, metric 1
                 MultiRecv

```

```

user@host> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30     *[Direct/0] 00:28:31
                 > via lt-1/2/0.1
10.0.0.2/32     *[Local/0] 00:28:32
                 Local via lt-1/2/0.1
10.0.1.0/30     *[OSPF/10] 00:27:38, metric 2
                 > to 10.0.0.1 via lt-1/2/0.1
                 to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30     *[Direct/0] 00:28:32
                 > via lt-1/2/0.4
10.0.2.2/32     *[Local/0] 00:28:32
                 Local via lt-1/2/0.4
224.0.0.5/32   *[OSPF/10] 00:28:34, metric 1
                 MultiRecv

```

```

user@host> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30     *[OSPF/10] 00:28:23, metric 2
                 > to 10.0.2.2 via lt-1/2/0.3
                 to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30     *[Direct/0] 00:29:13
                 > via lt-1/2/0.5

```

```

10.0.1.1/32      *[Local/0] 00:29:15
                  Local via lt-1/2/0.5
10.0.2.0/30     *[Direct/0] 00:29:14
                  > via lt-1/2/0.3
10.0.2.1/32     *[Local/0] 00:29:15
                  Local via lt-1/2/0.3
224.0.0.5/32    *[OSPF/10] 00:29:16, metric 1
                  MultiRecv

```

From LS1, Ping LS3

```
user@host> set cli logical-system LS1
```

```

user@host:LS1> ping 10.0.2.1
PING 10.0.2.1 (10.0.2.1): 56 data bytes
64 bytes from 10.0.2.1: icmp_seq=0 ttl=64 time=1.215 ms
64 bytes from 10.0.2.1: icmp_seq=1 ttl=64 time=1.150 ms
64 bytes from 10.0.2.1: icmp_seq=2 ttl=64 time=1.134 ms

```

From LS3, Ping LS1

```
user@host> set cli logical-system LS3
```

```

user@host:LS3> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.193 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.114 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.190 ms

```

RELATED DOCUMENTATION

[Logical Systems and Tenant Systems User Guide for Security Devices](#)

[Example: Creating an Interface on a Logical System](#)

[Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches](#)

[Example: Configuring a Conditional OSPF Default Route Policy on Logical Systems](#)

[Example: Configuring an OSPF Default Route Policy on Logical Systems](#)

[Example: Configuring an OSPF Import Policy on Logical Systems](#)

17

CHAPTER

Troubleshooting Network Issues

IN THIS CHAPTER

- [Troubleshooting Network Issues | 737](#)
-

Troubleshooting Network Issues

IN THIS SECTION

- [Working with Problems on Your Network | 737](#)
- [Isolating a Broken Network Connection | 738](#)
- [Identifying the Symptoms of a Broken Network Connection | 740](#)
- [Isolating the Causes of a Network Problem | 742](#)
- [Taking Appropriate Action for Resolving the Network Problem | 743](#)
- [Evaluating the Solution to Check Whether the Network Problem Is Resolved | 745](#)
- [Checklist for Tracking Error Conditions | 747](#)
- [Configure Routing Protocol Process Tracing | 749](#)
- [Configure Routing Protocol Tracing for a Specific Routing Protocol | 752](#)
- [Monitor Trace File Messages Written in Near-Real Time | 755](#)
- [Stop Trace File Monitoring | 756](#)

Working with Problems on Your Network

IN THIS SECTION

- [Problem | 737](#)
- [Solution | 738](#)

Problem

Description

This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

Solution


Table 4: Checklist for Working with Problems on Your Network

Tasks	Command or Action
<i>Isolating a Broken Network Connection</i>	
1. <i>Identifying the Symptoms of a Broken Network Connection</i>	ping (ip-address hostname) show route (ip-address hostname) traceroute (ip-address hostname)
1. <i>Isolating the Causes of a Network Problem</i>	show < configuration interfaces protocols route >
1. <i>Taking Appropriate Action for Resolving the Network Problem</i>	[edit] delete routing options static route <i>destination-prefix</i> commit and-quit show route <i>destination-prefix</i>
1. <i>Evaluating the Solution to Check Whether the Network Problem Is Resolved</i>	show route (ip-address hostname) ping (ip-address hostname) count 3 traceroute (ip-address hostname)

Isolating a Broken Network Connection

By applying the standard four-step process illustrated in [Figure 44 on page 738](#), you can isolate a failed node in the network. Note that the functionality described in this section is not supported in versions 15.1X49, 15.1X49-D30, or 15.1X49-D40.


Figure 44: Process for Diagnosing Problems in Your Network



Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic in your approach to the maintenance and monitoring of your network. To prepare for problems on your network, understand how the network functions under normal conditions, have records of baseline network activity, and carefully observe the behavior of your network during a problem situation.

Figure 45 on page 739 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 45: Network with a Problem



The network in Figure 45 on page 739 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (R1) in AS 65001 announces aggregated prefixes 100.100.0/24 to the AS 65002 network. The problem in this network is that R6 does not have access to R5 because of a loop between R2 and R6.

To isolate a failed connection in your network, follow the steps in these topics:

- *Isolating the Causes of a Network Problem*

- *Taking Appropriate Action for Resolving the Network Problem*
- *Taking Appropriate Action for Resolving the Network Problem*
- *Evaluating the Solution to Check Whether the Network Problem Is Resolved*

Identifying the Symptoms of a Broken Network Connection

IN THIS SECTION

- Problem | 740
- Solution | 740

Problem

Description

The symptoms of a problem in your network are usually quite obvious, such as the failure to reach a remote host.

Solution

To identify the symptoms of a problem on your network, start at one end of your network and follow the routes to the other end, entering all or one of the following Junos OS command-line interfaces (CLI) operational mode commands:

```
user@host> ping (ip-address | host-name)
user@host> show route (ip-address | host-name)
user@host> traceroute (ip-address | host-name)
```

Sample Output

```
user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
```

```

 4  5  00 0054 e2db  0 0000 01 01 a8c6 10.1.26.2 10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS Len  ID Flg off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2de  0 0000 01 01 a8c3 10.1.26.2 10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS Len  ID Flg off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2e2  0 0000 01 01 a8bf 10.1.26.2 10.0.0.5

^C
--- 10.0.0.5 ping statistics ---
 3 packets transmitted, 0 packets received, 100% packet loss

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[IS-IS/165] 00:02:39, metric 10
                    > to 10.1.26.1 via so-0/0/2.0

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.649 ms  0.521 ms  0.490 ms
 2  10.1.26.2 (10.1.26.2)  0.521 ms  0.537 ms  0.507 ms
 3  10.1.26.1 (10.1.26.1)  0.523 ms  0.536 ms  0.514 ms
 4  10.1.26.2 (10.1.26.2)  0.528 ms  0.551 ms  0.523 ms
 5  10.1.26.1 (10.1.26.1)  0.531 ms  0.550 ms  0.524 ms

```

Meaning

The sample output shows an unsuccessful ping command in which the packets are being rejected because the time to live is exceeded. The output for the `show route` command shows the interface (10.1.26.1) that you can examine further for possible problems. The `traceroute` command shows the loop between 10.1.26.1 (R2) and 10.1.26.2 (R6), as indicated by the continuous repetition of the two interface addresses.

Isolating the Causes of a Network Problem

IN THIS SECTION

- Problem | 742
- Solution | 742

Problem

Description

A particular symptom can be the result of one or more causes. Narrow down the focus of your search to find each individual cause of the unwanted behavior.

Solution

To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

```
user@host> show < configuration | bgp | interfaces | isis | ospf | route
>
```

Your particular problem may require the use of more than just the commands listed above. See the appropriate command reference for a more exhaustive list of commonly used operational mode commands.

Sample Output

```
user@R6> show interfaces terse
Interface           Admin Link Proto Local           Remote
so-0/0/0            up   up
so-0/0/0.0          up   up   inet 10.1.56.2/30
                   iso
so-0/0/2            up   up
so-0/0/2.0          up   up   inet 10.1.26.2/30
                   iso
so-0/0/3            up   up
```

```
so-0/0/3.0          up    up    inet 10.1.36.2/30
                   iso
[...Output truncated...]
```

The following sample output is from R2:

```
user@R2> show route 10.0.0.5

inet.0: 22 destinations, 25 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[Static/5] 00:16:21
                    > to 10.1.26.2 via so-0/0/2.0
                    [BGP/170] 3d 20:23:35, MED 5, localpref 100
                    AS path: 65001 I
                    > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows that all interfaces on R6 are up. The output from R2 shows that a static route [Static/5] configured on R2 points to R6 (10.1.26.2) and is the preferred route to R5 because of its low preference value. However, the route is looping from R2 to R6, as indicated by the missing reference to R5 (10.1.15.2).

Taking Appropriate Action for Resolving the Network Problem

IN THIS SECTION

- Problem | 744
- Solution | 744

Problem

Description

The appropriate action depends on the type of problem you have isolated. In this example, a static route configured on R2 is deleted from the [routing-options] hierarchy level. Other appropriate actions might include the following:

Solution

- Check the local router's configuration and edit it if appropriate.
- Troubleshoot the intermediate router.
- Check the remote host configuration and edit it if appropriate.
- Troubleshoot routing protocols.
- Identify additional possible causes.

To resolve the problem in this example, enter the following Junos OS CLI commands:

```
[edit]
user@R2# delete routing-options static route destination-
prefix
user@R2# commit and-quit
user@R2# show route destination-prefix
```

Sample Output

```
[edit]
user@R2# delete routing-options static route 10.0.0.5/32

[edit]
user@R2# commit and-quit
commit complete
Exiting configuration mode

user@R2> show route 10.0.0.5

inet.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.5/32      *[BGP/170] 3d 20:26:17, MED 5, localpref 100
                  AS path: 65001 I
                  > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows the static route deleted from the [routing-options] hierarchy and the new configuration committed. The output for the `show route` command now shows the BGP route as the preferred route, as indicated by the asterisk (*).

Evaluating the Solution to Check Whether the Network Problem Is Resolved

IN THIS SECTION

- [Problem | 745](#)
- [Solution | 746](#)

Problem

Description

If the problem is solved, you are finished. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In relation to the network in *Isolating a Broken Network Connection*, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you have reason to believe that the problem is related to a known issue, such as a recent change in configuration.

Solution

To evaluate the solution, enter the following Junos OS CLI commands:

```

user@host> show route (ip-address | host-name)
user@host> ping (ip-address | host-name)
user@host> traceroute (ip-address | host-name)

```

Sample Output

```

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[BGP/170] 00:01:35, MED 5, localpref 100, from 10.0.0.2
                    AS path: 65001 I
                    > to 10.1.26.1 via so-0/0/2.0

user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=253 time=0.866 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=253 time=0.837 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=253 time=0.796 ms
^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.796/0.833/0.866/0.029 ms

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1 10.1.26.1 (10.1.26.1) 0.629 ms 0.538 ms 0.497 ms
 2 10.1.12.1 (10.1.12.1) 0.534 ms 0.538 ms 0.510 ms
 3 10.0.0.5 (10.0.0.5) 0.776 ms 0.705 ms 0.672 ms

```

Meaning

The sample output shows that there is now a connection between R6 and R5. The `show route` command shows that the BGP route to R5 is preferred, as indicated by the asterisk (*). The `ping` command is successful and the `traceroute` command shows that the path from R6 to R5 is through R2 (10.1.26.1), and then through R1 (10.1.12.1).

Checklist for Tracking Error Conditions

IN THIS SECTION

- Problem | 747
- Solution | 747

Problem

Description

[Table 5 on page 747](#) provides links and commands for configuring routing protocol daemon tracing, Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS) protocol, and Open Shortest Path First (OSPF) protocol tracing to diagnose error conditions.

Solution

Table 5: Checklist for Tracking Error Conditions

Tasks	Command or Action
Configure Routing Protocol Process Tracing	
1. <i>Configure Routing Protocol Process Tracing</i>	[edit] edit routing-options traceoptions <i>filename</i> size <i>size</i> files <i>number</i> show con log <i>filename</i>
1. <i>Configure Routing Protocol Tracing for a Specific Routing Protocol</i>	[edit] edit protocol <i>protocol-name</i> trace <i>filename</i> size <i>size</i> files <i>number</i> show con log <i>filename</i>
1. <i>Monitor Trace File Messages Written in Near-Real Time</i>	monitor start <i>filename</i>
1. <i>Stop Trace File Monitoring</i>	monitor stop <i>filename</i>

Table 5: Checklist for Tracking Error Conditions (Continued)

Tasks	Command or Action
Configure BGP-Specific Options	
1. Display Detailed BGP Protocol Information	[edit] edit protocol bgp traceoptions send detail show commit run show log <i>filename</i>
1. Display Sent or Received BGP Packets	[edit] edit protocol bgp traceoptions send (send receive) show commit run show log
1. Diagnose BGP Session Establishment Problems	[edit] edit protocol bgp set traceoptions send detail show commit run show log <i>filename</i>
Configure IS-IS-Specific Options	
1. Displaying Detailed IS-IS Protocol Information	[edit] edit protocol isis traceoptions send detail show commit run show log <i>filename</i>
1. Displaying Sent or Received IS-IS Protocol Packets	[edit] edit protocols isis traceoptions send (send receive) show commit run show log
1. Analyzing IS-IS Link-State PDUs in Detail	[edit] edit protocols isis traceoptions send detail show commit run show log <i>filename</i>
Configure OSPF-Specific Options	
1. Diagnose OSPF Session Establishment Problems	[edit] edit protocols ospf traceoptions send detail show commit run show log <i>filename</i>
1. Analyze OSPF Link-State Advertisement Packets in Detail	[edit] edit protocols ospf traceoptions send update detail show commit run show log

Configure Routing Protocol Process Tracing

IN THIS SECTION

- [Action | 749](#)
- [Meaning | 751](#)

Action

To configure routing protocol process (rpd) tracing, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the file, file size, number, and flags:

```
[edit routing-options traceoptions]
user@host# set file filename size size file number
[edit routing-options traceoptions]
user@host# set flag flag
```

For example:

```
[edit routing-options traceoptions]
user@host# set file daemonlog size 10240 files 10
[edit routing-options traceoptions]
user@host# set flag general
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit routing-options traceoptions]
user@host# show
file daemonlog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```



NOTE: Some traceoptions flags generate an extensive amount of information. Tracing can also slow down the operation of routing protocols. Delete the traceoptions configuration if you no longer require it.

1. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit routing-options traceoptions]
user@pro4-a# run show log daemonlog
Sep 17 14:17:31 trace_on: Tracing to "/var/log/daemonlog" started
Sep 17 14:17:31 Tracing flags enabled: general
Sep 17 14:17:31 inet_routerid_notify: Router ID: 10.255.245.44
Sep 17 14:17:31 inet_routerid_notify: No Router ID assigned
Sep 17 14:17:31 Initializing LSI globals
Sep 17 14:17:31 LSI initialization complete
Sep 17 14:17:31 Initializing OSPF instances
Sep 17 14:17:31 Reinitializing OSPFv2 instance master
Sep 17 14:17:31 OSPFv2 instance master running
[...Output truncated...]
```

Meaning

Table 6 on page 751 lists tracing flags and example output for Junos-supported routing protocol daemon tracing.

Table 6: Routing Protocol Daemon Tracing Flags

Tracing Flag	Description	Example Output
all	All operations	Not available.
general	Normal operations and routing table change	Not available.
normal	Normal operations	Not available.
policy	Policy operations and actions	Nov 29 22:19:58 export: Dest 10.0.0.0 proto Static Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0 Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0 Nov 29 22:19:58 export: Dest 10.10.10.0 proto IS-IS
route	Routing table changes	Nov 29 22:23:59 Nov 29 22:23:59 rtlist_walker_job: rt_list walk for RIB inet.0 started with 42 entries Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) start Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) done Nov 29 22:23:59 rtlist_walker_job: rt_list walk for inet.0 ended with 42 entries Nov 29 22:23:59 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 CHANGE route/user af 2 addr 172.16.0.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 172.17.0.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 10.149.3.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:24:19 trace_on: Tracing to "/var/log/rpdlog" started Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 10.10.218.0 nhop-type unicast nhop 10.10.10.29 Nov 29 22:24:19 RELEASE 10.10.218.0 255.255.255.0 gw 10.10.10.29,10.10.10.33 BGP pref 170/-101 metric so-1/1/0.0,so-1/1/1.0 <Release Delete Int Ext> as 65401 Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 172.18.0.0 nhop-type unicast nhop 10.10.10.33
state	State transitions	Not available.

Table 6: Routing Protocol Daemon Tracing Flags *(Continued)*

Tracing Flag	Description	Example Output
task	Interface transactions and processing	<pre>Nov 29 22:50:04 foreground dispatch running job task_collect for task Scheduler Nov 29 22:50:04 task_collect_job: freeing task MGMT_Listen (DELETED) Nov 29 22:50:04 foreground dispatch completed job task_collect for task Scheduler Nov 29 22:50:04 background dispatch running job rt_static_update for task RT Nov 29 22:50:04 task_job_delete: delete background job rt_static_update for task RT Nov 29 22:50:04 background dispatch completed job rt_static_update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 background dispatch returned job Flash update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT Nov 29 22:50:04 background dispatch completed job Flash update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT</pre>
timer	Timer usage	<pre>Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 1 timer Nov 29 22:52:07 task_timer_hiprio_dispatch: running high priority timer queue Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 2 timers</pre>

Configure Routing Protocol Tracing for a Specific Routing Protocol

IN THIS SECTION

- Action | 752
- Meaning | 754

Action

To configure routing protocol tracing for a specific routing protocol, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocol protocol-name traceoptions
```

2. Configure the file, file size, number, and flags:

```
[edit protocols protocol name traceoptions]
user@host# set file filename size size files
number
[edit protocols protocol name traceoptions]
user@host# set flag flag
```

For example:

```
[edit protocols ospf traceoptions]
user@host# set file ospflog size 10240 files 10
[edit protocols ospf traceoptions]
user@host# set flag general
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospflog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit protocols ospf traceoptions]
user@pro4-a# run show log ospflog
Sep 17 14:23:10 trace_on: Tracing to "/var/log/ospflog" started
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) start
Sep 17 14:23:10 OSPF: multicast address 224.0.0.5/32, route ignored
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) done
Sep 17 14:23:10 CHANGE 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Delete Int>
Sep 17 14:23:10 CHANGE 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 ADD 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 CHANGE 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Delete Int>
Sep 17 14:23:10 CHANGE 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 ADD 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 rt_close: 4/4 routes proto OSPF
[...Output truncated...]
```

Meaning

[Table 7 on page 754](#) lists standard tracing options that are available globally or that can be applied to specific protocols. You can also configure tracing for a specific BGP peer or peer group. For more information, see the *Junos System Basics Configuration Guide*.

Table 7: Standard Trace Options for Routing Protocols

Tracing Flag	Description
all	All operations

Table 7: Standard Trace Options for Routing Protocols (*Continued*)

Tracing Flag	Description
general	Normal operations and routing table changes
normal	Normal operations
policy	Policy operations and actions
route	Routing table changes
state	State transitions
task	Interface transactions and processing
timer	Timer usage

Monitor Trace File Messages Written in Near-Real Time

IN THIS SECTION

- [Purpose | 755](#)
- [Action | 756](#)

Purpose

To monitor messages in near-real time as they are being written to a trace file.

Action

To monitor messages in near-real time as they are being written to a trace file, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> monitor start filename
```

Sample Output

command-name

```
user@host> monitor start isis
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
```

Stop Trace File Monitoring

IN THIS SECTION

- [Action | 757](#)
- [Sample Output | 757](#)

Action

To stop monitoring a trace file in near-real time, use the following Junos OS CLI operational mode command after you have started monitoring:

```
user@host          monitor stop filename
```

Sample Output

```
user@host> monitor start isis
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
monitor stop isis
user@host>
```

18

CHAPTER

Verifying and Monitoring OSPF

IN THIS CHAPTER

- [Verifying and Monitoring OSPF Configuration | 759](#)
-

Verifying and Monitoring OSPF Configuration

IN THIS SECTION

- [Verifying an OSPF Configuration | 759](#)
- [Tracing OSPF Protocol Traffic | 765](#)
- [Example: Tracing OSPF Protocol Traffic | 767](#)

Verifying an OSPF Configuration

IN THIS SECTION

- [Verifying OSPF-Enabled Interfaces | 759](#)
- [Verifying OSPF Neighbors | 761](#)
- [Verifying the Number of OSPF Routes | 762](#)
- [Verifying Reachability of All Hosts in an OSPF Network | 764](#)

To verify an OSPF configuration, perform these tasks:

Verifying OSPF-Enabled Interfaces

IN THIS SECTION

- [Purpose | 760](#)
- [Action | 760](#)
- [Meaning | 760](#)

Purpose

Verify that OSPF is running on a particular interface and that the interface is in the desired area.

Action

From the CLI, enter the **show ospf interface** command.

Sample Output

command-name

```
user@host> show ospf interface
```

Intf	State	Area	DR ID	BDR ID	Nbrs
at-5/1/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
ge-2/3/0.0	DR	0.0.0.0	192.168.4.16	192.168.4.15	1
lo0.0	DR	0.0.0.0	192.168.4.16	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-6/0/2.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Meaning

The output shows a list of the device interfaces that are configured for OSPF. Verify the following information:

- Each interface on which OSPF is enabled is listed.
- Under **Area**, each interface shows the area for which it was configured.
- Under **Intf** and **State**, the device loopback (**lo0.0**) interface and LAN interface that are linked to the OSPF network's designated router (DR) are identified.
- Under **DR ID**, the IP address of the OSPF network's designated router appears.
- Under **State**, each interface shows a state of **PtToPt** to indicate a point-to-point connection. If the state is **Waiting**, check the output again after several seconds. A state of **Down** indicates a problem.
- The designated router addresses always show a state of **DR**.

Verifying OSPF Neighbors

IN THIS SECTION

- Purpose | 761
- Action | 761
- Meaning | 761

Purpose

OSPF neighbors are interfaces that have an immediate adjacency. On a point-to-point connection between the device and another router running OSPF, verify that each router has a single OSPF neighbor.

Action

From the CLI, enter the **show ospf neighbor** command.

Sample Output

command-name

```
user@host> show ospf neighbor
Address      Intf          State      ID           Pri  Dead
192.168.254.225 fxp3.0       2Way      10.250.240.32 128  36
192.168.254.230 fxp3.0       Full      10.250.240.8  128  38
192.168.254.229 fxp3.0       Full      10.250.240.35 128  33
10.1.1.129     fxp2.0       Full      10.250.240.12 128  37
10.1.1.131     fxp2.0       Full      10.250.240.11 128  38
10.1.2.1       fxp1.0       Full      10.250.240.9  128  32
10.1.2.81      fxp0.0       Full      10.250.240.10 128  33
```

Meaning

The output shows a list of the device's OSPF neighbors and their addresses, interfaces, states, router IDs, priorities, and number of seconds allowed for inactivity ("dead" time). Verify the following information:

- Each interface that is immediately adjacent to the device is listed.
- The device's own loopback address and the loopback addresses of any routers with which the device has an immediate adjacency are listed.
- Under **State**, each neighbor shows a state of **Full**. Because full OSPF connectivity is established over a series of packet exchanges between clients, the OSPF link might take several seconds to establish. During that time, the state might be displayed as **Attempt**, **Init**, or **2way**, depending on the stage of negotiation.

If, after 30 seconds, the state is not **Full**, the OSPF configuration between the neighbors is not functioning correctly.

Verifying the Number of OSPF Routes

IN THIS SECTION

- Purpose | 762
- Action | 763
- Meaning | 764


Purpose

Verify that the OSPF routing table has entries for the following:

- Each subnetwork reachable through an OSPF link
- Each loopback address reachable on the network

For example, Figure 1 shows a sample network with an OSPF topology.

Figure 46: Sample OSPF Network Topology



In this topology, OSPF is being run on all interfaces. Each segment in the network is identified by an address with a /24 prefix, with interfaces on either end of the segment being identified by unique IP addresses.

Action

From the CLI, enter the **show ospf route** command.

Sample Output

command-name

```
user@host> show ospf route
Prefix          Path  Route  NH  Metric  NextHop  Nexthop
                Type  Type   Type
10.10.10.1/24   Intra Network IP  1      ge-0/0/2.0  10.0.21.1
10.10.10.2/24   Intra Network IP  1      ge-0/0/2.0  10.0.21.1
10.10.10.4/24   Intra Network IP  1      ge-0/0/1.0  10.0.13.1
10.10.10.5/24   Intra Network IP  1      ge-0/0/2.0  10.0.21.1
10.10.10.6/24   Intra Network IP  1      ge-0/0/1.0  10.0.13.1
10.10.10.10/24  Intra Network IP  1      ge-0/0/2.0  10.0.21.1
10.10.10.11/24  Intra Network IP  1      ge-0/0/1.0  10.0.13.1
10.10.10.13/24  Intra Network IP  1      ge-0/0/1.0
10.10.10.16/24  Intra Network IP  1      ge-0/0/1.0  10.0.13.1
10.10.10.19/24  Intra Network IP  1      ge-0/0/1.0  10.0.13.1
10.10.10.20/24  Intra Network IP  1      ge-0/0/2.0  10.0.21.1
10.10.10.21/24  Intra Network IP  1      ge-0/0/2.0
192.168.5.1     Intra Router  IP  1      ge-0/0/2.0  10.0.21.1
192.168.5.2     Intra Router  IP  1      lo0
192.168.5.3     Intra Router  IP  1      ge-0/0/1.0  10.0.13.1
```

192.168.5.4	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1
192.168.5.5	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1
192.168.5.6	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1
192.168.5.7	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1
192.168.5.8	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1
192.168.5.9	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1

Meaning

The output lists each route, sorted by IP address. Routes are shown with a route type of **Network**, and loopback addresses are shown with a route type of **Router**.

For the example shown in Figure 1, verify that the OSPF routing table has 21 entries, one for each network segment and one for each router's loopback address.

Verifying Reachability of All Hosts in an OSPF Network

IN THIS SECTION

- Purpose | 764
- Action | 764
- Meaning | 765

Purpose

By using the traceroute tool on each loopback address in the network, verify that all hosts in the network are reachable from each device.

Action

For each device in the OSPF network:

1. In the J-Web interface, select **Troubleshoot>Traceroute**.
2. In the Host Name box, type the name of a host for which you want to verify reachability from the device.
3. Click **Start**. Output appears on a separate page.

Sample Output

command-name

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

Meaning

Each numbered row in the output indicates a routing “hop” in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop, for each traceroute packet. To ensure that the OSPF network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is likely not reachable. In this case, verify the routes with the **show ospf route** command.

For information about **show ospf route**, see ["Verifying the Number of OSPF Routes" on page 759](#)

Tracing OSPF Protocol Traffic

Tracing operations record detailed messages about the operation of OSPF. You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace.

You can specify the following OSPF protocol-specific trace options:

- **database-description**—All database description packets, which are used in synchronizing the OSPF topological database
- **error**—OSPF error packets
- **event**—OSPF state transitions
- **flooding**—Link-state flooding packets
- **graceful-restart**—Graceful-restart events

- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable
- **ldp-synchronization**—Synchronization events between OSPF and LDP
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database
- **nsr-synchronization**—Nonstop routing synchronization events
- **on-demand**—Trace demand circuit extensions
- **packet-dump**—Dump the contents of selected packet types
- **packets**—All OSPF packets
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events
- **spf**—Shortest path first (SPF) calculations

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



NOTE: Use the **detail** flag modifier with caution as it might cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the OSPF protocol using the **traceoptions flag** statement included at the **[edit protocols ospf]** hierarchy level:

- **all**—All tracing operations

- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



NOTE: Use the trace flag **all** with caution as it might cause the CPU to become very busy.

Example: Tracing OSPF Protocol Traffic

IN THIS SECTION

- [Requirements | 767](#)
- [Overview | 767](#)
- [Configuration | 769](#)
- [Verification | 774](#)

This example shows how to trace OSPF protocol traffic.

Requirements

This example assumes that OSPF is properly configured and running in your network, and you want to trace OSPF protocol traffic for debugging purposes.

Overview

You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace. All files are placed

in a directory on the routing device's hard disk. On M Series and T Series routers, trace files are stored in the `/var/log` directory.

This example shows a few configurations that might be useful when debugging OSPF protocol issues. The verification output displayed is specific to each configuration.



TIP: To keep track of your log files, create a meaningful and descriptive name so it is easy to remember the content of the trace file. We recommend that you place global routing protocol tracing output in the file **routing-log**, and OSPF tracing output in the file **ospf-log**.

In the first example, you globally enable tracing operations for all routing protocols that are actively running on your routing device to the file `routing-log`. With this configuration, you keep the default settings for the trace file size and the number of trace files. After enabling global tracing operations, you enable tracing operations to provide detailed information about OSPF packets, including link-state advertisements, requests, and updates, database description packets, and hello packets to the file `ospf-log`, and you configure the following options:

- **size**—Specifies the maximum size of each trace file, in KB, MB, or GB. In this example, you configure 10 KB as the maximum size. When the file reaches its maximum size, it is renamed with a `.0` extension. When the file again reaches its maximum size, it is renamed with a `.1` extension, and the newly created file is renamed with a `.0` extension. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option. You specify **k** for KB, **m** for MB, and **g** for GB. By default, the trace file size is 128 KB. The file size range is 10 KB through the maximum file size supported on your system.
- **files**—Specifies the maximum number of trace files. In this example, you configure a maximum of 5 trace files. When a trace file reaches its maximum size, it is renamed with a `.0` extension, then a `.1` extension, and so on until the maximum number of trace files is reached. When the maximum number of files is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option. By default, there are 10 files. The range is 2 through 1000 files.

In the second example, you trace all SPF calculations to the file `ospf-log` by including the **spf** flag. You keep the default settings for the trace file size and the number of trace files.

In the third example, you trace the creation, receipt, and retransmission of all LSAs to the file `ospf-log` by including the **lsa-request**, **lsa-update**, and **lsa-ack** flags. You keep the default settings for the trace file size and the number of trace files.

Configuration

IN THIS SECTION

- [Configuring Global Tracing Operations and Tracing OSPF Packet Information | 769](#)
- [Tracing SPF Calculations | 772](#)
- [Tracing Link-State Advertisements | 773](#)

Configuring Global Tracing Operations and Tracing OSPF Packet Information

CLI Quick Configuration

To quickly enable global tracing operations for all routing protocols actively running on your routing device and to trace detailed information about OSPF packets, copy the following commands and paste them into the CLI.

```
[edit]
set routing-options traceoptions file routing-log
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions file files 5 size 10k
set protocols ospf traceoptions flag lsa-ack
set protocols ospf traceoptions flag database-description
set protocols ospf traceoptions flag hello
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-request
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in the [CLI User Guide](#).

To configure global routing tracing operations and tracing operations for OSPF packets:

1. Configure tracing at the routing options level to collect information about the active routing protocols on your routing device.

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the filename for the global trace file.

```
[edit routing-options traceoptions]
user@host# set file routing-log
```

3. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

4. Configure the maximum number of trace files.

```
[edit protocols ospf traceoptions]
user@host# set file files 5
```

5. Configure the maximum size of each trace file.

```
[edit protocols ospf traceoptions]
user@host# set file size 10k
```

6. Configure tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-ack
user@host# set flag database-description
user@host# set flag hello
```

```
user@host# set flag lsa-update
user@host# set flag lsa-request
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results

Confirm your configuration by entering the **show routing-options** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
traceoptions {
  file routing-log;
}
```

```
user@host# show protocols ospf
traceoptions {
  file ospf-log size 10k files 5;
  flag lsa-ack;
  flag database-description;
  flag hello;
  flag lsa-update;
  flag lsa-request;
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options** and the **show protocols ospf3** commands.

Tracing SPF Calculations

CLI Quick Configuration

To quickly trace SPF calculations, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag spf
```

Step-by-Step Procedure

To configure SPF tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the SPF tracing flag.

```
[edit protocols ospf traceoptions]
user@host# set flag spf
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log ;
  flag spf;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Tracing Link-State Advertisements

CLI Quick Configuration

To quickly trace the creation, receipt, and retransmission of all LSAs, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag lsa-request
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-ack
```

Step-by-Step Procedure

To configure link-state advertisement tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the link-state advertisement tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-request
user@host# set flag lsa-update
user@host# set flag lsa-ack
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log;
  flag lsa-request;
  flag lsa-update;
  flag lsa-ack;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

IN THIS SECTION

- [Verifying Trace Operations | 775](#)

Confirm that the configuration is working properly.

Verifying Trace Operations

Purpose

Verify that the Trace options field displays the configured trace operations, and verify that the Trace file field displays the location on the routing device where the file is saved, the name of the file to receive the output of the tracing operation, and the size of the file.

Action

From operational mode, enter the **show ospf overview extensive** command for OSPFv2, and enter the **show ospf3 overview extensive** command for OSPFv3.

RELATED DOCUMENTATION

[Understanding OSPF Configurations | 14](#)

[Tracing and Logging Junos OS Operations](#)

[Example: Tracing Global Routing Protocol Operations](#)

19

CHAPTER

Configuration Statements and Operational Commands

IN THIS CHAPTER

- [Junos CLI Reference Overview | 777](#)
-

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)