



---

Junos<sup>®</sup> OS

# Multitopology Routing Feature Guide for Routing Devices

Release

14.1



---

Published: 2014-05-13

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, 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<sup>®</sup> OS Multitopology Routing Feature Guide for Routing Devices*

14.1

Copyright © 2014, 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 <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	ix
	Documentation and Release Notes . . . . .	ix
	Supported Platforms . . . . .	ix
	Using the Examples in This Manual . . . . .	ix
	Merging a Full Example . . . . .	x
	Merging a Snippet . . . . .	x
	Documentation Conventions . . . . .	xi
	Documentation Feedback . . . . .	xiii
	Requesting Technical Support . . . . .	xiii
	Self-Help Online Tools and Resources . . . . .	xiii
	Opening a Case with JTAC . . . . .	xiv
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Multitopology Routing . . . . .</b>	<b>3</b>
	Multitopology Routing Overview . . . . .	4
	Routing Table Naming Conventions for Multitopology Routing . . . . .	6
	Filter-Based Forwarding Support . . . . .	7
<b>Chapter 2</b>	<b>Multitopology Routing Standards . . . . .</b>	<b>9</b>
	Multitopology Routing Standards . . . . .	9
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 3</b>	<b>Multitopology Routing . . . . .</b>	<b>13</b>
	Example: Configuring Multitopology Routing Based on Applications . . . . .	13
	Understanding Multitopology Routing Based on Applications . . . . .	13
	Example: Configuring Multitopology Routing Based on Applications . . . . .	15
	Example: Configuring Multitopology Routing Based on a Multicast Source . . . . .	39
	Understanding Multitopology Routing Based on a Multicast Source . . . . .	39
	Example: Configuring Multitopology Routing Based on a Multicast Source . . . . .	41
<b>Chapter 4</b>	<b>Multitopology Routing Configuration Statements . . . . .</b>	<b>67</b>
	community (Protocols BGP) . . . . .	68
	rib (Multitopology Routing) . . . . .	69
	topologies (Multitopology Routing) . . . . .	70
	topology (Filter-Based Forwarding) . . . . .	71
	topology (Multitopology Routing) . . . . .	72
	topology (OSPF) . . . . .	73
	topology (OSPF Interface) . . . . .	74
	topology-id . . . . .	75

<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 5</b>	<b>Multitopology Routing Operational Commands</b>	<b>79</b>
	show (ospf   ospf3) interface	80
<b>Part 4</b>	<b>Index</b>	
	Index	89

# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Multitopology Routing</b>	<b>3</b>
	Figure 1: Voice and Video Routing Topologies Enabled on a Subset of Links	5
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 3</b>	<b>Multitopology Routing</b>	<b>13</b>
	Figure 2: MTR-OSPF Area Boundary	14
	Figure 3: BGP Route Resolution	14
	Figure 4: Route Resolution for MTR	15
	Figure 5: Multitopology OSPF and IBGP for Designating Links Belonging to Voice and Video Services	16
	Figure 6: Core Links Configured to Prefer Specified Routing Topologies	39
	Figure 7: Core Links Configured to Prefer Specified Routing Topologies	41
	Figure 8: Multi-topology OSPF and BGP for Designating Links Belonging to Voice and Video Services	42



# List of Tables

	<b>About the Documentation . . . . . ix</b>
	Table 1: Notice Icons . . . . . xi
	Table 2: Text and Syntax Conventions . . . . . xii
<b>Part 1</b>	<b>Overview</b>
<b>Chapter 1</b>	<b>Introduction to Multitopology Routing . . . . . 3</b>
	Table 3: Examples of Routing Tables for Custom Topologies . . . . . 7
<b>Part 3</b>	<b>Administration</b>
<b>Chapter 5</b>	<b>Multitopology Routing Operational Commands . . . . . 79</b>
	Table 4: show (ospf   ospf3) interface Output Fields . . . . . 81



# About the Documentation

- Documentation and Release Notes on page ix
- Supported Platforms on page ix
- Using the Examples in This Manual on page ix
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiii

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

## Supported Platforms

---

For the features described in this document, the following platforms are supported:

- T Series
- MX Series
- M Series

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

## Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

## Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

## Documentation Conventions

Table 1 on page xi defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast</b>   <b>multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members</b> [ <i>community-ids</i> ]
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

---

#### GUI Conventions

---

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>

- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes:  
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:  
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:  
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

## PART 1

# Overview

- [Introduction to Multitopology Routing on page 3](#)
- [Multitopology Routing Standards on page 9](#)



## CHAPTER 1

# Introduction to Multitopology Routing

- [Multitopology Routing Overview on page 4](#)

## Multitopology Routing Overview

---

Service providers and enterprises can use multitopology routing (MTR) to engineer traffic flow across a network. MTR can be used with direct and static routes, OSPF, and BGP. MTR can be configured for unicast and multicast IP, using the Junos<sup>®</sup> operating system (Junos OS). With basic unicast IP, an IBGP core runs on top of OSPF to direct traffic based on application types, such as voice or video. For multicast, Protocol Independent Multicast (PIM) is used, in conjunction with multitopology OSPF and BGP, to direct multicast traffic over particular paths based on traffic characteristics.


In a network carrying multiple traffic types, you often need to direct different types of application traffic over multiple links depending on their link characteristics. For example, voice traffic requires links that are less likely to incur latency, jitter, or packet loss. File traffic, on the other hand, requires links that have large amounts of available bandwidth. MTR is a way to direct traffic to follow specified paths. You can use MTR to extend a traditional MPLS network into a segment where only IP forwarding is supported. With MTR, each traffic type is handled in its own conceptually incongruent topology, and yet runs on top of the same underlying network. You can configure separate topologies to share the same network links as needed. MTR uses a combination of control plane (routing) and forwarding plane (firewall filters). Each topology uses the unified control plane to make routing decisions for traffic associated with that topology. In addition, each topology has a separate forwarding table and, in effect, a dedicated forwarding plane for each topology. This forwarding plane not only directs traffic using its own forwarding table, but also simultaneously handles sophisticated functionality, such as queuing for class of service (CoS), that can be applied to a topology. As traffic enters the router, fields within a packet determine to which topology the traffic belongs.

Multitopology routing enables you to configure class-based forwarding for different types of traffic, such as voice, video, and data. Each type of traffic is defined by a topology that is used to create a new routing table for that topology. MTR provides the ability to generate forwarding tables based on the resolved entries in the routing tables for the custom topologies you create. In this way, packets of different classes can be routed independently from one another.

One way to manage traffic flow is to group links into specific routing topologies based on application requirements. Each routing topology can be thought of as a set of contiguous links. MTR provides a way for you to manage each set of links uniquely by directing traffic types to flow over specified links. This solution uses a combination of routing (control plane) and firewall filtering (forwarding plane) configurations.

[Figure 1 on page 5](#) shows a network with two topologies configured: voice (dotted lines) and video (dashed lines). Note there is also a default routing topology (solid line).

### Figure 1: Voice and Video Routing Topologies Enabled on a Subset of Links



You can configure MTR for BGP, OSPF, and static routes. When a routing topology is created, it has its own forwarding table.

Packet forwarding uses firewall filters to examine packets as they enter the router over an interface. These filters determine whether a specific topology or the default forwarding table should be used to make packet forwarding decisions. If applicable, firewall filters evaluate packet attributes, such as destination IP address, Differentiated Services code points (DSCPs), or next-level protocol headers, to determine which topology to use. In fact, any item in a packet that is recognized by a firewall filter can be used to direct the packet next-hop lookup to use a particular topology. Once the packet is directed to use a topology, the destination IP address must be in the topology forwarding table. Otherwise, the packet is dropped.

The following topics provide background information about multitopology routing:

- [Routing Table Naming Conventions for Multitopology Routing on page 6](#)
- [Filter-Based Forwarding Support on page 7](#)

## Routing Table Naming Conventions for Multitopology Routing

Routing topologies have their own routing tables, similar to any other routing table created by default or by a **rib-group** configuration with a few differences. The routing table name indicates that the routing table is associated with a topology by prepending a colon (:) to the name. For example, a routing topology named **voice** has a routing table named **:voice.inet.0**. When routing topologies are configured under **routing-options**, a new routing table for each topology is created.

Each routing protocol creates a routing table based on the topology name, the instance name, and the purpose of the table. A routing table for each topology uses the following format:

***logical-system-name/routing-instance-name:topology-name.protocol.identifier***

The routing instance string is included only if the instance is not the master. The logical system string is included only if the logical system identifier has a value other than 0 (zero). Each routing table for a topology includes a colon (:) before the topology name that also separates the routing-instance name from the topology name. **protocol** is the

protocol family, which can be **inet** or **inet6**. *identifier* is a positive integer that specifies the instance of the routing table. [Table 3 on page 7](#) shows specific examples of routing tables for various topologies.

**Table 3: Examples of Routing Tables for Custom Topologies**

Name of Routing Table	Description
<code>:voice.inet.0</code>	Master instance, voice topology, unicast IPv4 routes
<code>:voice.inet6.0</code>	Master instance, voice topology, unicast IPv6 routes
<code>:voice.inet.3</code>	Master instance, voice topology, ingress label-switched paths (LSPs)
<code>private_1/:voice.inet.0</code>	Logical system private, voice topology, unicast IPv4 routes
<code>customer-A:voice.inet.0</code>	Virtual-router customer-A, voice topology, unicast IPv4 routes
<code>customer-B:voice.inet.3</code>	Virtual-router customer-B, voice topology, ingress LSPs
<code>customer-A:voice.mpls.0</code>	Virtual-router customer-A, voice topology, unicast carrier-of-carriers IPv4 routes

To run multitopology routing (MTR), you must configure IP routing. MTR supports OSPF version 2 (OSPFv2), static routes, and BGP. You must configure an interior gateway protocol (IGP), such as OSPFv2 or static routing. Configure BGP to add routes learned through BGP to the appropriate custom topologies.

MTR is also supported on logical systems and the virtual router routing instance. No other routing instance type is supported on MTR.

**Filter-Based Forwarding Support**

By default, the ingress interface forwards traffic to the default topology for each configured routing instance. MTR supports filter-based forwarding, which enables you to match traffic on the ingress interface with a specific type of forwarding class and then forward that traffic to the specified topology. You can further define how traffic is handled for each forwarding class by configuring additional firewall filters that match traffic for such values as the IP precedence field or the Differentiated Services code point (DSCP).

- Related Documentation**
- [Example: Configuring Multitopology Routing Based on Applications on page 13](#)
  - [Example: Configuring Multitopology Routing Based on a Multicast Source on page 39](#)



## CHAPTER 2

# Multitopology Routing Standards

- [Multitopology Routing Standards on page 9](#)

### Multitopology Routing Standards

---

Multitopology routing is defined in the following document:

- RFC 4915, *Multi-Topology (MT) Routing in OSPF*



## PART 2

# Configuration

- [Multitopology Routing on page 13](#)
- [Multitopology Routing Configuration Statements on page 67](#)



## CHAPTER 3

# Multitopology Routing

- [Example: Configuring Multitopology Routing Based on Applications on page 13](#)
- [Example: Configuring Multitopology Routing Based on a Multicast Source on page 39](#)

### Example: Configuring Multitopology Routing Based on Applications

---

- [Understanding Multitopology Routing Based on Applications on page 13](#)
- [Example: Configuring Multitopology Routing Based on Applications on page 15](#)

### Understanding Multitopology Routing Based on Applications

Multitopology routing (MTR) enables you to configure class-based forwarding for different types of traffic, such as voice, video, and data. Each type of traffic is defined by a topology that is used to create a new routing table for that topology. MTR provides the ability to generate forwarding tables based on the resolved entries in the routing tables for the custom topologies you create. In this way, packets of different classes can be routed independently from one another.

To run MTR, you must configure IP routing. MTR supports OSPFv2, static routes, and BGP. You must configure an interior gateway protocol (IGP), such as OSPFv2 or static routing. Configure BGP to add routes learned through BGP to the appropriate custom topologies. MTR also supports filter-based forwarding, which enables you to match traffic on the ingress interface with a specific type of forwarding class and then forward that traffic to the specified topology.


#### OSPF in MTR

OSPF in MTR uses a single instance of OSPF to carry connectivity and IP reachability information for different topologies. That information is used to calculate shortest-path-first (SPF) trees and routing tables. OSPF for MTR supports protocol extensions that include metrics that correspond to different topologies for link and prefix reachability information. The type-of-service (TOS) metric field is used to advertise the topology-specific metric for links and prefixes belonging to that topology. The TOS field is redefined as MT-ID in the payload of router, summary, and Type 5 and Type 7 AS-external link-state advertisements (LSAs).

Under MTR, each OSPF interface continues to belong to a single area. Therefore, by default, all topologies share the same area boundaries. As a result, attributes of an area, such as stubbiness, are independent of the topology. By default, all topologies configured

for OSPF are enabled on all interfaces. However, you can disable one or more configured topologies on an interface. You can thus allocate an interface for a specific topology. In [Figure 2 on page 14](#), Area 51 includes an interface that is uniquely allocated to voice traffic, and Area 0 includes an interface that is uniquely allocated to data traffic. Each topology thus corresponds to a different OSPF area that shares a boundary.

**Figure 2: MTR-OSPF Area Boundary**




### BGP in MTR

BGP in MTR provides the ability to resolve BGP routes against configured topologies. An inbound policy is used to select routes for inclusion in the appropriate routing tables for the topologies. The default behavior for virtual private networks (VPNs) that use MPLS for forwarding packets over the backbone and that use BGP for distributing routes over the backbone is to place BGP route updates in the **bgp.l3vpn** routing table.

[Figure 3 on page 14](#) shows a BGP peer operating in an environment that conforms with the requirements in RFC 2547, *BGP/MPLS VPNs*. The figure shows how a BGP peer configured for MTR performs secondary route resolution.

**Figure 3: BGP Route Resolution**




The BGP peer in a standard VPN topology places prefixes for routes it learns in the **bgp.l3vpn** routing table, which does not result in automatic updates to the forwarding table. Under BGP in MTR, when BGP receives a route from a peer, it attempts to resolve that route against a route in the **inet.0** routing table. If the route is resolved, it is placed in that table, which generates a forwarding state. If you have configured a community

target identifier that matches the import policy for the topology, routing and forwarding states are added to the tables for the topology.

Because MTR provides support for BGP to perform secondary route resolution, as [Figure 4 on page 15](#) shows, MTR is able to create two distinct network paths for each type of traffic. Each router advertises BGP routes that need to be resolved against the IGP routes for each topology. Based on the IGP metrics configured for each topology, for all routes that originate from Router 4 (R4), the upper path between R1 and R4, which traverses R2 and R3, is selected for voice traffic, whereas the lower path between R1 and R4, which traverses R5 and R6, is selected for data traffic.

**Figure 4: Route Resolution for MTR**



### Example: Configuring Multitopology Routing Based on Applications

This example shows how to use multitopology routing (MTR) to choose a topology path based on an application, either voice or video.

- [Requirements on page 15](#)
- [Overview on page 15](#)
- [Configuration on page 16](#)
- [Verification on page 33](#)

#### Requirements

This example requires that Junos OS Release 9.0 or later is running on the provider core devices.

#### Overview

In this example, the network is running OSPF and internal BGP (IBGP) in the core, but not MPLS. Even without traffic engineering, voice traffic uses one set of links, and video traffic uses a different set of links. This traffic might or might not be destined for the same IP address. In some cases, both applications traverse the same link. The solution uses MTR-based OSPF and BGP, along with firewall filters to direct different traffic types over designated links. The routers use a fairly similar set of configurations, which reduces complexity and improves network management.


The OSPF topologies are defined to support each service offering over the OSPF area. The links of a topology must be contiguous, consistent with a typical OSPF area. IBGP routes in each routing topology automatically use the associated OSPF topology routing table for protocol next-hop route resolution. No special route resolution configurations

are required. In this solution, multiple topologies can be configured over the same link. However, traffic in each application service class cannot traverse links unless they are configured for the topology designated for that service. Figure 5 on page 16 shows a diagram of this case. Contiguous paths for routing the voice topology are shown with dotted lines, and paths for routing the video topology are shown with dashed lines.

For a complete set of configurations for all of the devices in the topology, see “CLI Quick Configuration” on page 16. The remainder of the example focuses on Device CE1 and Device PE1.

Figure 5 on page 16 shows the sample topology.

**Figure 5: Multitopology OSPF and IBGP for Designating Links Belonging to Voice and Video Services**



### Configuration

- [Configuring Device CE1 on page 23](#)
- [Configuring Device PE1 on page 26](#)

#### 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 CE1

```
set interfaces fe-0/1/0 fastether-options loopback
set interfaces fe-0/1/0 unit 0 family inet address 11.19.130.1/24
set interfaces fe-0/1/0 unit 0 family inet address 11.19.131.1/24
set interfaces fe-0/1/0 unit 0 family inet address 11.19.132.1/24
set interfaces fe-1/2/0 unit 1 description to-PE1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 97 family inet address 10.255.165.97/32 primary
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.1
set protocols bgp group ebgp export set_community
set protocols bgp group ebgp export inject_directs
set protocols bgp group ebgp peer-as 100
set protocols bgp group ebgp neighbor 10.0.0.2
set policy-options policy-statement inject_directs term a from protocol direct
```

```

set policy-options policy-statement inject_directs term a from interface fe-0/1/0.0
set policy-options policy-statement inject_directs term a then next policy
set policy-options policy-statement inject_directs term a then accept
set policy-options policy-statement inject_directs term b then reject
set policy-options policy-statement set_community term a from route-filter 11.19.130.0/24
  exact
set policy-options policy-statement set_community term a from route-filter 11.19.131.0/24
  exact
set policy-options policy-statement set_community term a then community add voice
set policy-options policy-statement set_community term a then accept
set policy-options policy-statement set_community term b from route-filter 11.19.132.0/24
  exact
set policy-options policy-statement set_community term b from route-filter 11.19.133.0/24
  exact
set policy-options policy-statement set_community term b then community add video
set policy-options policy-statement set_community term b then accept
set policy-options policy-statement set_community term default then accept
set policy-options community video members target:50:50
set policy-options community voice members target:40:40
set routing-options autonomous-system 101

```

**Device CE2**

```

set interfaces fe-0/1/1 fastether-options loopback
set interfaces fe-0/1/1 unit 0 family inet address 11.19.140.1/24
set interfaces fe-0/1/1 unit 0 family inet address 11.19.141.1/24
set interfaces fe-0/1/1 unit 0 family inet address 11.19.142.1/24
set interfaces fe-1/2/0 unit 46 description to-PE2
set interfaces fe-1/2/0 unit 46 family inet address 10.0.0.46/30
set interfaces lo0 unit 20 family inet address 10.255.165.20/32 primary
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.46
set protocols bgp group ebgp export set_community
set protocols bgp group ebgp export inject_directs
set protocols bgp group ebgp peer-as 100
set protocols bgp group ebgp neighbor 10.0.0.45
set policy-options policy-statement inject_directs term a from protocol direct
set policy-options policy-statement inject_directs term a from interface fe-0/1/1.0
set policy-options policy-statement inject_directs term a then next policy
set policy-options policy-statement inject_directs term a then accept
set policy-options policy-statement inject_directs term b then reject
set policy-options policy-statement set_community term a from route-filter 11.19.140.0/24
  exact
set policy-options policy-statement set_community term a from route-filter 11.19.141.0/24
  exact
set policy-options policy-statement set_community term a then community add voice
set policy-options policy-statement set_community term a then accept
set policy-options policy-statement set_community term b from route-filter 11.19.142.0/24
  exact
set policy-options policy-statement set_community term b from route-filter 11.19.143.0/24
  exact
set policy-options policy-statement set_community term b then community add video
set policy-options policy-statement set_community term b then accept
set policy-options policy-statement set_community term default then accept
set policy-options community video members target:50:50
set policy-options community voice members target:40:40
set routing-options autonomous-system 102

```

Device PE1

```
set interfaces fe-1/2/0 unit 2 description to-CE1
set interfaces fe-1/2/0 unit 2 family inet filter input ef_path
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 6 description to-P1
set interfaces fe-1/2/1 unit 6 family inet filter input ef_path
set interfaces fe-1/2/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-1/2/2 unit 9 description to-P3
set interfaces fe-1/2/2 unit 9 family inet filter input ef_path
set interfaces fe-1/2/2 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 93 family inet address 10.255.165.93/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.93
set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.2
set protocols bgp group ebgp family inet unicast topology voice community target:40:40
set protocols bgp group ebgp family inet unicast topology video community target:50:50
set protocols bgp group ebgp peer-as 101
set protocols bgp group ebgp neighbor 10.0.0.1
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/2.9 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/2.9 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/2.9 topology video
set protocols ospf area 0.0.0.0 interface lo0.93 passive
set policy-options policy-statement nhs then next-hop self
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term ef then topology voice
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 2 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 6 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 9 classifiers inet-precedence default
```

Device PE2

```
set interfaces fe-1/2/0 unit 38 description to-P2
set interfaces fe-1/2/0 unit 38 family inet filter input ef_path
set interfaces fe-1/2/0 unit 38 family inet address 10.0.0.38/30
set interfaces fe-1/2/1 unit 42 description to-P4
```

```

set interfaces fe-1/2/1 unit 42 family inet filter input ef_path
set interfaces fe-1/2/1 unit 42 family inet address 10.0.0.42/30
set interfaces fe-1/2/2 unit 45 description to-CE2
set interfaces fe-1/2/2 unit 45 family inet filter input ef_path
set interfaces fe-1/2/2 unit 45 family inet address 10.0.0.45/30
set interfaces lo0 unit 203 family inet address 10.255.165.203/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.203
set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.45
set protocols bgp group ebgp family inet unicast topology voice community target:40:40
set protocols bgp group ebgp family inet unicast topology video community target:50:50
set protocols bgp group ebgp peer-as 102
set protocols bgp group ebgp neighbor 10.0.0.46
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/0.38 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/0.38 topology video metric 200
set protocols ospf area 0.0.0.0 interface fe-1/2/0.38 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/1.42 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.42 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.42 topology video
set protocols ospf area 0.0.0.0 interface lo0.203 passive
set policy-options policy-statement nhs then next-hop self
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term ef then topology voice
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 38 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 42 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 45 classifiers inet-precedence default

```

**Device P1**

```

set interfaces fe-1/2/0 unit 5 description to-PE1
set interfaces fe-1/2/0 unit 5 family inet filter input ef_path
set interfaces fe-1/2/0 unit 5 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 13 description to-P3
set interfaces fe-1/2/1 unit 13 family inet filter input ef_path
set interfaces fe-1/2/1 unit 13 family inet address 10.0.0.13/30
set interfaces fe-1/2/2 unit 17 description to-P4
set interfaces fe-1/2/2 unit 17 family inet filter input ef_path
set interfaces fe-1/2/2 unit 17 family inet address 10.0.0.17/30

```

```

set interfaces fe-1/2/3 unit 33 description to-P2
set interfaces fe-1/2/3 unit 33 family inet filter input ef_path
set interfaces fe-1/2/3 unit 33 family inet address 10.0.0.33/30
set interfaces lo0 unit 99 family inet address 10.255.165.99/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.99
set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/3.33 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/3.33 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/3.33 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/2.17 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/2.17 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/2.17 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.13 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.13 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.13 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5 topology video disable
set protocols ospf area 0.0.0.0 interface lo0.99 passive
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term ef then topology voice
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 5 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 13 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 17 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/3 unit 33 classifiers inet-precedence default

```

**Device P2**

```

set interfaces fe-1/2/0 unit 22 description to-P3
set interfaces fe-1/2/0 unit 22 family inet filter input ef_path
set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30
set interfaces fe-1/2/1 unit 25 description to-P4
set interfaces fe-1/2/1 unit 25 family inet filter input ef_path
set interfaces fe-1/2/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 34 description to-P1
set interfaces fe-1/2/2 unit 34 family inet filter input ef_path
set interfaces fe-1/2/2 unit 34 family inet address 10.0.0.34/30
set interfaces fe-1/2/3 unit 37 description to-PE2
set interfaces fe-1/2/3 unit 37 family inet filter input ef_path
set interfaces fe-1/2/3 unit 37 family inet address 10.0.0.37/30

```

```

set interfaces lo0 unit 113 family inet address 10.255.165.113/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.113
set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/2.34 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/2.34 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/2.34 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.22 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/0.22 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.22 topology video metric 20
set protocols ospf area 0.0.0.0 interface fe-1/2/1.25 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.25 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.25 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/3.37 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/3.37 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/3.37 topology video metric 200
set protocols ospf area 0.0.0.0 interface lo0.113 passive
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term ef then topology voice
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 22 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 25 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 34 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/3 unit 37 classifiers inet-precedence default

```

**Device P3**

```

set interfaces fe-1/2/0 unit 10 description to-PE1
set interfaces fe-1/2/0 unit 10 family inet filter input ef_path
set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-1/2/1 unit 14 description to-P1
set interfaces fe-1/2/1 unit 14 family inet filter input ef_path
set interfaces fe-1/2/1 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/2/2 unit 21 description to-P2
set interfaces fe-1/2/2 unit 21 family inet filter input ef_path
set interfaces fe-1/2/2 unit 21 family inet address 10.0.0.21/30
set interfaces fe-1/2/3 unit 29 description to-P4
set interfaces fe-1/2/3 unit 29 family inet filter input ef_path
set interfaces fe-1/2/3 unit 29 family inet address 10.0.0.29/30
set interfaces lo0 unit 111 family inet address 10.255.165.111/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.111

```

```

set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/3.29 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/3.29 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/3.29 topology video
set protocols ospf area 0.0.0.0 interface fe-1/2/2.21 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/2.21 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/2.21 topology video metric 20
set protocols ospf area 0.0.0.0 interface fe-1/2/1.14 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.14 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.14 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10 topology video
set protocols ospf area 0.0.0.0 interface lo0.111 passive
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term ef then topology voice
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 10 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 14 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 21 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/3 unit 29 classifiers inet-precedence default

```

**Device P4**

```

set interfaces fe-1/2/0 unit 18 description to-P1
set interfaces fe-1/2/0 unit 18 family inet filter input ef_path
set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30
set interfaces fe-1/2/1 unit 26 description to-P2
set interfaces fe-1/2/1 unit 26 family inet filter input ef_path
set interfaces fe-1/2/1 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/2 unit 30 description to-P3
set interfaces fe-1/2/2 unit 30 family inet filter input ef_path
set interfaces fe-1/2/2 unit 30 family inet address 10.0.0.30/30
set interfaces fe-1/2/3 unit 41 description to-PE2
set interfaces fe-1/2/3 unit 41 family inet filter input ef_path
set interfaces fe-1/2/3 unit 41 family inet address 10.0.0.41/30
set interfaces lo0 unit 95 family inet address 10.255.165.95/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.95
set protocols bgp group ibgp family inet unicast topology voice community target:40:40
set protocols bgp group ibgp family inet unicast topology video community target:50:50
set protocols bgp group ibgp neighbor 10.255.165.93

```

```

set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology voice topology-id 126
set protocols ospf topology video topology-id 52
set protocols ospf area 0.0.0.0 interface fe-1/2/2.30 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/2.30 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/2.30 topology video
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18 topology video metric 20
set protocols ospf area 0.0.0.0 interface fe-1/2/1.26 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.26 topology voice disable
set protocols ospf area 0.0.0.0 interface fe-1/2/1.26 topology video disable
set protocols ospf area 0.0.0.0 interface fe-1/2/3.41 metric 10
set protocols ospf area 0.0.0.0 interface fe-1/2/3.41 topology voice
set protocols ospf area 0.0.0.0 interface fe-1/2/3.41 topology video
set protocols ospf area 0.0.0.0 interface lo0.95 passive
set routing-options autonomous-system 100
set routing-options topologies family inet topology voice
set routing-options topologies family inet topology video
set firewall family inet filter ef_path term ef from forwarding-class expedited-forwarding
set firewall family inet filter ef_path term video from source-address 11.19.132.0/24
set firewall family inet filter ef_path term video from source-address 11.19.133.0/24
set firewall family inet filter ef_path term video from source-address 11.19.142.0/24
set firewall family inet filter ef_path term video from source-address 11.19.144.0/24
set firewall family inet filter ef_path term video then topology video
set firewall family inet filter ef_path term default then accept
set class-of-service interfaces fe-1/2/0 unit 18 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/1 unit 26 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/2 unit 30 classifiers inet-precedence default
set class-of-service interfaces fe-1/2/3 unit 41 classifiers inet-precedence default

```

### Configuring Device CE1

#### 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 CE1:

1. Configure the interfaces.

For demonstration purposes, the example places an Ethernet interface into loopback mode and configures several addresses on this loopback interface. The addresses are then announced to the network.

```

[edit interfaces]
user@CE1# set fe-0/1/0 fastether-options loopback
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.130.1/24
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.131.1/24
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.132.1/24

```

```
user@CE1# set fe-1/2/0 unit 1 description to-PE1
user@CE1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@CE1# set lo0 unit 97 family inet address 10.255.165.97/32 primary
```

2. Configure the external BGP (EBGP) connection to Device PE1.

```
[edit protocols bgp group ebgp]
user@CE1# set type external
user@CE1# set local-address 10.0.0.1
user@CE1# set peer-as 100
user@CE1# set neighbor 10.0.0.2
```

3. Configure the routing policy that announces the addresses that are configured on interface fe-0/1/0.

```
[edit policy-options policy-statement inject_directs]
user@CE1# set term a from protocol direct
user@CE1# set term a from interface fe-0/1/0.0
user@CE1# set term a then next policy
user@CE1# set term a then accept
user@CE1# set term b then reject
```

4. Configure the routing policy that tags voice routes with the video community attribute, and video routes with the voice community attribute.

```
[edit policy-options policy-statement set_community]
user@CE1# set term a from route-filter 11.19.130.0/24 exact
user@CE1# set term a from route-filter 11.19.131.0/24 exact
user@CE1# set term a then community add voice
user@CE1# set term a then accept
```

```
user@CE1# set term b from route-filter 11.19.132.0/24 exact
user@CE1# set term b from route-filter 11.19.133.0/24 exact
user@CE1# set term b then community add video
user@CE1# set term b then accept
```

```
user@CE1# set term default then accept
```

```
[edit policy-options community]
user@CE1# set video members target:50:50
user@CE1# set voice members target:40:40
```

5. Apply the **set\_community** export policy so that direct routes are exported from the routing table into BGP.

Apply the **inject\_directs** export policy to announce the addresses that are configured on interface fe-0/1/0.

```
[edit protocols bgp group ebgp]
user@CE1# set export set_community
user@CE1# set export inject_directs
```

6. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@CE1# set autonomous-system 101
```

**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@CE1# show interfaces
fe-0/1/0 {
  fastether-options {
    loopback;
  }
  unit 0 {
    family inet {
      address 11.19.130.1/24;
      address 11.19.131.1/24;
      address 11.19.132.1/24;
    }
  }
}
fe-1/2/0 {
  unit 1 {
    description to-PE1;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 97 {
    family inet {
      address 10.255.165.97/32 {
        primary;
      }
    }
  }
}

user@CE1# show protocols
bgp {
  group ebgp {
    type external;
    local-address 10.0.0.1;
    export [ set_community inject_directs ];
    peer-as 100;
    neighbor 10.0.0.2;
  }
}

user@CE1# show policy-options
policy-statement inject_directs {
  term a {
    from {
      protocol direct;
      interface fe-0/1/0.0;
    }
    then {
      next policy;
    }
  }
}

```

```
        accept;
    }
}
term b {
    then reject;
}
}
policy-statement set_community {
    term a {
        from {
            route-filter 11.19.130.0/24 exact;
            route-filter 11.19.131.0/24 exact;
        }
        then {
            community add voice;
            accept;
        }
    }
    term b {
        from {
            route-filter 11.19.132.0/24 exact;
            route-filter 11.19.133.0/24 exact;
        }
        then {
            community add video;
            accept;
        }
    }
    term default {
        then accept;
    }
}
community video members target:50:50;
community voice members target:40:40;

user@CE1# show routing-options
autonomous-system 101;
```

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

### ***Configuring Device PE1***

**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 PE1:

1. Configure the interfaces.

The forwarding plane uses a firewall filter to indicate which topology forwarding table traffic should use. In this case, you must configure a firewall filter on all interfaces related to routing topologies. In general, all multitopology OSPF interfaces in the core where topologies are configured have input firewall filters. In addition, the ingress interfaces, where traffic from a CE device enters a PE device toward the core, have firewall filters configured. This configuration on Device PE1 shows a

firewall filter applied to the ingress interface (connected to the CE device) and to the two core-facing interfaces (connected to Device P1 and Device P3).

```
[edit interfaces]
```

```
user@PE1# set fe-1/2/0 unit 2 description to-CE1
user@PE1# set fe-1/2/0 unit 2 family inet filter input ef_path
user@PE1# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30
```

```
user@PE1# set fe-1/2/1 unit 6 description to-P1
user@PE1# set fe-1/2/1 unit 6 family inet filter input ef_path
user@PE1# set fe-1/2/1 unit 6 family inet address 10.0.0.6/30
```

```
user@PE1# set fe-1/2/2 unit 9 description to-P3
user@PE1# set fe-1/2/2 unit 9 family inet filter input ef_path
user@PE1# set fe-1/2/2 unit 9 family inet address 10.0.0.9/30
```

```
user@PE1# set lo0 unit 93 family inet address 10.255.165.93/32 primary
```

2. Configure the autonomous system (AS) number.

```
[edit routing-options]
```

```
user@PE1# set autonomous-system 100
```

3. Configure BGP.

```
[edit protocols bgp group ibgp]
```

```
user@PE1# set type internal
user@PE1# set local-address 10.255.165.93
user@PE1# set neighbor 10.255.165.111
user@PE1# set neighbor 10.255.165.203
user@PE1# set neighbor 10.255.165.113
user@PE1# set neighbor 10.255.165.95
user@PE1# set neighbor 10.255.165.99
```

```
[edit protocols bgp group ebgp]
```

```
user@PE1# set type external
user@PE1# set local-address 10.0.0.2
user@PE1# set peer-as 101
user@PE1# set neighbor 10.0.0.1
```

4. Configure a next-hop self routing policy to ensure that the IBGP devices use the loopback address on Device PE1 as the next-hop address on all IBGP route advertisements.

This way, Device PE1 serves as the gateway router for EBGp routes.

```
[edit policy-options policy-statement nhs]
```

```
user@PE1# set then next-hop self
```

5. Apply the next-hop self policy to the IBGP sessions.

```
[edit protocols bgp group ibgp]
```

```
user@PE1# set export nhs
```

6. Configure the voice and video topologies, which enable you to use these topologies with OSPF and BGP.

The names **voice** and **video** are local to the router. The names are not propagated beyond this router. However, for management purposes, a consistent naming scheme across routers in a multitopology environment is convenient.

```
[edit routing-options topologies family inet]
user@PE1# set topology voice
user@PE1# set topology video
```

7. Apply the community tags to identify the voice and video topologies by configuring a routing topology name and BGP community value.

In Junos OS, multitopology support for BGP is based on the community value in a BGP route. This configuration determines the association between a topology and one or more community values and populates the topology routing tables. Arriving BGP updates that have a matching community value are replicated in the associated topology routing table. You decide which BGP community values are associated with a given topology.

This configuration causes BGP updates that are received with community value **target:40:40** to be added into topology routing table **:voice.inet.0** (in addition to the default routing table **inet.0**). Updates that are received with community value **target:50:50** are added into topology routing table **:video.inet.0** (in addition to the default routing table **inet.0**).

```
[edit protocols bgp group ibgp family inet unicast]
user@PE1# set topology voice community target:40:40
user@PE1# set topology video community target:50:50
```

```
[edit protocols bgp group ebgp family inet unicast]
user@PE1# set topology voice community target:40:40
user@PE1# set topology video community target:50:50
```

8. Enable and disable multitopology OSPF on particular interfaces.

Enable multitopology OSPF designations only on desired interfaces, as shown in [Figure 5 on page 16](#). On interface fe-1/2/1.6 facing Device P1, enable the voice topology, and disable the video topology. On interface fe-1/2/2.9 facing Device P3, enable the video topology, and disable the voice topology.

When a topology ID is configured under OSPF, the topology is automatically enabled on all interfaces under OSPF. To disable a topology or to add a metric, you must add an explicit configuration.

For readability purposes, each topology is configured under each desired OSPF interface even though this default behavior occurs when the topology ID is configured. Configure higher metric values on a link to make the link less preferred than another available link.

```
[edit protocols ospf ]
user@PE1# set topology voice topology-id 126
user@PE1# set topology video topology-id 52
```

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface fe-1/2/1.6 metric 10
user@PE1# set interface fe-1/2/1.6 topology video disable
user@PE1# set interface fe-1/2/1.6 topology voice
```

```

user@PE1# set interface fe-1/2/2.9 metric 10
user@PE1# set interface fe-1/2/2.9 topology voice disable
user@PE1# set interface fe-1/2/2.9 topology video

```

```

user@PE1# set interface lo0.93 passive

```

9. Configure the firewall filter.

After routing topologies are configured, traffic must go through a firewall filter to make use of routing topology forwarding tables. For basic routing topologies, where traffic is first entering the core network, apply an input firewall filter to the ingress interface. Additionally, add firewall filters to interfaces where multitopology OSPF is configured. All routers must use the same firewall filter to associate packets with a topology to ensure consistent forwarding and to avoid routing loops or packet loss.

The forwarding plane handles traffic as it enters the router and exits out a particular interface. To inspect traffic and use a specified topology forwarding table to perform next-hop lookups, configure an input firewall filter on each interface where routing topology support is desired. Use a regular firewall filter to identify packet characteristics.

In general, for application-level differentiation, it is convenient to use DiffServ code points (DSCPs). When there is a firewall filter match, the firewall instructs the route lookup to use a particular topology forwarding table. Packet attributes are identified in the **from** clause, followed by a **then** clause indicating the topology forwarding table for use in forwarding next-hop lookups. This configuration notifies the router which traffic uses a routing topology forwarding table and which traffic uses the default forwarding table. The last term, which is named **default**, specifies the use of the default forwarding table.

These firewall configurations show source addresses and DSCPs used to sort voice, video, and default traffic. DSCPs are practical because you can set them at or near a CE device and because the information is intact across the network. For instance, here class of service (CoS) is configured for expedited traffic. DSCPs are also practical when the same IP address is used for different applications.

```

[edit firewall family inet filter ef_path]
user@PE1# set term ef from forwarding-class expedited-forwarding
user@PE1# set term ef then topology voice

```

```

user@PE1# set term video from source-address 11.19.132.0/24
user@PE1# set term video from source-address 11.19.133.0/24
user@PE1# set term video from source-address 11.19.142.0/24
user@PE1# set term video from source-address 11.19.144.0/24
user@PE1# set term video then topology video

```

```

user@PE1# set term default then accept

```

10. Enable CoS on the interfaces.

```

[edit class-of-service interfaces]
user@PE1# set fe-1/2/0 unit 2 classifiers inet-precedence default
user@PE1# set fe-1/2/1 unit 6 classifiers inet-precedence default

```

```
user@PE1# set fe-1/2/2 unit 9 classifiers inet-precedence default
```

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

```
user@PE1# show interfaces
fe-1/2/0 {
  unit 2 {
    description to-CE1;
    family inet {
      filter {
        input ef_path;
      }
      address 10.0.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 6 {
    description to-P1;
    family inet {
      filter {
        input ef_path;
      }
      address 10.0.0.6/30;
    }
  }
}
fe-1/2/2 {
  unit 9 {
    description to-P3;
    family inet {
      filter {
        input ef_path;
      }
      address 10.0.0.9/30;
    }
  }
}
lo0 {
  unit 93 {
    family inet {
      address 10.255.165.93/32 {
        primary;
      }
    }
  }
}

user@PE1# show protocols
bgp {
  group ibgp {
    type internal;
```

```

local-address 10.255.165.93;
family inet {
    unicast {
        topology voice {
            community target:40:40;
        }
        topology video {
            community target:50:50;
        }
    }
}
export nhs;
neighbor 10.255.165.111;
neighbor 10.255.165.203;
neighbor 10.255.165.113;
neighbor 10.255.165.95;
neighbor 10.255.165.99;
}
group ebgp {
    type external;
    local-address 10.0.0.2;
    family inet {
        unicast {
            topology voice {
                community target:40:40;
            }
            topology video {
                community target:50:50;
            }
        }
    }
    peer-as 101;
    neighbor 10.0.0.1;
}
}
ospf {
    topology voice topology-id 126;
    topology video topology-id 52;
    area 0.0.0.0 {
        interface fe-1/2/1.6 {
            metric 10;
            topology video disable;
            topology voice;
        }
        interface fe-1/2/2.9 {
            metric 10;
            topology voice disable;
            topology video;
        }
        interface lo0.93 {
            passive;
        }
    }
}
}
user@PE1# show policy-options

```

```
policy-statement nhs {
  then {
    next-hop self;
  }
}

user@PE1# show routing-options
autonomous-system 100;
topologies {
  family inet {
    topology voice;
    topology video;
  }
}

user@PE1# show firewall
family inet {
  filter ef_path {
    term ef {
      from {
        forwarding-class expedited-forwarding;
      }
      then topology voice;
    }
    term video {
      from {
        source-address {
          11.19.132.0/24;
          11.19.133.0/24;
          11.19.142.0/24;
          11.19.144.0/24;
        }
      }
      then topology video;
    }
    term default {
      then accept;
    }
  }
}

user@PE1# show class-of-service
interfaces {
  fe-1/2/0 {
    unit 2 {
      classifiers {
        inet-precedence default;
      }
    }
    unit 6 {
      classifiers {
        inet-precedence default;
      }
    }
    unit 9 {
      classifiers {
        inet-precedence default;
      }
    }
  }
}
```

```

    }
  }
}

```

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

### Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Interfaces on page 33](#)
- [Verifying the Routes on page 34](#)
- [Checking the Resolving BGP Next Hops on page 35](#)
- [Examining the Protocol Next Hop on page 36](#)
- [Verifying the OSPF Neighbor on page 36](#)
- [Checking the Router LSA on page 37](#)
- [Checking How Traffic Traverses the Network on page 38](#)

#### Verifying the OSPF Interfaces

**Purpose** Verify that the OSPF interfaces are configured to belong to one or more topologies.

**Action** From operational mode, enter the **show (ospf | ospf3) interface interface-name detail** command.

```

user@PE1> show ospf interface fe-1/2/1.6 detail

```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-1/2/1.6	DR	0.0.0.0	10.255.165.93	10.255.165.99	1

```

Type: LAN, Address: 10.0.0.6, Mask: 255.255.255.252, MTU: 1500, Cost: 10
DR addr: 10.0.0.6, BDR addr: 10.0.0.5, Priority: 128
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None
Topology default (ID 0) -> Cost: 10
Topology video (ID 52) -> Disabled, Cost: 10
Topology voice (ID 126) -> Cost: 10

```

**Meaning** This output shows that the voice topology was added to the fe-1/2/1.6 interface on Device PE1. The topology name is voice, and the MT-ID is 126. The video topology is disabled on this interface. The cost of the interface is 10.

The Router-LSA originated and flooded by the router includes all relevant topology information for specific interfaces, such as MT-ID and metric. If MTR is not configured on an OSPF interface, then the Router-LSA does not include any topology information for that interface. OSPF neighbors might or might not support multitopology OSPF. That is, a particular link is not used to calculate OSPF routes for a topology unless routers at both ends of the link announce that link as part of the topology. If multitopology OSPF is not supported in neighboring OSPF routers or is not configured to do so, then topology information in LSAs received by the neighbor is ignored.

**Verifying the Routes**

**Purpose** Make sure that the routes are in the expected routing tables and that the expected communities are attached to the routes.

**Action** From operational mode, enter the **show route detail** command on Device PE1.

```

user@PE1> show route 11.19.130.0/24 detail

inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 812
        Address: 0xb9f064c
        Next-hop reference count: 22
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.2, selected
        Session Id: 0x600004
        State: <Active Ext>
        Local AS: 100 Peer AS: 101
        Age: 3d 21:44:07
        Task: BGP_101.10.0.0.1+51873
        Announcement bits (3): 0-KRT 3-BGP_RT_Background 4-Resolve tree
3
        AS path: 101 I
        Communities: target:40:40
        Accepted
        Localpref: 100
        Router ID: 10.255.165.97
        Secondary Tables: :voice.inet.0

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 812
        Address: 0xb9f064c
        Next-hop reference count: 22
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.2, selected
        Session Id: 0x600004
        State: <Secondary Active IndepResolution Ext>
        Local AS: 100 Peer AS: 101
        Age: 3d 21:44:07
        Task: BGP_101.10.0.0.1+51873
        Announcement bits (2): 0-KRT 1-Resolve tree 1
        AS path: 101 I
        Communities: target:40:40
        Accepted
        Localpref: 100
        Router ID: 10.255.165.97
        Primary Routing Table inet.0

```

**Meaning** This output shows BGP route 11.19.130.0/24 with community value target:40:40. Because the route matches the criteria for the voice topology, it is added to both the default and voice topology routing tables (**inet.0** and **:voice.inet.0**). Device PE1 learns the route from Device CE1 through EBGp and then injects the route into IBGP.

*Checking the Resolving BGP Next Hops*

**Purpose** Check the protocol next hop and forwarding next hop.

**Action** From operational mode, enter the **show route detail** command on Device PE2.

```

user@PE2> show route 11.19.130.0/24 detail
inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
           Next hop type: Indirect
           Address: 0xb9f0e04
           Next-hop reference count: 12
           Source: 10.255.165.93
           Next hop type: Router, Next hop index: 262153
           Next hop: 10.0.0.37 via fe-1/2/0.38
           Session Id: 0x700004
           Next hop: 10.0.0.41 via fe-1/2/1.42, selected
           Session Id: 0x700005
           Protocol next hop: 10.255.165.93
           Indirect next hop: bb8c000 262154 INH Session ID: 0x700007
           State: <Active Int Ext>
           Local AS: 100 Peer AS: 100
           Age: 3d 4:27:40 Metric2: 30
           Task: BGP_100.10.255.165.93+179
           Announcement bits (3): 0-KRT 3-BGP_RT_Background 4-Resolve tree
3
           AS path: 101 I
           Communities: target:40:40
           Accepted
           Localpref: 100
           Router ID: 10.255.165.93
           Secondary Tables: :voice.inet.0

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
           Next hop type: Indirect
           Address: 0xb9f0f34
           Next-hop reference count: 6
           Source: 10.255.165.93
           Next hop type: Router, Next hop index: 1188
           Next hop: 10.0.0.37 via fe-1/2/0.38, selected
           Session Id: 0x700004
           Protocol next hop: 10.255.165.93
           Indirect next hop: bb8c1d8 262177 INH Session ID: 0x700007
           State: <Secondary Active IndepResolution Int Ext>
           Local AS: 100 Peer AS: 100
           Age: 3d 2:00:20 Metric2: 30
           Task: BGP_100.10.255.165.93+179
           Announcement bits (2): 0-KRT 1-Resolve tree 1
           AS path: 101 I
           Communities: target:40:40
           Accepted
           Localpref: 100
           Router ID: 10.255.165.93
           Primary Routing Table inet.0

```

**Meaning** A typical IBGP core has BGP routes with protocol next hops that resolve using the underlying IGP routes. IBGP routes in a topology routing table have protocol next-hop IP addresses. By default, the same topology routing table is used to look up and resolve the protocol next-hop IP address to a forwarding next hop. This output from Device PE2 shows the same BGP route as seen in the previous example: 11.19.130.0/24. The route is being shown from a different perspective, that is, from Device PE2 as an IBGP route. Similarly, this IBGP route is added to both **inet.0** and **:voice.inet.0** on Device PE2. However, while each route has the same protocol next hop, each route has a different forwarding next hop (ge-0/0/3.0 instead of ge-0/1/4.0). The reason for this difference is when the protocol next-hop IP address 10.255.165.93 is resolved, it uses the corresponding routing table (**inet.0** or **:voice.inet.0**) to look up the protocol next hop.

#### *Examining the Protocol Next Hop*

**Purpose** Check the protocol next hop and forwarding next hop.

**Action** From operational mode, enter the **show route** command on Device PE2.

```
user@PE2> show route 10.255.165.93
inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 04:37:26, metric 30
                  > to 10.0.0.37 via fe-1/2/0.38
                  to 10.0.0.41 via fe-1/2/1.42

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 02:10:04, metric 30
                  > to 10.0.0.37 via fe-1/2/0.38

:video.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 02:03:16, metric 30
                  > to 10.0.0.41 via fe-1/2/1.42
```

**Meaning** This output from Device PE2 shows the protocol next hop of 11.19.130.0/24, which is IP address 10.255.165.93, thus further demonstrating how IBGP route 11.19.130.0/24 resolves its protocol next hop. The forwarding next hops of 10.255.165.93 match the IBGP forwarding next hops of route 11.19.130.0/24 as shown in the previous example. Observe here that the IP address 10.255.165.93 is also in routing table **:video.inet.0**. This address is the loopback address of Device PE1, and as such, resides in all three routing tables. This example also shows how traffic entering Device PE2 destined to 11.19.130.0/24 exits out different interfaces depending on its associated topology. The actual traffic is marked in such a way that a firewall filter can direct the traffic to use a particular topology routing table.

#### *Verifying the OSPF Neighbor*

**Purpose** Make sure that the expected topologies are enabled on the OSPF neighbor.

**Action** From operational mode, enter the **show ospf neighbor 10.0.0.21 extensive** command on Device P2.

```
user@P2> show ospf neighbor 10.0.0.21 extensive
Address          Interface          State    ID                Pri  Dead
10.0.0.21        fe-1/2/0.22        Full    10.255.165.111   128   39
Area 0.0.0.0, opt 0x52, DR 10.0.0.22, BDR 10.0.0.21
Up 3d 06:09:50, adjacent 3d 06:09:50
Topology default (ID 0) -> Bidirectional
Topology video (ID 52) -> Bidirectional
```

**Meaning** This Device P2 output shows OSPF neighbor PE2 (10.0.0.21), where multitopology OSPF default and video are multitopology OSPF participants. The **Bidirectional** flag shows that the neighbor is configured using the same multitopology OSPF ID.

### *Checking the Router LSA*

**Purpose** Check the links where video and voice topologies are enabled.

**Action** From operational mode, enter the **show ospf database lsa-id 10.255.165.203 extensive** command on Device P2.

```
user@P2> show ospf database lsa-id 10.255.165.203 extensive

OSPF database, Area 0.0.0.0
Type      ID                Adv Rtr           Seq             Age  Opt  Cksum  Len
Router    10.255.165.203    10.255.165.203    0x800000063     1552 0x22 0xdff3  80
bits 0x0, link count 3
id 10.255.165.203, data 255.255.255.255, Type Stub (3)
Topology count: 2, Default metric: 0
Topology video (ID 52) -> Metric: 0
Topology voice (ID 126) -> Metric: 0
id 10.0.0.38, data 10.0.0.38, Type Transit (2)
Topology count: 2, Default metric: 10
Topology video (ID 52) -> Metric: 200
Topology voice (ID 126) -> Metric: 10
id 10.0.0.42, data 10.0.0.42, Type Transit (2)
Topology count: 1, Default metric: 10
Topology video (ID 52) -> Metric: 10
Topology default (ID 0)
Type: Transit, Node ID: 10.0.0.42
Metric: 10, Bidirectional
Type: Transit, Node ID: 10.0.0.38
Metric: 10, Bidirectional
Topology video (ID 52)
Type: Transit, Node ID: 10.0.0.42
Metric: 10, Bidirectional
Type: Transit, Node ID: 10.0.0.38
Metric: 200, Bidirectional
Topology voice (ID 126)
Type: Transit, Node ID: 10.0.0.38
Metric: 10, Bidirectional
Aging timer 00:34:08
Installed 00:25:49 ago, expires in 00:34:08, sent 00:25:47 ago
Last changed 3d 01:45:51 ago, Change count: 10
```

**Meaning** This Device P2 output shows the Router-LSA originated by Device PE2. The LSA shows links where video and voice topologies are enabled (in addition to the default topology).

**Checking How Traffic Traverses the Network**

**Purpose** Make sure that the expected paths are used.

**Action** From operational mode, enter the **traceroute** command on Device CE1.

The first example output shows a traceroute over the voice topology goes from Device CE1 to Device CE2 where DSCPs are set. The routes are resolved over **:voice.inet.0**. This traceroute path follows the voice path CE1-PE1-P1-P2-PE2-CE2

```
user@CE1> traceroute 11.19.140.1 source 11.19.130.1 tos 160
```

```
traceroute to 11.19.140.1 (11.19.140.1) from 11.19.130.1, 30 hops max, 40 byte packets
```

```
 1 10.0.0.2 (10.0.0.2)  2.015 ms  1.924 ms  1.770 ms
 2 10.0.0.5 (10.0.0.5)  1.890 ms  1.010 ms  0.974 ms
 3 10.0.0.34 (10.0.0.34) 0.986 ms  1.031 ms  0.973 ms
 4 10.0.0.38 (10.0.0.38) 1.213 ms  1.065 ms  1.154 ms
 5 11.19.140.1 (11.19.140.1) 1.696 ms  4.286 ms  1.332 ms
```

This output shows a traceroute from Device CE1 to Device CE2 for voice where no DSCPs are set. The routes are resolved over **inet.0**, and the resulting path is different from the previous case where the DSCPs are set. This traceroute path follows the default path CE1-PE1-P4-PE2-CE2.

```
user@CE1> traceroute 11.19.140.1 source 11.19.130.1
```

```
traceroute to 11.19.140.1 (11.19.140.1) from 11.19.130.1, 30 hops max, 40 byte packets
```

```
 1 10.0.0.2 (10.0.0.2)  1.654 ms  1.710 ms  1.703 ms
 2 10.0.0.5 (10.0.0.5)  1.790 ms  1.045 ms  0.975 ms
 3 10.0.0.18 (10.0.0.18) 0.989 ms  1.041 ms  0.983 ms
 4 10.0.0.42 (10.0.0.42) 0.994 ms  1.036 ms  1.002 ms
 5 11.19.140.1 (11.19.140.1) 1.329 ms  2.248 ms  2.225 ms
```

This output shows a traceroute from Device CE1 to Device CE2 for video traffic where the firewall filter is based on the destination address. The routes are resolved over **:video.inet.0**. This traceroute follows the video path CE1-PE1-P3-P4-PE2-CE2.

```
user@CE1> traceroute 11.19.142.1 source 11.19.132.1
```

```
traceroute to 11.19.142.1 (11.19.142.1) from 11.19.132.1, 30 hops max, 40 byte packets
```

```
 1 10.0.0.2 (10.0.0.2)  1.126 ms  1.300 ms  0.995 ms
 2 10.0.0.10 (10.0.0.10) 0.981 ms  1.018 ms  0.991 ms
 3 10.0.0.30 (10.0.0.30) 0.997 ms  1.886 ms  1.952 ms
 4 10.0.0.42 (10.0.0.42) 1.800 ms  1.038 ms  0.980 ms
 5 11.19.142.1 (11.19.142.1) 1.367 ms  1.352 ms  1.328 ms
```

This output shows a traceroute from Device CE1 to Device CE2 for video where DSCPs are set. The DSCP bits are directing Device PE1 to use the topology table **:voice.inet.0**. Because there is no entry in the voice routing table for video routes, traffic is dropped.

```
user@CE1> traceroute 11.19.142.1 source 11.19.132.1 tos 160
```

```
traceroute to 11.19.142.1 (11.19.142.1) from 11.19.132.1, 30 hops max, 40 byte
```

packets

1 10.0.0.2 (10.0.0.2) 1.135 ms !N 1.007 ms !N 0.954 ms !N

**Related Documentation**

- [Example: Configuring Multitopology Routing Based on a Multicast Source on page 39](#)

## Example: Configuring Multitopology Routing Based on a Multicast Source

- [Understanding Multitopology Routing Based on a Multicast Source on page 39](#)
- [Example: Configuring Multitopology Routing Based on a Multicast Source on page 41](#)


### Understanding Multitopology Routing Based on a Multicast Source

Protocol Independent Multicast (PIM), in conjunction with multitopology routing extensions to OSPF (multitopology OSPF) and BGP, can direct multicast traffic over particular paths based on traffic characteristics.

Junos OS provides a mechanism whereby multicast traffic traverses user-specified topology paths based on the sender's source address. Multitopology routing (MTR) is used for OSPF, BGP, and route resolution over the specified topology routing tables. OSPF and BGP independently populate the routing table used by PIM. Firewall filters are not required because the multicast forwarding plane uses the multicast tree after it has been built.

[Figure 6 on page 39](#) shows a diagram of routing topology paths, where the dashed lines are associated with multicast group A (topology red), and the dotted lines are associated with multicast group B (topology blue).

**Figure 6: Core Links Configured to Prefer Specified Routing Topologies**



Two copies of the same stream enter Device PE1 and then traverse separate paths over the internal BGP (IBGP) core.

This solution leverages Junos OS features that allow particular routing tables to perform route resolution using specified routing tables.

The configuration includes a combination of the following features:

- BGP communities
- Separate IBGP next hops belonging to user-specified OSPF routing topologies
- Route resolution over user-specified topology routing tables
- A separate routing table (**inet.2**) for multicast protocols

Commonly, networks use a separate routing table for multicast. In Junos OS, the multicast routing table is **inet.2**. Routing topologies are grouped based on BGP communities. Each group represents a set of IP addresses associated with multicast servers and receivers. Primarily, the group must be related to the set of servers because the multicast receivers initiate tree creation toward these servers. Multicast traffic directed downstream toward receivers uses the previously created PIM tree, and therefore the forwarding plane does not need to know about routing topologies.

PIM uses the **inet.2** routing table for lookups of multicast source addresses. These IP addresses used for tree creation are IP unicast addresses. The customer edge (CE) routers, nearest to the multicast servers, announce the multicast source IP addresses to the provider edge (PE) routers using external BGP (EBGP). They are announced with both **family inet** unicast and **family inet multicast**, thus causing the BGP route to be added to the default routing table **inet.0** and to **inet.2**.

Both versions of the route are injected by the PE router into IBGP. Each BGP route injected into IBGP has a specific protocol next hop. Junos OS provides the flexibility to set the protocol next hop when exporting the route into IBGP. For instance, a next-hop self can be set with an export policy configuration. You can also set the protocol next hop to a route associated with a specified topology routing table.

Keeping in mind that an EBGP route can have a community associated with a routing topology, you can conveniently configure a policy to use this community to designate which protocol next hop should be set when exporting the IBGP route into **inet.2**. As such, a specific protocol next-hop IP address is required for each topology on each router injecting IBGP routes. You can configure multiple secondary loopback IP addresses on a router to be used as protocol next-hop addresses.

A group of BGP routes associated with a routing topology use the same unique protocol next hop. For instance, if you configure a PE router to handle two routing topologies, you would also configure two unique nonprimary addresses under loopback interface lo0. Next, associate each nonprimary loopback IP address with a topology for inclusion in the associated topology routing table. Configure the loopback IP address and topology under an OSPF interface statement. You must specifically disable all other topologies known to OSPF for two reasons. First, the loopback address specific to a topology must reside in only one topology routing table. Second, once the topology is added to OSPF, the topology defaults to being enabled on all subsequent interfaces under OSPF.


You can specify up to two routing tables in the resolution configuration. A key element to this solution is that the protocol next-hop address resides in only one topology table. That is, the protocol next hop belongs to a remote PE secondary loopback address and is injected into only one topology table. The route resolution scheme first checks the

topology table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme then checks the second topology table. Hence, only one topology table is used for each protocol next-hop address.

Links can support all routing topologies to provide a backup path should a primary multicast path fail. You can configure specific OSPF link metrics on topologies to identify paths and build trees to different servers. When a multicast tree gets built with PIM join messages directed toward the source, it follows the most preferred path. A multicast tree to a different multicast source (in a different routing topology) can create another tree along a different path.

Figure 7 on page 41 shows an example of two trees using different paths over different topologies. It shows Server A using the multicast tree with the dashed line as its path and Server B using the multicast tree with the dotted line as its path.

Figure 7: Core Links Configured to Prefer Specified Routing Topologies



### Example: Configuring Multitopology Routing Based on a Multicast Source

This example shows how to use multitopology routing (MTR) to provide redundancy for multicast traffic over separate network paths. That is, two multicast sources send the same multicast stream, yet for redundancy purposes in the case of link failure, the two streams use disjoint paths.



**NOTE:** Note there is no standard defined at this time for using MTR extensions to PIM.

- [Requirements on page 42](#)
- [Overview on page 42](#)
- [Configuration on page 42](#)
- [Verification on page 60](#)

## Requirements

This example requires that Junos OS Release 9.0 or later is running on the provider core devices.


## Overview

Assume that each source providing redundant multicast streams, S1 and S2, have different IP subnet addresses. Each source sends multicast traffic using different groups: G1 and G2. Further, assume that S1 and S2 are attached to the same customer edge (CE) device and use BGP to announce routes to the provider edge (PE) router.

For a complete set of configurations for all of the devices in the topology, see “[CLI Quick Configuration](#)” on page 42. The remainder of the example focuses on Device CE1 and Device PE1.

Figure 8 on page 42 shows the sample topology.

**Figure 8: Multi-topology OSPF and BGP for Designating Links Belonging to Voice and Video Services**



## Configuration

- [Configuring Device CE1 on page 49](#)
- [Configuring Device PE1 on page 53](#)

### 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 CE1

```
set interfaces fe-0/1/0 fastether-options loopback
set interfaces fe-0/1/0 unit 0 family inet address 11.19.130.1/24
set interfaces fe-0/1/0 unit 0 family inet address 11.19.131.1/24
set interfaces fe-0/1/0 unit 0 family inet address 11.19.132.1/24
set interfaces ge-1/2/0 unit 1 description to-PE1
set interfaces ge-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 97 family inet address 10.255.165.97/32 primary
set protocols bgp group ebgp type external
```

```

set protocols bgp group ebgp local-address 10.0.0.1
set protocols bgp group ebgp family inet unicast
set protocols bgp group ebgp family inet multicast
set protocols bgp group ebgp export set_community
set protocols bgp group ebgp export inject_directs
set protocols bgp group ebgp peer-as 100
set protocols bgp group ebgp neighbor 10.0.0.2
set protocols pim interface fe-0/1/0.0 mode sparse
set protocols pim interface ge-1/2/0.1 mode sparse
set policy-options policy-statement inject_directs term a from protocol direct
set policy-options policy-statement inject_directs term a from interface fe-0/1/0.0
set policy-options policy-statement inject_directs term a then next policy
set policy-options policy-statement inject_directs term a then accept
set policy-options policy-statement inject_directs term b then reject
set policy-options policy-statement set_community term a from route-filter 11.19.130.0/24
    exact
set policy-options policy-statement set_community term a from route-filter 11.19.131.0/24
    exact
set policy-options policy-statement set_community term a then community add red
set policy-options policy-statement set_community term a then accept
set policy-options policy-statement set_community term b from route-filter 11.19.132.0/24
    exact
set policy-options policy-statement set_community term b from route-filter 11.19.133.0/24
    exact
set policy-options policy-statement set_community term b then community add blue
set policy-options policy-statement set_community term b then accept
set policy-options policy-statement set_community term default then accept
set policy-options community blue members target:50:50
set policy-options community red members target:40:40
set routing-options interface-routes rib-group inet if-rib
set routing-options static route 10.0.0.0/16 next-hop 10.0.0.2
set routing-options rib-groups inet.2 import-rib inet.0
set routing-options rib-groups if-rib import-rib inet.0
set routing-options rib-groups if-rib import-rib inet.2
set routing-options rib-groups if-rib import-policy inject_directs
set routing-options autonomous-system 101

```

Device CE2

```

set interfaces fe-0/1/1 unit 0
set interfaces ge-1/2/0 unit 46 description to-PE2
set interfaces ge-1/2/0 unit 46 family inet address 10.0.0.46/30
set interfaces lo0 unit 20 family inet address 10.255.165.20/32 primary
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.46
set protocols bgp group ebgp peer-as 100
set protocols bgp group ebgp neighbor 10.0.0.45
set routing-options autonomous-system 102

```

Device PE1

```

set interfaces ge-1/2/0 unit 2 description to-CE1
set interfaces ge-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 6 description to-P1
set interfaces ge-1/2/1 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/2 unit 9 description to-P3
set interfaces ge-1/2/2 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 93 family inet address 10.255.165.93/32 primary
set interfaces lo0 unit 93 family inet address 1.1.1.30/32

```

```
set interfaces lo0 unit 93 family inet address 2.2.2.30/32
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.93
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp export nhs_test
set protocols bgp group ibgp export nhs_inet0_self
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.2
set protocols bgp group ebgp family inet unicast
set protocols bgp group ebgp family inet multicast
set protocols bgp group ebgp peer-as 101
set protocols bgp group ebgp neighbor 10.0.0.1
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52
set protocols ospf area 0.0.0.0 interface ge-1/2/1.6 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/1.6 topology blue metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/1.6 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/2.9 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/2.9 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/2.9 topology blue
set protocols ospf area 0.0.0.0 interface lo0.93 passive
set protocols ospf area 0.0.0.0 interface 1.1.1.30 topology red
set protocols ospf area 0.0.0.0 interface 1.1.1.30 topology blue disable
set protocols ospf area 0.0.0.0 interface 2.2.2.30 topology blue
set protocols ospf area 0.0.0.0 interface 2.2.2.30 topology red disable
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/0.2 mode sparse
set protocols pim interface ge-1/2/1.6 mode sparse
set protocols pim interface ge-1/2/2.9 mode sparse
set policy-options policy-statement nhs_inet0_self term a from protocol bgp
set policy-options policy-statement nhs_inet0_self term a from rib inet.0
set policy-options policy-statement nhs_inet0_self term a then next-hop self
set policy-options policy-statement nhs_test term a from protocol bgp
set policy-options policy-statement nhs_test term a from community red
set policy-options policy-statement nhs_test term a then next-hop 1.1.1.30
set policy-options policy-statement nhs_test term a then next policy
set policy-options policy-statement nhs_test term a then accept
set policy-options policy-statement nhs_test term b from protocol bgp
set policy-options policy-statement nhs_test term b from community blue
set policy-options policy-statement nhs_test term b then next-hop 2.2.2.30
set policy-options policy-statement nhs_test term b then next policy
set policy-options policy-statement nhs_test term b then accept
set policy-options policy-statement nhs_test term c then next-hop self
set policy-options community blue members target:50:50
set policy-options community red members target:40:40
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options resolution rib inet.2 resolution-ribs :red.inet.0
set routing-options resolution rib inet.2 resolution-ribs :blue.inet.0
set routing-options topologies family inet topology red
```

```
set routing-options topologies family inet topology blue
```

#### Device PE2

```
set interfaces ge-1/2/0 unit 38 description to-P2
set interfaces ge-1/2/0 unit 38 family inet address 10.0.0.38/30
set interfaces ge-1/2/1 unit 42 description to-P4
set interfaces ge-1/2/1 unit 42 family inet address 10.0.0.42/30
set interfaces ge-1/2/2 unit 45 description to-CE2
set interfaces ge-1/2/2 unit 45 family inet address 10.0.0.45/30
set interfaces lo0 unit 203 family inet address 10.255.165.203/32 primary
set interfaces lo0 unit 203 family inet address 1.1.1.40/32
set interfaces lo0 unit 203 family inet address 2.2.2.40/32
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.203
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp export nhs_test
set protocols bgp group ibgp export nhs_inet0_self
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols bgp group ebgp type external
set protocols bgp group ebgp local-address 10.0.0.45
set protocols bgp group ebgp family inet unicast
set protocols bgp group ebgp family inet multicast
set protocols bgp group ebgp peer-as 102
set protocols bgp group ebgp neighbor 10.0.0.46
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52
set protocols ospf area 0.0.0.0 interface ge-1/2/0.38 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/0.38 topology blue metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/0.38 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/1.42 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/1.42 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/1.42 topology blue
set protocols ospf area 0.0.0.0 interface lo0.203 passive
set protocols ospf area 0.0.0.0 interface 1.1.1.40 topology red
set protocols ospf area 0.0.0.0 interface 1.1.1.40 topology blue disable
set protocols ospf area 0.0.0.0 interface 2.2.2.40 topology red disable
set protocols ospf area 0.0.0.0 interface 2.2.2.40 topology blue
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/0.38 mode sparse
set protocols pim interface ge-1/2/1.42 mode sparse
set protocols pim interface ge-1/2/2.45 mode sparse
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement nhs_inet0_self term a from protocol bgp
set policy-options policy-statement nhs_inet0_self term a from rib inet.0
set policy-options policy-statement nhs_inet0_self term a then next-hop self
set policy-options policy-statement nhs_test term a from protocol bgp
set policy-options policy-statement nhs_test term a from community red
set policy-options policy-statement nhs_test term a then next-hop 1.1.1.40
set policy-options policy-statement nhs_test term a then next policy
set policy-options policy-statement nhs_test term a then accept
set policy-options policy-statement nhs_test term b from protocol bgp
set policy-options policy-statement nhs_test term b from community blue
```

```
set policy-options policy-statement nhs_test term b then next-hop 2.2.2.40
set policy-options policy-statement nhs_test term b then next policy
set policy-options policy-statement nhs_test term b then accept
set policy-options policy-statement nhs_test term c then next-hop self
set policy-options community blue members target:50:50
set policy-options community red members target:40:40
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options resolution rib inet.2 resolution-ribs :red.inet.0
set routing-options resolution rib inet.2 resolution-ribs :blue.inet.0
set routing-options topologies family inet topology red
set routing-options topologies family inet topology blue
```

Device P1

```
set interfaces ge-1/2/0 unit 5 description to-PE1
set interfaces ge-1/2/0 unit 5 family inet address 10.0.0.5/30
set interfaces ge-1/2/1 unit 13 description to-P3
set interfaces ge-1/2/1 unit 13 family inet address 10.0.0.13/30
set interfaces ge-1/2/2 unit 17 description to-P4
set interfaces ge-1/2/2 unit 17 family inet address 10.0.0.17/30
set interfaces ge-1/2/3 unit 33 description to-P2
set interfaces ge-1/2/3 unit 33 family inet address 10.0.0.33/30
set interfaces lo0 unit 99 family inet address 10.255.165.99/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.99
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52
set protocols ospf area 0.0.0.0 interface ge-1/2/3.33 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/3.33 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/3.33 topology blue metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/2.17
set protocols ospf area 0.0.0.0 interface ge-1/2/1.13
set protocols ospf area 0.0.0.0 interface ge-1/2/0.5 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/0.5 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/0.5 topology blue metric 1
set protocols ospf area 0.0.0.0 interface lo0.99 passive
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/3.33 mode sparse
set protocols pim interface ge-1/2/2.17 mode sparse
set protocols pim interface ge-1/2/1.13 mode sparse
set protocols pim interface ge-1/2/0.5 mode sparse
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options topologies family inet topology red
set routing-options topologies family inet topology blue
```

Device P2

```
set interfaces ge-1/2/0 unit 22 description to-P3
set interfaces ge-1/2/0 unit 22 family inet address 10.0.0.22/30
set interfaces ge-1/2/1 unit 25 description to-P4
```

```

set interfaces ge-1/2/1 unit 25 family inet address 10.0.0.25/30
set interfaces ge-1/2/2 unit 34 description to-P1
set interfaces ge-1/2/2 unit 34 family inet address 10.0.0.34/30
set interfaces ge-1/2/3 unit 37 description to-PE2
set interfaces ge-1/2/3 unit 37 family inet address 10.0.0.37/30
set interfaces lo0 unit 113 family inet address 10.255.165.113/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.113
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52
set protocols ospf area 0.0.0.0 interface ge-1/2/2.34 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/2.34 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/2.34 topology blue metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols ospf area 0.0.0.0 interface ge-1/2/1.25
set protocols ospf area 0.0.0.0 interface ge-1/2/3.37 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/3.37 topology red
set protocols ospf area 0.0.0.0 interface ge-1/2/3.37 topology blue metric 1
set protocols ospf area 0.0.0.0 interface lo0.113 passive
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/2.34 mode sparse
set protocols pim interface ge-1/2/0.22 mode sparse
set protocols pim interface ge-1/2/1.25 mode sparse
set protocols pim interface ge-1/2/3.37 mode sparse
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options topologies family inet topology red
set routing-options topologies family inet topology blue

```

Device P3

```

set interfaces ge-1/2/0 unit 10 description to-PE1
set interfaces ge-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces ge-1/2/1 unit 14 description to-P1
set interfaces ge-1/2/1 unit 14 family inet address 10.0.0.14/30
set interfaces ge-1/2/2 unit 21 description to-P2
set interfaces ge-1/2/2 unit 21 family inet address 10.0.0.21/30
set interfaces ge-1/2/3 unit 29 description to-P4
set interfaces ge-1/2/3 unit 29 family inet address 10.0.0.29/30
set interfaces lo0 unit 111 family inet address 10.255.165.111/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.111
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.95
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52

```

```

set protocols ospf area 0.0.0.0 interface ge-1/2/3.29 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/3.29 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/3.29 topology blue
set protocols ospf area 0.0.0.0 interface ge-1/2/2.21
set protocols ospf area 0.0.0.0 interface ge-1/2/1.14
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10 topology blue
set protocols ospf area 0.0.0.0 interface lo0.111 passive
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/3.29 mode sparse
set protocols pim interface ge-1/2/2.21 mode sparse
set protocols pim interface ge-1/2/1.14 mode sparse
set protocols pim interface ge-1/2/0.10 mode sparse
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options topologies family inet topology red
set routing-options topologies family inet topology blue

```

**Device P4**

```

set interfaces ge-1/2/0 unit 18 description to-P1
set interfaces ge-1/2/0 unit 18 family inet address 10.0.0.18/30
set interfaces ge-1/2/1 unit 26 description to-P2
set interfaces ge-1/2/1 unit 26 family inet address 10.0.0.26/30
set interfaces ge-1/2/2 unit 30 description to-P3
set interfaces ge-1/2/2 unit 30 family inet address 10.0.0.30/30
set interfaces ge-1/2/3 unit 41 description to-PE2
set interfaces ge-1/2/3 unit 41 family inet address 10.0.0.41/30
set interfaces lo0 unit 95 family inet address 10.255.165.95/32 primary
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.165.95
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet multicast
set protocols bgp group ibgp neighbor 10.255.165.93
set protocols bgp group ibgp neighbor 10.255.165.113
set protocols bgp group ibgp neighbor 10.255.165.203
set protocols bgp group ibgp neighbor 10.255.165.111
set protocols bgp group ibgp neighbor 10.255.165.99
set protocols ospf topology red topology-id 126
set protocols ospf topology blue topology-id 52
set protocols ospf area 0.0.0.0 interface ge-1/2/2.30 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/2.30 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/2.30 topology blue
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols ospf area 0.0.0.0 interface ge-1/2/1.26
set protocols ospf area 0.0.0.0 interface ge-1/2/3.41 metric 10
set protocols ospf area 0.0.0.0 interface ge-1/2/3.41 topology red metric 1
set protocols ospf area 0.0.0.0 interface ge-1/2/3.41 topology blue
set protocols ospf area 0.0.0.0 interface lo0.95 passive
set protocols pim rib-group inet mcast-rib
set protocols pim interface ge-1/2/2.30 mode sparse
set protocols pim interface ge-1/2/0.18 mode sparse
set protocols pim interface ge-1/2/1.26 mode sparse
set protocols pim interface ge-1/2/3.41 mode sparse
set routing-options rib-groups mcast-rib import-rib inet.2
set routing-options autonomous-system 100
set routing-options topologies family inet topology red

```

```
set routing-options topologies family inet topology blue
```

### Configuring Device CE1

**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 CE1:

1. Configure the interfaces.

For demonstration purposes, the example places an Ethernet interface into loopback mode and configures several addresses on this loopback interface. The addresses are then announced to the network as direct routes. These routes simulate a group of BGP routes with communities attached.

```
[edit interfaces]
user@CE1# set fe-0/1/0 fastether-options loopback
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.130.1/24
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.131.1/24
user@CE1# set fe-0/1/0 unit 0 family inet address 11.19.132.1/24
```

```
user@CE1# set ge-1/2/0 unit 1 description to-PE1
user@CE1# set ge-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@CE1# set lo0 unit 97 family inet address 10.255.165.97/32 primary
```

2. Configure the external BGP (EBGP) connection to Device PE1.

The CE router nearest to the multicast servers announces the multicast source IP addresses to the PE routers using EBGP. The source addresses are announced with both **family inet unicast** and **family inet multicast**, thus causing the BGP route to be added to the default routing table, **inet.0**, and to the multicast routing table, **inet.2**. Both sets of routes are injected by the PE router into IBGP.

```
[edit protocols bgp group ebgp]
user@CE1# set type external
user@CE1# set local-address 10.0.0.1
user@CE1# set family inet unicast
user@CE1# set family inet multicast
user@CE1# set peer-as 100
user@CE1# set neighbor 10.0.0.2
```

3. Configure PIM on the interfaces.

```
[edit protocols pim]
user@CE1# set interface fe-0/1/0.0 mode sparse
user@CE1# set interface ge-1/2/0.1 mode sparse
```

4. Configure the routing policy that announces the addresses that are configured on interface fe-0/1/0.

```
[edit policy-options policy-statement inject_directs]
user@CE1# set term a from protocol direct
user@CE1# set term a from interface fe-0/1/0.0
user@CE1# set term a then next policy
```

```
user@CE1# set term a then accept
user@CE1# set term b then reject
```

5. Configure the routing policy that tags some routes with the red community attribute and other routes with the blue community attribute.

The CE router advertises routes through EBGp to the PE router. These routes are advertised as BGP **family inet multicast** routes with communities set for two different groups. Policies identify the two groups of BGP routes.

```
[edit policy-options policy-statement set_community term a]
user@CE1# set from route-filter 11.19.130.0/24 exact
user@CE1# set from route-filter 11.19.131.0/24 exact
user@CE1# set then community add red
user@CE1# set then accept
```

```
[edit policy-options policy-statement set_community term b]
user@CE1# set from route-filter 11.19.132.0/24 exact
user@CE1# set from route-filter 11.19.133.0/24 exact
user@CE1# set then community add blue
user@CE1# set then accept
```

```
[edit policy-options policy-statement set_community term default]
user@CE1# set then accept
```

```
[edit policy-options]
user@CE1# set community blue members target:50:50
user@CE1# set community red members target:40:40
```

6. Apply the **set\_community** export policy so that the direct routes are exported into BGP.

Apply the **inject\_directs** export policy to announce the addresses that are configured on interface fe-0/1/0.

```
[edit protocols bgp group ebgp]
user@CE1# set export set_community
user@CE1# set export inject_directs
```

7. Use **rib-groups** to simulate a group of BGP routes with communities attached and announced as multicast routes.

This configuration creates a multicast routing table and causes PIM to use the multicast routing table **inet.2**.

```
[edit routing-options]
user@CE1# set interface-routes rib-group inet if-rib
```

```
user@CE1# set static route 10.0.0.0/16 next-hop 10.0.0.2
```

```
[edit routing-options rib-groups]
user@CE1# set inet.2 import-rib inet.0
user@CE1# set if-rib import-rib inet.0
user@CE1# set if-rib import-rib inet.2
user@CE1# set if-rib import-policy inject_directs
```

8. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@CE1# set autonomous-system 101
```

**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@CE1# show interfaces
fe-0/1/0 {
  fastether-options {
    loopback;
  }
  unit 0 {
    family inet {
      address 11.19.130.1/24;
      address 11.19.131.1/24;
      address 11.19.132.1/24;
    }
  }
}
ge-1/2/0 {
  unit 1 {
    description to-PE1;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 97 {
    family inet {
      address 10.255.165.97/32 {
        primary;
      }
    }
  }
}

user@CE1# show protocols
bgp {
  group ebgp {
    type external;
    local-address 10.0.0.1;
    family inet {
      unicast;
      multicast;
    }
    export [ set_community inject_directs ];
    peer-as 100;
    neighbor 10.0.0.2;
  }
}
pim {
```

```
interface fe-0/1/0.0 {
  mode sparse;
}
interface ge-1/2/0.1 {
  mode sparse;
}
}

user@CE1# show policy-options
policy-statement inject_directs {
  term a {
    from {
      protocol direct;
      interface fe-0/1/0.0;
    }
    then {
      next policy;
      accept;
    }
  }
  term b {
    then reject;
  }
}
policy-statement set_community {
  term a {
    from {
      route-filter 11.19.130.0/24 exact;
      route-filter 11.19.131.0/24 exact;
    }
    then {
      community add red;
      accept;
    }
  }
  term b {
    from {
      route-filter 11.19.132.0/24 exact;
      route-filter 11.19.133.0/24 exact;
    }
    then {
      community add blue;
      accept;
    }
  }
  term default {
    then accept;
  }
}
community blue members target:50:50;
community red members target:40:40;

user@CE1# show routing-options
interface-routes {
  rib-group inet if-rib;
}
static {
```

```

    route 10.0.0.0/16 next-hop 10.0.0.2;
  }
  rib-groups {
    inet.2 {
      import-rib inet.0;
    }
    if-rib {
      import-rib [ inet.0 inet.2 ];
      import-policy inject_directs;
    }
  }
  autonomous-system 101;

```

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

### Configuring Device PE1

**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 PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-1/2/0 unit 2 description to-CE1
user@PE1# set ge-1/2/0 unit 2 family inet address 10.0.0.2/30

```

```

user@PE1# set ge-1/2/1 unit 6 description to-P1
user@PE1# set ge-1/2/1 unit 6 family inet address 10.0.0.6/30

```

```

user@PE1# set ge-1/2/2 unit 9 description to-P3
user@PE1# set ge-1/2/2 unit 9 family inet address 10.0.0.9/30

```

```

user@PE1# set lo0 unit 93 family inet address 10.255.165.93/32 primary

```

2. Configure secondary addresses, 1.1.1.30 and 2.2.2.30.

A specific protocol next-hop IP address is required for each topology on each router injecting IBGP routes. You can configure multiple secondary loopback IP addresses on a router to be used as protocol next-hop addresses. This configuration shows nonprimary IP addresses 1.1.1.30/32 and 2.2.2.30/32 configured on loopback interface lo0 for use in the red and blue topologies, respectively.

A group of BGP routes associated with a routing topology use the same unique protocol next hop. For instance, if you configure a PE router to handle two routing topologies, then you would also configure two unique nonprimary addresses under loopback interface lo0.

```

[edit interfaces]
user@PE1# set lo0 unit 93 family inet address 1.1.1.30/32
user@PE1# set lo0 unit 93 family inet address 2.2.2.30/32

```

3. Associate each nonprimary loopback IP address with a topology for inclusion in the associated topology routing table.

Configure the loopback IP address and topology under an OSPF interface statement. You must specifically disable all other topologies known to OSPF for two reasons. First, the loopback address specific to a topology must reside in only one topology routing table. Second, once the topology is added to OSPF, the topology defaults to being enabled on all subsequent interfaces under OSPF.

The Device PE1 configuration places the loopback address 1.1.1.30/32 into the OSPF database as a stub route under this router's OSPF Router-LSA. It belongs to the red and default topologies, but not to the blue topology. The loopback address 1.1.1.30/32 is installed in the remote core routers' topology routing tables **inet.0** and **:red.inet.0**, (but not in **:blue.inet.0**). Use a similar configuration for the blue loopback address 2.2.2.30/32.

```
[edit protocols ospf]
user@PE1# set topology red topology-id 126
user@PE1# set topology blue topology-id 52

[edit protocols ospf area 0.0.0.0]
user@PE1# set interface 1.1.1.30 topology red
user@PE1# set interface 1.1.1.30 topology blue disable

user@PE1# set interface 2.2.2.30 topology blue
user@PE1# set interface 2.2.2.30 topology red disable
```

4. Enable OSPF on the interfaces, and configure specific OSPF link metrics on topologies to identify paths and build trees to different servers.

Links can support all routing topologies to provide backup should a primary multicast path fail.

When a multicast tree gets built through PIM join messages directed toward the source, it follows the most preferred path. A multicast tree to a different multicast source (in a different routing topology) can create another tree along a different path.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-1/2/1.6 metric 10
user@PE1# set interface ge-1/2/1.6 topology blue metric 1
user@PE1# set interface ge-1/2/1.6 topology red

user@PE1# set interface ge-1/2/2.9 metric 10
user@PE1# set interface ge-1/2/2.9 topology red metric 1
user@PE1# set interface ge-1/2/2.9 topology blue

user@PE1# set interface lo0.93 passive
```

5. Create the multicast routing table **inet.2**, and configure PIM to use the **inet.2** routing table.

Set up a separate routing table for multicast lookups. It is populated with routes from **inet.2**. The **inet.2** routing table is populated by routes of type multicast.

```
[edit routing-options]
user@PE1# set rib-groups mcast-rib import-rib inet.2
```

6. Configure PIM to use the routes in **inet.2**.

```
[edit protocols pim]
user@PE1# set rib-group inet mcast-rib
```

7. Enable PIM on the interfaces.

```
[edit protocols pim]
user@PE1# set interface ge-1/2/0.2 mode sparse
user@PE1# set interface ge-1/2/1.6 mode sparse
user@PE1# set interface ge-1/2/2.9 mode sparse
```

8. Configure the router to perform route resolution on protocol next hops using specified routing tables.

The protocol next hop is used to determine the forwarding next-hop interface out of which to forward PIM join messages. This configuration directs **inet.2** route resolution to use topology routing tables **:red.inet.0** and **:blue.inet.0** for protocol next-hop IP address lookups.

You can specify up to two routing tables in the resolution configuration. A key element to this solution is that the protocol next-hop address resides in only one topology routing table. That is, the protocol next hop belongs to a remote PE secondary loopback address and is injected into only one topology routing table. The route resolution scheme first checks routing table **:red.inet.0** for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks routing table **:blue.inet.0**. Hence, only one topology routing table is used for each protocol nexthop address.

```
[edit routing-options resolution rib inet.2]
user@PE1# set resolution-ribs :red.inet.0
user@PE1# set resolution-ribs :blue.inet.0
```

9. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set autonomous-system 100
```

10. Configure BGP.

```
[edit protocols bgp group ibgp]
user@PE1# set type internal
user@PE1# set local-address 10.255.165.93
user@PE1# set family inet unicast
user@PE1# set family inet multicast
user@PE1# set neighbor 10.255.165.111
user@PE1# set neighbor 10.255.165.203
user@PE1# set neighbor 10.255.165.113
user@PE1# set neighbor 10.255.165.95
user@PE1# set neighbor 10.255.165.99
```

```
[edit protocols bgp group ebgp]
user@PE1# set type external
user@PE1# set local-address 10.0.0.2
user@PE1# set family inet unicast
user@PE1# set family inet multicast
user@PE1# set peer-as 101
user@PE1# set neighbor 10.0.0.1
```

11. Set the protocol next hop when exporting EBGp routes into IBGP.

Configure the ingress Device PE1 router to set the BGP route's protocol next-hop address when exporting the route into IBGP.

BGP uses an export policy to set the next hop when injecting the EBGp routes into IBGP.

This configuration is an export policy where there are three possibilities of next hops being set. Route 1.1.1.30 is associated with the red topology. Route 2.2.2.30 is associated with the blue topology. For the default next-hop self policy, the primary loopback address 10.255.165.93 on Device PE1 is used.

The **nhs\_test** policy sets the protocol next-hop based on the community in the BGP update.

```
[edit policy-options]
user@PE1# set community blue members target:50:50
user@PE1# set community red members target:40:40
```

```
[edit policy-options policy-statement nhs_test term a]
user@PE1# set from protocol bgp
user@PE1# set from community red
user@PE1# set then next-hop 1.1.1.30
user@PE1# set then next policy
user@PE1# set then accept
```

```
[edit policy-options policy-statement nhs_test term b]
user@PE1# set from protocol bgp
user@PE1# set from community blue
user@PE1# set then next-hop 2.2.2.30
user@PE1# set then next policy
user@PE1# set then accept
user@PE1# set policy-options policy-statement nhs_test term c then next-hop self
```

```
[edit policy-options policy-statement nhs_inet0_self term a]
user@PE1# set from protocol bgp
user@PE1# set from rib inet.0
user@PE1# set then next-hop self
```

12. Apply the next-hop self policies to the IBGP sessions.

```
[edit protocols bgp group ibgp]
user@PE1# set export nhs_test
user@PE1# set export nhs_inet0_self
```

13. Configure the voice and video topologies, which enable you to use these topologies with OSPF and BGP.

The names **voice** and **video** are local to the router. The names are not propagated beyond this router. However, for management purposes, a consistent naming scheme across routers in a multitopology environment is convenient.

```
[edit routing-options topologies family inet]
user@PE1# set topology red
user@PE1# set topology blue
```

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

```

user@PE1# show interfaces
ge-1/2/0 {
  unit 2 {
    description to-CE1;
    family inet {
      address 10.0.0.2/30;
    }
  }
}
ge-1/2/1 {
  unit 6 {
    description to-P1;
    family inet {
      address 10.0.0.6/30;
    }
  }
}
ge-1/2/2 {
  unit 9 {
    description to-P3;
    family inet {
      address 10.0.0.9/30;
    }
  }
}
lo0 {
  unit 93 {
    family inet {
      address 10.255.165.93/32 {
        primary;
      }
      address 1.1.1.30/32;
      address 2.2.2.30/32;
    }
  }
}

user@PE1# show protocols
bgp {
  group ibgp {
    type internal;
    local-address 10.255.165.93;
    family inet {
      unicast;
      multicast;
    }
    export [ nhs_test nhs_inet0_self ];
    neighbor 10.255.165.111;
    neighbor 10.255.165.203;
    neighbor 10.255.165.113;
    neighbor 10.255.165.95;
  }
}

```

```
        neighbor 10.255.165.99;
    }
    group ebgp {
        type external;
        local-address 10.0.0.2;
        family inet {
            unicast;
            multicast;
        }
        peer-as 101;
        neighbor 10.0.0.1;
    }
}
ospf {
    topology red topology-id 126;
    topology blue topology-id 52;
    area 0.0.0.0 {
        interface ge-1/2/1.6 {
            metric 10;
            topology blue metric 1;
            topology red;
        }
        interface ge-1/2/2.9 {
            metric 10;
            topology red metric 1;
            topology blue;
        }
        interface lo0.93 {
            passive;
        }
        interface 1.1.1.30 {
            topology red;
            topology blue disable;
        }
        interface 2.2.2.30 {
            topology blue;
            topology red disable;
        }
    }
}
pim {
    rib-group inet mcast-rib;
    interface ge-1/2/0.2 {
        mode sparse;
    }
    interface ge-1/2/1.6 {
        mode sparse;
    }
    interface ge-1/2/2.9 {
        mode sparse;
    }
}
user@PE1# show policy-options
policy-statement nhs_inet0_self {
    term a {
```

```

        from {
            protocol bgp;
            rib inet.0;
        }
        then {
            next-hop self;
        }
    }
}
policy-statement nhs_test {
    term a {
        from {
            protocol bgp;
            community red;
        }
        then {
            next-hop 1.1.1.30;
            next policy;
            accept;
        }
    }
    term b {
        from {
            protocol bgp;
            community blue;
        }
        then {
            next-hop 2.2.2.30;
            next policy;
            accept;
        }
    }
    term c {
        then {
            next-hop self;
        }
    }
}
community blue members target:50:50;
community red members target:40:40;

user@PE1# show routing-options
rib-groups {
    mcast-rib {
        import-rib inet.2;
    }
}
autonomous-system 100;
resolution {
    rib inet.2 {
        resolution-ribs [ :red.inet.0 :blue.inet.0 ];
    }
}
topologies {
    family inet {
        topology red;
    }
}

```

```
        topology blue;
    }
}
```

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

## Verification

---

Confirm that the configuration is working properly.

- [Checking the IBGP routes in inet.2 on page 60](#)
- [Verifying the Routes on page 61](#)
- [Checking the Resolving BGP Next Hops on page 62](#)
- [Examining the Protocol Next Hop on page 63](#)
- [Verifying the OSPF Neighbor on page 64](#)
- [Checking the Router LSA on page 64](#)
- [Checking How Traffic Traverses the Network on page 65](#)

### *Checking the IBGP routes in inet.2*

**Purpose** Make sure that the routes injected into IBGP by Device PE1 have next hops that are based on the topology to which they belong.

**Action** From operational mode, enter the **show route table extensive** command.

```
user@PE1> show route 11.19.130.0/24 table inet.2 extensive
inet.2: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 93e9768
  Flags: Nexthop Change
  Nexthop: 1.1.1.30
  Localpref: 100
  AS path: [100] 101 I
  Communities: target:40:40
Path 11.19.130.0 from 10.0.0.1 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 1180
    Address: 0x94003ec
    Next-hop reference count: 16
    Source: 10.0.0.1
    Next hop: 10.0.0.1 via lt-1/2/0.2, selected
    Session Id: 0x380004
    State: <Active Ext>
    Local AS: 100 Peer AS: 101
    Age: 22
    Validation State: unverified
    Task: BGP_101.10.0.0.1+58346
    Announcement bits (1): 0-BGP_RT_Background
    AS path: 101 I
    Communities: target:40:40
    Accepted
    Localpref: 100
    Router ID: 10.255.165.97
```

**Meaning** This output shows an IBGP route in the **inet.2** routing table, as seen from Device PE1. The route was originally injected into IBGP by Device PE1, where the next hop was set based on the topology to which the route belonged. The BGP community value determined the topology association.

The route 11.19.130/24 belongs to the red topology because it has a community value of target:40:40. The protocol next hop is 1.1.1.30, and the forwarding next hop is ge-1/2/1.42.

### *Verifying the Routes*

**Purpose** Make sure that the routes are in the expected routing tables and that the expected communities are attached to the routes.

**Action** From operational mode, enter the **show route detail** command on Device PE1.

```
user@PE1> show route 11.19.130.0/24 detail
```

```
inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Router, Next hop index: 812
            Address: 0xb9f064c
            Next-hop reference count: 22
            Source: 10.0.0.1
            Next hop: 10.0.0.1 via fe-1/2/0.2, selected
            Session Id: 0x600004
            State: <Active Ext>
            Local AS: 100 Peer AS: 101
            Age: 3d 21:44:07
            Task: BGP_101.10.0.0.1+51873
            Announcement bits (3): 0-KRT 3-BGP_RT_Background 4-Resolve tree
3
            AS path: 101 I
            Communities: target:40:40
            Accepted
            Localpref: 100
            Router ID: 10.255.165.97
            Secondary Tables: :voice.inet.0

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Router, Next hop index: 812
            Address: 0xb9f064c
            Next-hop reference count: 22
            Source: 10.0.0.1
            Next hop: 10.0.0.1 via fe-1/2/0.2, selected
            Session Id: 0x600004
            State: <Secondary Active IndepResolution Ext>
            Local AS: 100 Peer AS: 101
            Age: 3d 21:44:07
            Task: BGP_101.10.0.0.1+51873
            Announcement bits (2): 0-KRT 1-Resolve tree 1
            AS path: 101 I
            Communities: target:40:40
            Accepted
            Localpref: 100
```

```
Router ID: 10.255.165.97
Primary Routing Table inet.0
```

**Meaning** This output shows BGP route 11.19.130.0/24 with community value target:40:40. Because the route matches the criteria for the voice topology, it is added to both the default and voice topology routing tables (**inet.0** and **:voice.inet.0**). Device PE1 learns the route from Device CE1 through EBGp and then injects the route into IBGP.

### *Checking the Resolving BGP Next Hops*

**Purpose** Check the protocol next hop and forwarding next hop.

**Action** From operational mode, enter the **show route detail** command on Device PE2.

```
user@PE2> show route 11.19.130.0/24 detail
inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Indirect
            Address: 0xb9f0e04
            Next-hop reference count: 12
            Source: 10.255.165.93
            Next hop type: Router, Next hop index: 262153
            Next hop: 10.0.0.37 via fe-1/2/0.38
            Session Id: 0x700004
            Next hop: 10.0.0.41 via fe-1/2/1.42, selected
            Session Id: 0x700005
            Protocol next hop: 10.255.165.93
            Indirect next hop: bb8c000 262154 INH Session ID: 0x700007
            State: <Active Int Ext>
            Local AS: 100 Peer AS: 100
            Age: 3d 4:27:40 Metric2: 30
            Task: BGP_100.10.255.165.93+179
            Announcement bits (3): 0-KRT 3-BGP_RT_Background 4-Resolve tree
3
      AS path: 101 I
      Communities: target:40:40
      Accepted
      Localpref: 100
      Router ID: 10.255.165.93
      Secondary Tables: :voice.inet.0

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
11.19.130.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Indirect
            Address: 0xb9f0f34
            Next-hop reference count: 6
            Source: 10.255.165.93
            Next hop type: Router, Next hop index: 1188
            Next hop: 10.0.0.37 via fe-1/2/0.38, selected
            Session Id: 0x700004
            Protocol next hop: 10.255.165.93
            Indirect next hop: bb8c1d8 262177 INH Session ID: 0x700007
            State: <Secondary Active IndepResolution Int Ext>
            Local AS: 100 Peer AS: 100
            Age: 3d 2:00:20 Metric2: 30
            Task: BGP_100.10.255.165.93+179
```

```

Announcement bits (2): 0-KRT 1-Resolve tree 1
AS path: 101 I
Communities: target:40:40
Accepted
Localpref: 100
Router ID: 10.255.165.93
Primary Routing Table inet.0

```

**Meaning** A typical IBGP core has BGP routes with protocol next hops that resolve using the underlying IGP routes. IBGP routes in a topology routing table have protocol next-hop IP addresses. By default, the same topology routing table is used to look up and resolve the protocol next-hop IP address to a forwarding next hop. This output from Device PE2 shows the same BGP route as seen in the previous example: 11.19.130.0/24. The route is being shown from a different perspective, that is, from Device PE2 as an IBGP route. Similarly, this IBGP route is added to both **inet.0** and **:voice.inet.0** on Device PE2. However, while each route has the same protocol next hop, each route has a different forwarding next hop (ge-0/0/3.0 instead of ge-0/1/4.0). The reason for this difference is when the protocol next-hop IP address 10.255.165.93 is resolved, it uses the corresponding routing table (**inet.0** or **:voice.inet.0**) to look up the protocol next hop.

#### *Examining the Protocol Next Hop*

**Purpose** Check the protocol next hop and forwarding next hop.

**Action** From operational mode, enter the **show route** command on Device PE2.

```

user@PE2> show route 10.255.165.93
inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 04:37:26, metric 30
                   > to 10.0.0.37 via fe-1/2/0.38
                   to 10.0.0.41 via fe-1/2/1.42

:voice.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 02:10:04, metric 30
                   > to 10.0.0.37 via fe-1/2/0.38

:video.inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.165.93/32    *[OSPF/10] 3d 02:03:16, metric 30
                   > to 10.0.0.41 via fe-1/2/1.42

```

**Meaning** This output from Device PE2 shows the protocol next hop of 11.19.130.0/24, which is IP address 10.255.165.93, thus further demonstrating how IBGP route 11.19.130.0/24 resolves its protocol next hop. The forwarding next hops of 10.255.165.93 match the IBGP forwarding next hops of route 11.19.130.0/24 as shown in the previous example. Observe here that the IP address 10.255.165.93 is also in routing table **:video.inet.0**. This address is the loopback address of Device PE1, and as such, resides in all three routing tables. This example also shows how traffic entering Device PE2 destined to 11.19.130.0/24 exits different interfaces depending on its associated topology. The actual traffic is marked

in such a way that a firewall filter can direct the traffic to use a particular topology routing table.

### *Verifying the OSPF Neighbor*

**Purpose** Make sure that the expected topologies are enabled on the OSPF neighbor.

**Action** From operational mode, enter the **show (ospf | ospf3) neighbor extensive** command on Device P2.

```
user@P2> show ospf neighbor 10.0.0.21 extensive
Address          Interface          State    ID                Pri  Dead
10.0.0.21        fe-1/2/0.22        Full    10.255.165.111   128   39
Area 0.0.0.0, opt 0x52, DR 10.0.0.22, BDR 10.0.0.21
Up 3d 06:09:50, adjacent 3d 06:09:50
Topology default (ID 0) -> Bidirectional
Topology video (ID 52) -> Bidirectional
```

**Meaning** This Device P2 output shows OSPF neighbor PE2 (10.0.0.21), where multitopology OSPF default and video are participants. The **Bidirectional** flag shows that the neighbor is configured using the same multitopology OSPF ID.

### *Checking the Router LSA*

**Purpose** Check the links where video and voice topologies are enabled.

**Action** From operational mode, enter the **show ospf database extensive** command on Device P2.

```
user@P2> show ospf database lsa-id 10.255.165.203 extensive

OSPF database, Area 0.0.0.0
Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
Router    10.255.165.203    10.255.165.203  0x80000063  1552 0x22 0xdff3  80
bits 0x0, link count 3
id 10.255.165.203, data 255.255.255.255, Type Stub (3)
Topology count: 2, Default metric: 0
Topology video (ID 52) -> Metric: 0
Topology voice (ID 126) -> Metric: 0
id 10.0.0.38, data 10.0.0.38, Type Transit (2)
Topology count: 2, Default metric: 10
Topology video (ID 52) -> Metric: 200
Topology voice (ID 126) -> Metric: 10
id 10.0.0.42, data 10.0.0.42, Type Transit (2)
Topology count: 1, Default metric: 10
Topology video (ID 52) -> Metric: 10
Topology default (ID 0)
Type: Transit, Node ID: 10.0.0.42
Metric: 10, Bidirectional
Type: Transit, Node ID: 10.0.0.38
Metric: 10, Bidirectional
Topology video (ID 52)
Type: Transit, Node ID: 10.0.0.42
Metric: 10, Bidirectional
Type: Transit, Node ID: 10.0.0.38
Metric: 200, Bidirectional
Topology voice (ID 126)
```

```

Type: Transit, Node ID: 10.0.0.38
Metric: 10, Bidirectional
Aging timer 00:34:08
Installed 00:25:49 ago, expires in 00:34:08, sent 00:25:47 ago
Last changed 3d 01:45:51 ago, Change count: 10

```

**Meaning** This Device P2 output shows the Router-LSA originated by Device PE2. The LSA shows links where video and voice topologies are enabled (in addition to the default topology).

### *Checking How Traffic Traverses the Network*

**Purpose** Make sure that the expected paths are used.

**Action** From operational mode, enter the **traceroute** command on Device CE1.

The first example output shows that a traceroute over the voice topology goes from Device CE1 to Device CE2 where DSCPs are set. The routes are resolved over **:voice.inet.0**. This traceroute path follows the voice path CE1-PE1-P1-P2-PE2-CE2.

```

user@CE1> traceroute 11.19.140.1 source 11.19.130.1 tos 160
traceroute to 11.19.140.1 (11.19.140.1) from 11.19.130.1, 30 hops max, 40 byte
packets
 1 10.0.0.2 (10.0.0.2) 2.015 ms 1.924 ms 1.770 ms
 2 10.0.0.5 (10.0.0.5) 1.890 ms 1.010 ms 0.974 ms
 3 10.0.0.34 (10.0.0.34) 0.986 ms 1.031 ms 0.973 ms
 4 10.0.0.38 (10.0.0.38) 1.213 ms 1.065 ms 1.154 ms
 5 11.19.140.1 (11.19.140.1) 1.696 ms 4.286 ms 1.332 ms

```

This output shows a traceroute from Device CE1 to Device CE2 for voice where no DSCPs are set. The routes are resolved over **inet.0**, and the resulting path is different from the previous case where the DSCPs are set. This traceroute path follows the default path CE1-PE1-P4-PE2-CE2.

```

user@CE1> traceroute 11.19.140.1 source 11.19.130.1

traceroute to 11.19.140.1 (11.19.140.1) from 11.19.130.1, 30 hops max, 40 byte
packets
 1 10.0.0.2 (10.0.0.2) 1.654 ms 1.710 ms 1.703 ms
 2 10.0.0.5 (10.0.0.5) 1.790 ms 1.045 ms 0.975 ms
 3 10.0.0.18 (10.0.0.18) 0.989 ms 1.041 ms 0.983 ms
 4 10.0.0.42 (10.0.0.42) 0.994 ms 1.036 ms 1.002 ms
 5 11.19.140.1 (11.19.140.1) 1.329 ms 2.248 ms 2.225 ms

```

This output shows a traceroute from Device CE1 to Device CE2 for video traffic where the firewall filter is based on the destination address. The routes are resolved over **:video.inet.0**. This traceroute follows the video path CE1-PE1-P3-P4-PE2-CE2.

```

user@CE1> traceroute 11.19.142.1 source 11.19.132.1

traceroute to 11.19.142.1 (11.19.142.1) from 11.19.132.1, 30 hops max, 40 byte
packets
 1 10.0.0.2 (10.0.0.2) 1.126 ms 1.300 ms 0.995 ms
 2 10.0.0.10 (10.0.0.10) 0.981 ms 1.018 ms 0.991 ms
 3 10.0.0.30 (10.0.0.30) 0.997 ms 1.886 ms 1.952 ms
 4 10.0.0.42 (10.0.0.42) 1.800 ms 1.038 ms 0.980 ms
 5 11.19.142.1 (11.19.142.1) 1.367 ms 1.352 ms 1.328 ms

```

This output shows a traceroute from Device CE1 to Device CE2 for video where DSCPs are set. The DSCP bits are directing Device PE1 to use the topology table **:voice.inet.0**. Because there is no entry in the voice routing table for the video routes, traffic is dropped.

```
user@CE1> traceroute 11.19.142.1 source 11.19.132.1 tos 160
```

```
traceroute to 11.19.142.1 (11.19.142.1) from 11.19.132.1, 30 hops max, 40 byte packets
```

```
1  10.0.0.2 (10.0.0.2)  1.135 ms !N  1.007 ms !N  0.954 ms !N
```

#### Related Documentation

- [Example: Configuring Multitopology Routing Based on Applications on page 13](#)

## CHAPTER 4

# Multitopology Routing Configuration Statements

- [community \(Protocols BGP\) on page 68](#)
- [rib \(Multitopology Routing\) on page 69](#)
- [topologies \(Multitopology Routing\) on page 70](#)
- [topology \(Filter-Based Forwarding\) on page 71](#)
- [topology \(Multitopology Routing\) on page 72](#)
- [topology \(OSPF\) on page 73](#)
- [topology \(OSPF Interface\) on page 74](#)
- [topology-id on page 75](#)

## community (Protocols BGP)

<b>Syntax</b>	community { target <i>identifier</i> ; }
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit protocols bgp family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet   inet6) unicast topology <i>name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet   inet6) topology <i>name</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	Configure the community to identify the multitopology routes. BGP uses the target community identifier to install the routes it learns in the appropriate multitopology routing tables.
<b>Options</b>	<b>target <i>identifier</i></b> —Configure the destination to which the route is going.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li> <li>• <a href="#">Example: Configuring Origin Validation for BGP</a></li> </ul>

## rib (Multitopology Routing)


<b>Syntax</b>	<pre> rib <i>routing-table-name</i> {     static {         route <i>destination-prefix</i> {             <i>next-hop</i>;         }         <i>static-options</i>;     } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
<b>Release Information</b>	Statement support for multitopology routing introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a static route to install routes in the routing table for a specific topology.
<b>Options</b>	<p><b><i>routing-table-name</i></b>—Name of the routing table for a topology. Use the following format: <b><i>logical-system-name/routing-instance-name:topology-name.protocol.identifier</i></b>. Include the routing instance string only if the instance is not the master. The logical system string is included only if the logical system identifier has a value other than 0 (zero). Each routing table for a topology includes a colon (:) before the topology name. <b><i>protocol</i></b> is the protocol family, which can be <b>inet</b> or <b>inet6</b>. <b><i>identifier</i></b> is the positive integer that specifies the instance of the routing table. For example, to install IPv6 routes to the routing table for a topology named voice in the master instance, include <b>:voice.inet6.0</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Examples: Configuring Multiprotocol BGP</i></li> <li>• <i>Example: Configuring Martian Addresses</i></li> <li>• <i>static</i></li> </ul>

## topologies (Multitopology Routing)

---

<b>Syntax</b>	<pre>topologies {     family (inet   inet6) {         <b>topology</b> <i>topology-name</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Configure a topology for multitopology routing. Each topology creates a new routing table that is populated with direct routes from the topology.
<b>Options</b>	<b>family</b> —Configure the type of family address type.  <b>inet</b> —IPv4  <b>inet6</b> —IPv6  The remaining statement is explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li><li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source on page 39</a></li></ul>

## topology (Filter-Based Forwarding)

<b>Syntax</b>	<code>topology <i>topology-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit firewall family (inet   inet6) filter <i>filter-name</i> term <i>term-name</i> then],</p> <p>[edit firewall family (inet   inet6) filter <i>filter-name</i> term <i>term-name</i> then logical-system <i>logical-system-name</i>],</p> <p>[edit firewall family (inet   inet6) filter <i>filter-name</i> term <i>term-name</i> then logical-system <i>logical-system-name</i> routing-instance <i>routing-instance-name</i>],</p> <p>[edit firewall family (inet   inet6) filter <i>filter-name</i> term <i>term-name</i> then routing-instance <i>routing-instance-name</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	<p>Configure a topology for filter-based forwarding for multitopology routing. The firewall filter you apply to the ingress interface is used to look up traffic against the configured topology, and, if a route matches the conditions you configure for the term, the route is accepted and added to the routing table for the specific topology.</p> <p>There are multiple ways to configure a topology for filter-based forwarding, depending on the type of instance or logical system you want to specify for the forwarding class.</p>
<div>  <b>NOTE:</b> The options for logical system and routing instance precede the <b>topology</b> statement with the <b>then</b> statement. </div>	
<b>Options</b>	<i>topology-name</i> —Name of a topology against which you want to match traffic.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li> <li>• <i>Routing Policy Feature Guide for Routing Devices</i></li> </ul>

## topology (Multitopology Routing)

---

<b>Syntax</b>	<code>topology <i>topology-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>topologies</b> family (inet   inet6)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options <b>topologies</b> family (inet   inet6)],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options <b>topologies</b> family (inet   inet6)],</code> <code>[edit routing-options <b>topologies</b> family (inet   inet6)]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Configure the name of a topology configured to run multitopology routing.
<b>Options</b>	<b><i>topology-name</i></b> —Name of the topology. Include a string value that describes the type of traffic, such as voice or video. For IPv4 multicast traffic, include <b>ipv4-multicast</b> as the name.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li><li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source on page 39</a></li></ul>

## topology (OSPF)

<b>Syntax</b>	<pre> topology (default   ipv4-multicast   <i>name</i>) {     spf-options {         delay <i>milliseconds</i>;         holddown <i>milliseconds</i>;         rapid-runs <i>number</i>;     }     topology-id <i>number</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf],  [edit protocols ospf],  [edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.  Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Enable a topology for OSPF multitopology routing. You must first configure one or more topologies under the <b>[edit routing-options]</b> hierarchy level.</p>
<b>Options</b>	<p><b>default</b>—Name of the default topology. This topology is automatically created, and all routes that correspond to it are automatically added to the <b>inet.0</b> routing table. You can modify certain default parameters, such as for the SPF algorithm.</p> <p><b>ipv4-multicast</b>—Name of the topology for IPv4 multicast traffic.</p> <p><b><i>name</i></b>—Name of a topology you configured at the <b>[edit routing-options]</b> hierarchy level to create a topology for a specific type of traffic, such as voice or video.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source on page 39</a></li> </ul>

## topology (OSPF Interface)

---

<b>Syntax</b>	<code>topology (ipv4-multicast   <i>name</i>) {     metric <i>metric</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> ], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	<p>Configure interface-specific properties for multitopology OSPF, including topology-specific metric values for an interface.</p> <p>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 over those with higher path metrics. The default value for the OSPF metric for an interface is 1. You can modify the default value for an OSPF interface and configure a topology-specific metric for that interface. The topology-specific metric applies to routes advertised from the interface that belong only to that topology.</p>
<b>Default</b>	The default value of the topology metric is the same as the default metric value calculated by OSPF or the value configured for the OSPF metric.
<b>Options</b>	<p><b>ipv4-multicast</b>—Name of the topology for IPv4 multicast traffic.</p> <p><b><i>name</i></b>—Name of a topology created under the <b>[edit routing-options]</b> hierarchy level.</p> <p><b>metric <i>metric</i></b>—Cost of a route from an OSPF interface. You can specify a metric value for a topology that is different from the value specified for the interface.</p> <p><b>Range:</b> 1 through 65,535</p> <p><b>Default:</b> 1</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li><li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source on page 39</a></li></ul>

## topology-id

---

<b>Syntax</b>	<code>topology-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">topology name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology name</a>],</p> <p>[edit protocols ospf <a href="#">topology name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology name</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	Configure a topology identifier for a topology enabled for OSPF.
<b>Default</b>	The default identifier for the default topology is 0, and the default identifier for the topology for IPv4 multicast traffic is 1. These identifiers are predefined and cannot be modified.
<b>Options</b>	<p><b><i>number</i></b>—The integer value used to identify the topology.</p> <p><b>Range:</b> 32 through 127</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications on page 13</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source on page 39</a></li> <li>• <a href="#">topology on page 73</a></li> </ul>



## PART 3

# Administration

- [Multitopology Routing Operational Commands on page 79](#)



## CHAPTER 5

# Multitopology Routing Operational Commands

- `show (ospf | ospf3) interface`

## show (ospf | ospf3) interface

---

<b>Syntax</b>	<code>show (ospf   ospf3) interface</code> <code>&lt;brief   detail   extensive&gt;</code> <code>&lt;area <i>area-id</i>&gt;</code> <code>&lt;interface-name&gt;</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</code>
<b>Syntax (EX Series Switches and QFX Series)</b>	<code>show (ospf   ospf3) interface</code> <code>&lt;brief   detail   extensive&gt;</code> <code>&lt;area <i>area-id</i>&gt;</code> <code>&lt;interface-name&gt;</code> <code>&lt;instance <i>instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>area</b> option introduced in Junos OS Release 9.2. <b>area</b> option introduced in Junos OS Release 9.2 for EX Series switches. <b>realm</b> option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display the status of OSPF interfaces.
<b>Options</b>	<b>none</b> —Display standard information about the status of all OSPF interfaces for all routing instances  <b>brief   detail   extensive</b> —(Optional) Display the specified level of output.  <b>area <i>area-id</i></b> —(Optional) Display information about the interfaces that belong to the specified area.  <b><i>interface-name</i></b> —(Optional) Display information for the specified interface.  <b>instance <i>instance-name</i></b> —(Optional) Display all OSPF interfaces under the named routing instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b> —(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ospf interface brief on page 83</a> <a href="#">show ospf interface detail on page 83</a> <a href="#">show ospf3 interface detail on page 83</a>

[show ospf interface detail \(When Multiarea Adjacency Is Configured\) on page 83](#)  
[show ospf interface area area-id on page 84](#)  
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 85](#)  
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 85](#)

**Output Fields** Table 4 on page 81 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

**Table 4: show (ospf | ospf3) interface Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
<b>State</b>	State of the interface: <b>BDR</b> , <b>Down</b> , <b>DR</b> , <b>DROther</b> , <b>Loop</b> , <b>PtToPt</b> , or <b>Waiting</b> .	All levels
<b>Area</b>	Number of the area that the interface is in.	All levels
<b>DR ID</b>	Address of the area's designated router.	All levels
<b>BDR ID</b>	Backup designated router for a particular subnet.	All levels
<b>Nbrs</b>	Number of neighbors on this interface.	All levels
<b>Type</b>	Type of interface: <b>LAN</b> , <b>NBMA</b> , <b>P2MP</b> , <b>P2P</b> , or <b>Virtual</b> .	<b>detail extensive</b>
<b>Address</b>	IP address of the neighbor.	<b>detail extensive</b>
<b>Mask</b>	Netmask of the neighbor.	<b>detail extensive</b>
<b>Prefix-length</b>	(OSPFv3) IPv6 prefix length, in bits.	<b>detail extensive</b>
<b>OSPF3-Intf-Index</b>	(OSPFv3) OSPF version 3 interface index.	<b>detail extensive</b>
<b>MTU</b>	Interface maximum transmission unit (MTU).	<b>detail extensive</b>
<b>Cost</b>	Interface cost (metric).	<b>detail extensive</b>
<b>DR addr</b>	Address of the designated router.	<b>detail extensive</b>
<b>BDR addr</b>	Address of the backup designated router.	<b>detail extensive</b>
<b>Adj count</b>	Number of adjacent neighbors.	<b>detail extensive</b>
<b>Secondary</b>	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	<b>detail extensive</b>
<b>Flood Reduction</b>	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the <b>DoNotAge</b> bit set. As a result, LSAs are refreshed only when a change occurs.	<b>extensive</b>

Table 4: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Router priority used in designated router (DR) election on this interface.	detail extensive
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	extensive
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	extensive
Descriptor list	List of packet descriptors.	extensive
Hello	Configured value for the hello timer.	detail extensive
Dead	Configured value for the dead timer.	detail extensive
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> <li>• <b>MD5</b>—The MD5 mechanism is configured in accordance with RFC 2328.</li> <li>• <b>None</b>—No authentication method is configured.</li> <li>• <b>Password</b>—A simple password (RFC 2328) is configured.</li> </ul>	detail extensive
Topology	(Multiarea adjacency) Name of topology: <b>default</b> or <i>name</i> .	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: <b>in sync</b> , <b>in holddown</b> , and <b>not supported</b> .	extensive
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	extensive
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer.  If the state is not synchronized, and the hold time is not infinity, the <b>remaining</b> field displays the number of seconds that remain until the configured hold timer expires.	extensive
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	detail extensive
Active key ID	(OSPFv2 and MD5) Number from <b>0</b> to <b>255</b> that uniquely identifies an MD5 key.	detail extensive
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as <b>Start time 1970 Jan 01 00:00:00 PST</b> .	detail extensive
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

## Sample Output

### show ospf interface brief

```
user@host> show ospf interface brief
```

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

### show ospf interface detail

```
user@host> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-0/0/1.0	BDR	0.0.0.0	192.168.37.12	10.255.245.215	1

Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40  
 DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128  
 Hello 10, Dead 40, ReXmit 5, Not Stub

Interface	State	Area	DR ID	BDR ID	Nbrs
tl-0/2/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	0

Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604  
 Adj count 0  
 Hello 10, Dead 40, ReXmit 5, Not Stub  
 Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST  
 IPsec SA Name: sa

### show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
```

Interface	State	Area	DR-ID	BDR-ID	Nbrs
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64  
 OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1  
 Hello 10, Dead 40, ReXmit 5, Not Stub

### show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
```

```
regress@router> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0  
 DR addr: 127.0.0.1, Adj count: 0, Priority: 128  
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub  
 Auth type: None  
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0  
 DR addr: 10.255.245.2, Adj count: 0, Priority: 128  
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub  
 Auth type: None  
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	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: 4470, Cost: 1

```

Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      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: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

#### show ospf interface area area-id

```

user@host> show ospf interface area 1.1.1.1
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1

```

**show ospf interface extensive**  
**(When Flooding Reduction Is Enabled)**

```
user@host> show ospf interface extensive
Interface          State   Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0         PtToPt 0.0.0.0    0.0.0.0    0.0.0.0      0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

**show ospf interface extensive**  
**(When LDP Synchronization Is Configured)**

```
user@host> show ospf interface extensive
Interface          State   Area      DR ID      BDR ID
Nbrs
so-1/0/3.0         Down    0.0.0.0    0.0.0.0    0.0.0.0
0

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddown: 10 seconds, remaining: 1
```



## PART 4

# Index

- [Index on page 89](#)



# Index

## Symbols

#, comments in configuration statements.....	xii
( ), in syntax descriptions.....	xii
< >, in syntax descriptions.....	xii
[ ], in configuration statements.....	xii
{ }, in configuration statements.....	xii
(pipe), in syntax descriptions.....	xii

## B

BGP See secure interdomain routing	
origin AS validation.....	15, 41
secure interdomain routing.....	15
BGP in multitopology routing.....	14
braces, in configuration statements.....	xii
brackets	
angle, in syntax descriptions.....	xii
square, in configuration statements.....	xii

## C

comments, in configuration statements.....	xii
community statement	
Multitopology Routing.....	68
conventions	
text and syntax.....	xi
curly braces, in configuration statements.....	xii
customer support.....	xiii
contacting JTAC.....	xiii

## D

documentation	
comments on.....	xiii

## F

font conventions.....	xi
-----------------------	----

## M

manuals	
comments on.....	xiii
Multitopology Routing	
community statement.....	68
overview.....	13

## O

origin AS validation	
BGP.....	15, 41
OSPF	
interfaces, displaying.....	80
OSPF in multitopology routing.....	13

## P

parentheses, in syntax descriptions.....	xii
--	-----

## R

rib statement	
multitopology routing	
static routes.....	69

## S

secure interdomain routing.....	15, 41 See BGP
show (ospf   ospf3) interface command.....	80
support, technical See technical support	
syntax conventions.....	xi

## T

technical support	
contacting JTAC.....	xiii
topologies statement	
multitopology routing.....	70
topology statement	
filter-based forwarding	
multitopology routing.....	71
multitopology routing.....	72
OSPF.....	73
OSPF interface.....	74
topology-id statement	
multitopology routing.....	75

## V

validation statement	
usage guidelines.....	15, 41

