Junos® OS

Protocol-Independent Routing Properties Feature Guide for Routing Devices

Release
13.2

Published: 2013-07-23

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS Protocol-Independent Routing Properties Feature Guide for Routing Devices
13.2

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

Table of Contents

List of Figures

List of Tables

About the Documentation

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:

- ACX Series
- J Series
- SRX Series
- T Series
- MX Series
- M Series
- PTX 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.xsl;
       }
     }
   }
   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.xsl; }
   ```

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 xv 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.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the**configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the**[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric *metric*>;**
| (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.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		static {
		route default {
		nexthop *address*;
		retain;
		}
		}
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

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**, 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: https://www.juniper.net/alerts/
- 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

CHAPTER 1

Introduction to Protocol-Independent Routing Properties

Protocol-Independent Routing Properties Overview

In Junos OS, routing capabilities and features that are not specific to any particular routing protocol are collectively called protocol-independent routing properties. These features often interact with routing protocols. In many cases, you combine protocol-independent properties and routing policy to achieve a goal. For example, you define a static route using protocol-independent properties, and then, using a routing policy, you can redistribute the static route into a routing protocol, such as BGP, OSPF, or IS-IS.

Protocol-independent routing properties include:

- Static, aggregate, and generated routes
- Bidirectional Forwarding Detection on static routes
- Global preference
- Martian routes
- Routing tables and routing information base (RIB) groups

Related Documentation

- Examples: Configuring Static Routes on page 9
- Examples: Creating a Routing Table and Populating It with Routes on page 38

Routing Table Features in the Junos OS

The Junos OS maintains two databases for routing information:

- Routing table—Contains all the routing information learned by all routing protocols. (Some vendors refer to this kind of table as a routing information base [RIB].)
- Forwarding table—Contains the routes actually used to forward packets. (Some vendors refer to this kind of table as a forwarding information base [FIB].)

By default, the Junos OS maintains three routing tables: one for IP version 4 (IPv4) unicast routes, a second for multicast routes, and a third for MPLS. You can configure additional routing tables.

The Junos OS maintains separate routing tables for IPv4 and IP version 6 (IPv6) routes.

The Junos OS installs all active routes from the routing table into the forwarding table. The active routes are routes that are used to forward packets to their destinations. The Junos operating system kernel maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the component responsible for forwarding packets.

The Junos routing protocol process generally determines the active route by selecting the route with the lowest preference value. The Junos OS provides support for alternate and tiebreaker preferences, and some of the routing protocols, including BGP and MPLS, use these additional preferences.

You can add martian addresses and static, aggregate, and generated routes to the Junos routing tables, configuring the routes with one or more of the properties shown in Table 3 on page 4.

Table 3: Routing Table Route Properties

Description	Static	Aggregate	Generated
Destination address	X	X	X
Default route to the destination	X	X	X
IP address or interface of the next hop to the destination	X	–	–
Label-switched path (LSP) as next hop	X	–	–
Drop the packets, install a reject route for this destination, and send Internet Control Message Protocol (ICMP) unreachable messages	X	X	X
Drop the packets, install a reject route for this destination, but do not send ICMP unreachable messages	X	X	X
Cause packets to be received by the local router	X	–	–
Associate a metric value with the route	X	X	X
Type of route	X	X	X
Preference values	X	X	X
Additional preference values	X	X	X
Independent preference (**qualified-next-hop** statement)	X	–	–

Table 3: Routing Table Route Properties *(continued)*

Description	Static	Aggregate	Generated
BGP community information to associate with the route	X	X	X
Autonomous system (AS) path information to associate with the route	X	X	X
OSPF tag strings to associate with the route	X	X	X
Do not install active static routes into the forwarding table	X	–	–
Install the route into the forwarding table	X	–	–
Permanently retain a static route in the forwarding table	X	–	–
Include only the longest common leading sequences from the contributing AS paths	–	X	–
Include all AS numbers for a specific route	–	X	–
Retain an inactive route in the routing and forwarding tables	X	X	X
Remove an inactive route from the routing and forwarding tables	X	X	X
Active policy to associate with the route	–	X	X
Specify that a route is ineligible for readvertisement	X	–	–
Specify route to a prefix that is not a directly connected next hop	X	–	–

PART 2

Configuration

CHAPTER 2

Configuring Routing Tables and Static Routes

Examples: Configuring Static Routes

Understanding Basic Static Routing

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

You can specify options that define additional information about static routes that is included with the route when it is installed in the routing table. All static options are optional.

Example: Configuring a Basic Set of Static Routes

This example shows how to configure a basic set of static routes.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 1 on page 11 shows the sample network.

Figure 1: Customer Routes Connected to a Service Provider

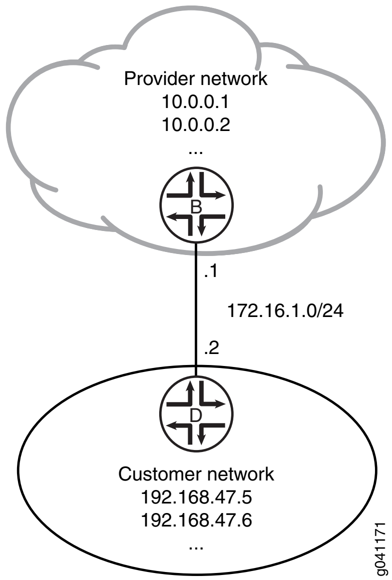

Configuration

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 B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
```

Step-by-Step Procedure

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

To configure basic static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```


2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. If you are done configuring Device B, commit the configuration.

```
[edit interfaces]
user@B# commit
```

4. On Device D, configure the interfaces.

```
[edit]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

5. On Device D, create a static route and set the next-hop address.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

6. If you are done configuring Device D, commit the configuration.

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

Results

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

Device B

```
user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
      address 10.0.0.2/32;
    }
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 next-hop 172.16.1.2;
}
```

Device D

```
user@D# show interfaces
ge-1/2/0 {
```



```
    unit 1 {
        description D->B;
        family inet {
            address 172.16.1.2/24;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.47.5/32;
            address 192.168.47.6/32;
        }
    }
}

user@D# show routing-options
static {
    route 0.0.0.0/0 next-hop 172.16.1.1;
}
```

Verification

Confirm that the configuration is working properly.

- Checking the Routing Tables on page 13
- Pinging the Remote Addresses on page 14

Checking the Routing Tables

Purpose Make sure that the static routes appear in the routing tables of Device B and Device D.

Action

```
user@B> show route
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32        *[Direct/0] 00:29:43
                    > via lo0.57
10.0.0.2/32        *[Direct/0] 00:29:43
                    > via lo0.57
172.16.1.0/24      *[Direct/0] 00:34:40
                    > via ge-1/2/0.0
172.16.1.1/32      *[Local/0] 00:34:40
                      Local via ge-1/2/0.0
192.168.47.0/24    *[Static/5] 00:31:23
                    > to 172.16.1.2 via ge-1/2/0.0

user@D> show route
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:31:24
                    > to 172.16.1.1 via ge-1/2/0.1
172.16.1.0/24      *[Direct/0] 00:35:21
                    > via ge-1/2/0.1
172.16.1.2/32      *[Local/0] 00:35:21
                      Local via ge-1/2/0.1
192.168.47.5/32    *[Direct/0] 00:35:22
                    > via lo0.2
192.168.47.6/32    *[Direct/0] 00:35:21
                    > via lo0.2
```

Meaning The static routes are in the routing tables.

Pinging the Remote Addresses

Purpose Verify that the static routes are working.

From Device B, ping one of the loopback interface addresses on Device D.

From Device D, ping one of the loopback interface addresses on Device B.

Action

```
user@B> ping 192.168.47.5
PING 192.168.47.5 (192.168.47.5): 56 data bytes
64 bytes from 192.168.47.5: icmp_seq=0 ttl=64 time=156.126 ms
64 bytes from 192.168.47.5: icmp_seq=1 ttl=64 time=120.393 ms
64 bytes from 192.168.47.5: icmp_seq=2 ttl=64 time=175.361 ms

user@D> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.315 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=31.819 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.268 ms
```

Example: Configuring IPv6 Static Routes

This example shows how to configure static routes when the interfaces have IPv6 addresses.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure a static default route of ::/0, using a next-hop address 2001:db8:0:1:2a0:a502:0:1da.

For demonstration purposes, some loopback interfaces are configured on Device A and Device E. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 2 on page 16 shows the sample network.

Figure 2: Customer Routes Connected to a Service Provider

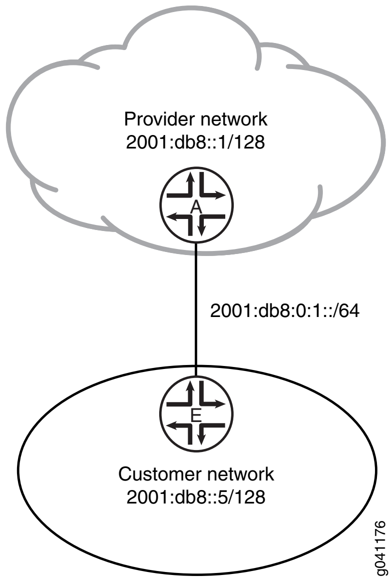

Configuration

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 A

```
set interfaces fe-1/2/0 unit 1 description to-E
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1:2a0:a502:0:1da/64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128 primary
set interfaces lo0 unit 1 family inet6 address 2001:db8::2/128
set interfaces lo0 unit 1 family inet6 address 2001:db8::3/128
set routing-options rib inet6.0 static route 2001:db8::5/128 next-hop
   2001:db8:0:1:2a0:a502:0:19da
```

Device E

```
set interfaces fe-1/2/0 unit 25 description to-A
set interfaces fe-1/2/0 unit 25 family inet6 address 2001:db8:0:1:2a0:a502:0:19da/64
set interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da
```

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 basic static routes:

1. On Device A, configure the interfaces.

```
[edit interfaces]
set interfaces fe-1/2/0 unit 1 description to-E
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1:2a0:a502:0:1da/64

set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128 primary
```

```
set interfaces lo0 unit 1 family inet6 address 2001:db8::2/128
set interfaces lo0 unit 1 family inet6 address 2001:db8::3/128
```

2. On Device A, create a static route to Device E's loopback address and set the next-hop address.

   This ensures that Device A has a route back to Device E.

```
[edit routing-options]
set rib inet6.0 static route 2001:db8::5/128 next-hop 2001:db8:0:1:2a0:a502:0:19da
```

3. If you are done configuring Device A, commit the configuration.

```
[edit interfaces]
user@A# commit
```

4. On Device E, configure the interfaces.

```
[edit]
set interfaces fe-1/2/0 unit 25 description to-A
set interfaces fe-1/2/0 unit 25 family inet6 address 2001:db8:0:1:2a0:a502:0:19da/64
set interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
```

5. On Device E, create a static default route and set the next-hop address.

```
[edit routing-options]
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da
```

6. If you are done configuring Device E, commit the configuration.

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

Results

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

Device A

```
user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-E;
    family inet6 {
      address 2001:db8:0:1:2a0:a502:0:1da/64;
    }
  }
}
lo0 {
  unit 1 {
    family inet6 {
      address 2001:db8::1/128 {
        primary;
      }
      address 2001:db8::2/128;
      address 2001:db8::3/128;
    }
  }
}
```

```
user@A# show routing-options
rib inet6.0 {
  static {
    route 2001:db8::5/128 next-hop 2001:db8:0:1:2a0:a502:0:19da;
  }
}
```

Device E

```
user@E# show interfaces
fe-1/2/0 {
  unit 25 {
    description to-A;
    family inet6 {
      address 2001:db8:0:1:2a0:a502:0:19da/64;
    }
  }
}
lo0 {
  unit 5 {
    family inet6 {
      address 2001:db8::5/128;
    }
  }
}

user@E# show routing-options
rib inet6.0 {
  static {
    route ::/0 next-hop 2001:db8:0:1:2a0:a502:0:1da;
  }
}
```

Verification

Confirm that the configuration is working properly.

- Checking the Routing Tables on page 18
- Pinging the Remote Addresses on page 19

Checking the Routing Tables

Purpose Make sure that the static routes appear in the routing tables of Device A and Device E.

Action

```
user@A> show route protocol static
inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128    *[Static/5] 00:27:46
                    > to 2001:db8:0:1:2a0:a502:0:19da via fe-1/2/0.1

user@E> show route protocol static
inet6.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0               *[Static/5] 00:19:11
                    > to 2001:db8:0:1:2a0:a502:0:1da via fe-1/2/0.25
```

Meaning The static routes are in the routing tables.

Pinging the Remote Addresses

Purpose Verify that the static routes are working.

From Device A, ping one of the loopback interface addresses on Device E.

From Device E, ping one of the loopback interface addresses on Device A.

Action

```
user@A> ping 2001:db8::5
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a502:0:1da --> 2001:db8::5
16 bytes from 2001:db8::5, icmp_seq=0 hlim=64 time=1.790 ms
16 bytes from 2001:db8::5, icmp_seq=1 hlim=64 time=1.529 ms
16 bytes from 2001:db8::5, icmp_seq=2 hlim=64 time=1.531 ms

user@E> ping 2001:db8::3
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a502:0:19da --> 2001:db8::3
16 bytes from 2001:db8::3, icmp_seq=0 hlim=64 time=2.146 ms
16 bytes from 2001:db8::3, icmp_seq=1 hlim=64 time=1.964 ms
16 bytes from 2001:db8::3, icmp_seq=2 hlim=64 time=1.550 ms
```

Example: Configuring Static Routing on Logical Systems

- Understanding Basic Static Routing on page 19
- Example: Configuring Static Routes Between Logical Systems Within the Same Router on page 20

Understanding Basic Static Routing

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

You can specify options that define additional information about static routes that is included with the route when it is installed in the routing table. All static options are optional.

Example: Configuring Static Routes Between Logical Systems Within the Same Router

This example shows how to configure static routes between logical systems. The logical systems are configured in a single physical router and are connected by logical tunnel interfaces.

- Requirements on page 20
- Overview on page 20
- Configuration on page 21
- Verification on page 24

Requirements

You must connect the logical systems by using logical tunnel (**lt**) interfaces. See *Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces*.

Overview

A static route is a hard-coded path in the device that specifies how the route gets to a certain subnet by using a certain path. Routers that are connected to stub networks are often configured to use static routes. A *stub network* is a network with no knowledge of other networks. Stub networks send non-local traffic by way of a single path, with the network aware only of a default route to non-local destinations. In this example, you configure Logical System LS1 with a static route to the 10.10.10.0/30 network and define the next-hop address as 192.168.10.2. You also configure Logical System LS1 with a static route to the 192.168.10.0/30 network and define a next-hop address of 10.10.10.1.

Figure 3 on page 21 shows the sample network.

Figure 3: Static Routes Between Logical Systems

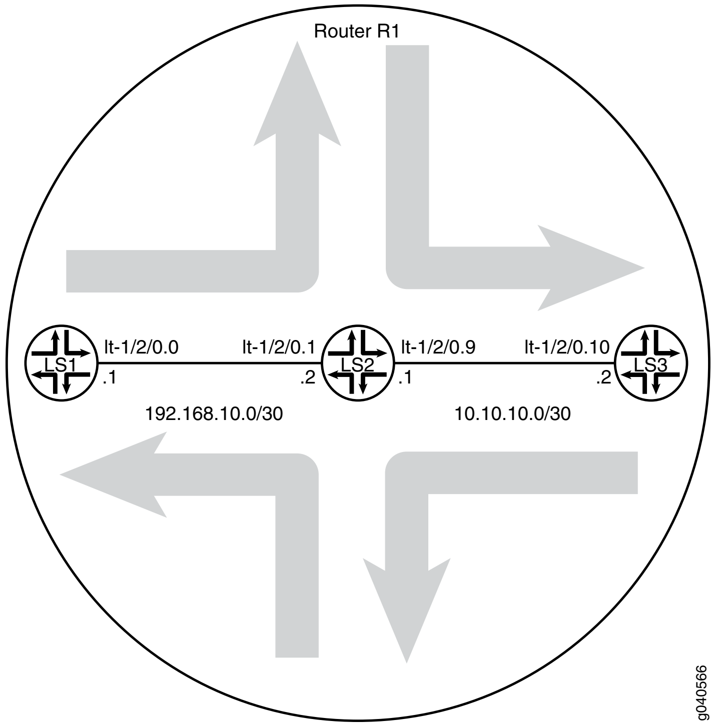

Configuration

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.

```
set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS2
set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 1
set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 192.168.10.1/30
set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 0
set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 192.168.10.2/30
set logical-systems LS2 interfaces lt-1/2/0 unit 9 description LS2->LS3
set logical-systems LS2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 9 peer-unit 10
set logical-systems LS2 interfaces lt-1/2/0 unit 9 family inet address 10.10.10.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 10 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 10 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 10 peer-unit 9
set logical-systems LS3 interfaces lt-1/2/0 unit 10 family inet address 10.10.10.2/30
set logical-systems LS1 routing-options static route 10.10.10.0/30 next-hop 192.168.10.2
set logical-systems LS3 routing-options static route 192.168.10.0/30 next-hop 10.10.10.1
```

Step-by-Step Procedure

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

To configure static routes between logical systems:

1. Run the **show interfaces terse** command to verify that the router has a logical tunnel (**lt**) interface.

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
so-0/0/0                up    down
so-0/0/1                up    down
so-0/0/2                up    down
so-0/0/3                up    down
gr-1/2/0                up    up
ip-1/2/0                up    up
lt-1/2/0                up    up
...
```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```
[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address
  192.168.10.1/30
```

3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 0
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address
  192.168.10.2/30
```

4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 description LS2->LS3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 peer-unit 10
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 family inet address
  10.10.10.1/30
```

5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.

```
[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 description LS3->LS2
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 encapsulation
  ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 peer-unit 9
```



```
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 family inet address
  10.10.10.2/30
```

6. Configure the static route on Logical System LS1 connecting to the 10.10.10.0/30 network.

```
[edit]
user@host# set logical-systems LS1 routing-options static route 10.10.10.0/30
  next-hop 192.168.10.2
```

7. Configure the static route on Logical System LS3 connecting to the 192.168.10.0/30 network.

```
[edit]
user@host# set logical-systems LS3 routing-options static route 192.168.10.0/30
  next-hop 10.10.10.1
```

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

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

Results

Confirm your configuration by issuing the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems
LS1 {
  interfaces {
    lt-1/2/0 {
      unit 0 {
        description LS1->LS2;
        encapsulation ethernet;
        peer-unit 1;
        family inet {
          address 192.168.10.1/30;
        }
      }
    }
  }
  routing-options {
    static {
      route 10.10.10.0/30 next-hop 192.168.10.2;
    }
  }
}
LS2 {
  interfaces {
    lt-1/2/0 {
      unit 1 {
        description LS2->LS1;
        encapsulation ethernet;
        peer-unit 0;
        family inet {
          address 192.168.10.2/30;
```

```
            }
          }
          unit 9 {
            description LS2->LS3;
            encapsulation ethernet;
            peer-unit 10;
            family inet {
              address 10.10.10.1/30;
            }
          }
        }
      }
    }
    LS3 {
      interfaces {
        lt-1/2/0 {
          unit 10 {
            description LS3->LS2;
            encapsulation ethernet;
            peer-unit 9;
            family inet {
              address 10.10.10.2/30;
            }
          }
        }
      }
      routing-options {
        static {
          route 192.168.10.0/30 next-hop 10.10.10.1;
        }
      }
    }
```

Verification

Confirm that the configuration is working properly.

- Verifying That the Logical Systems Are Up on page 24
- Verifying Connectivity Between the Logical Systems on page 25

Verifying That the Logical Systems Are Up

Purpose Make sure that the interfaces are properly configured.

Action

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
...
lt-1/2/0                up    up
lt-1/2/0.0              up    up   inet     192.168.10.1/30
lt-1/2/0.1              up    up   inet     192.168.10.2/30
lt-1/2/0.9              up    up   inet     10.10.10.1/30
lt-1/2/0.10             up    up   inet     10.10.10.2/30
...
```

Verifying Connectivity Between the Logical Systems

Purpose Make sure that the static routes appear in the routing tables of Logical Systems LS1 and LS3. Also, make sure that the logical systems can ping each other.

Action

```
user@host> show route logical-system LS1
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      *[Static/5] 18:43:25
                    > to 192.168.10.2 via lt-1/2/0.0
192.168.10.0/30    *[Direct/0] 18:43:25
                    > via lt-1/2/0.0
192.168.10.1/32    *[Local/0] 18:43:25
                      Local via lt-1/2/0.0

user@host> show route logical-system LS3
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      *[Direct/0] 23:11:21
                    > via lt-1/2/0.10
10.10.10.2/32      *[Local/0] 23:11:21
                      Local via lt-1/2/0.10
192.168.10.0/30    *[Static/5] 00:23:31
                    > to 10.10.10.1 via lt-1/2/0.10
```

From LS1, Ping LS3

```
user@host> set cli logical-system LS1

user@host:LS1> ping 10.10.10.2
PING 10.10.10.2 (10.10.10.2): 56 data bytes
64 bytes from 10.10.10.2: icmp_seq=0 ttl=63 time=1.263 ms
64 bytes from 10.10.10.2: icmp_seq=1 ttl=63 time=1.086 ms
64 bytes from 10.10.10.2: icmp_seq=2 ttl=63 time=1.077 ms
```

From LS3, Ping LS1

```
user@host> set cli logical-system LS3

user@host:LS3> ping 192.168.10.1
PING 192.168.10.1 (192.168.10.1): 56 data bytes
64 bytes from 192.168.10.1: icmp_seq=0 ttl=63 time=10.781 ms
64 bytes from 192.168.10.1: icmp_seq=1 ttl=63 time=1.167 ms
64 bytes from 192.168.10.1: icmp_seq=2 ttl=63 time=1.152 ms
```

Related Documentation

- *Introduction to Logical Systems*

Example: Configuring Static Route Preferences and Qualified Next Hops

Understanding Static Route Preferences and Qualified Next Hops

A static route destination address can have multiple next hops associated with it. In this case, multiple routes are inserted into the routing table, and route selection must occur. Because the primary criterion for route selection is the route preference, you can control the routes that are used as the primary route for a particular destination by setting the route preference associated with a particular next hop. The routes with a lower route preference are always used to route traffic. When you do not set a preferred route, traffic is alternated between routes in round-robin fashion.

In general, the default properties assigned to a static route apply to all the next-hop addresses configured for the static route. If, however, you want to configure two possible next-hop addresses for a particular route and have them treated differently, you can define one as a qualified next hop.

Qualified next hops allow you to associate one or more properties with a particular next-hop address. You can set an overall preference for a particular static route and then specify a different preference for the qualified next hop. For example, suppose two next-hop addresses (10.10.10.10 and 10.10.10.7) are associated with the static route 192.168.47.5/32. A general preference is assigned to the entire static route, and then a different preference is assigned to only the qualified next-hop address 10.10.10.7. For example:

```
route 192.168.47.5/32 {
  next-hop 10.10.10.10;
  qualified-next-hop 10.10.10.7 {
    preference 2;
  }
  preference 6;
}
```

In this example, the qualified next hop 10.10.10.7 is assigned the preference 2, and the next-hop 10.10.10.10 is assigned the preference 6.

NOTE: **The preference and metric options only apply to the qualified next hops. The qualified next-hop preference and metric override the route preference and metric (for that specific qualified next hop), similar to how the route preference overrides the default preference and metric (for that specific route).**

Example: Configuring Static Route Preferences and Qualified Next Hops

This example shows how to control static route selection.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, the static route 192.168.47.0/24 has two possible next hops. Because one link has higher bandwidth, this link is the preferred path. To enforce this preference, the **qualified-next-hop** statement is included in the configuration on both devices. See Figure 4 on page 27.

Figure 4: Controlling Static Route Selection

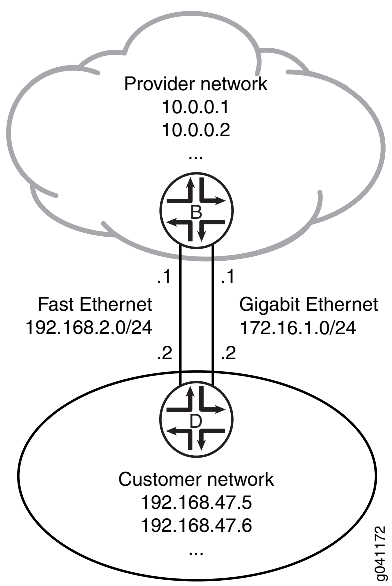

Configuration

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 B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces fe-1/2/1 unit 2 description secondary-B->D
set interfaces fe-1/2/1 unit 2 family inet address 192.168.2.1/24
```



```
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2 preference
  25
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces fe-1/2/1 unit 3 description secondary-D->B
set interfaces fe-1/2/1 unit 3 family inet address 192.168.2.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1 preference 25
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To control static route selection:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set fe-1/2/1 unit 2 description secondary-B->D
user@B# set fe-1/2/1 unit 2 family inet address 192.168.2.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, configure a static route to the customer network.

```
[edit routing-options static route 192.168.47.0/24]
user@B# set next-hop 172.16.1.2
```

3. On Device B, configure a backup route to the customer network.

```
[edit routing options static route 192.168.47.0/24]
user@B# set qualified-next-hop 192.168.2.2 preference 25
```

4. On Device D, configure the interfaces.

```
[edit interfaces]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set fe-1/2/1 unit 3 description secondary-D->B
user@D# set fe-1/2/1 unit 3 family inet address 192.168.2.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

5. On Device D, configure a static default route to external networks.

```
[edit routing options static route 0.0.0.0/0]
user@D# set next-hop 172.16.1.1
```

6. On Device D, configure a backup static default route to external networks.

```
[edit routing options static route 0.0.0.0/0]
```



```
user@D# set qualified-next-hop 192.168.2.1 preference 25
```

Results Confirm your configuration by issuing the **show interfaces** 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@B# show interfaces
ge-1/2/0 {
   unit 0 {
      description B->D;
      family inet {
         address 172.16.1.1/24;
      }
   }
}
fe-1/2/1 {
   unit 2 {
      description secondary-B->D;
      family inet {
         address 192.168.2.1/24;
      }
   }
}
lo0 {
   unit 57 {
      family inet {
         address 10.0.0.1/32;
         address 10.0.0.2/32;
      }
   }
}

user@B# show routing-options
static {
   route 192.168.47.0/24 {
      next-hop 172.16.1.2;
      qualified-next-hop 192.168.2.2 {
         preference 25;
      }
   }
}

user@D# show interfaces
ge-1/2/0 {
   unit 1 {
      description D->B;
      family inet {
         address 172.16.1.2/24;
      }
   }
}
fe-1/2/1 {
   unit 3 {
      description secondary-D->B;
      family inet {
         address 192.168.2.2/24;
```

```
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.47.5/32;
            address 192.168.47.6/32;
        }
    }
}
```

```
user@D# show routing-options
static {
    route 0.0.0.0/0 {
        next-hop 172.16.1.1;
        qualified-next-hop 192.168.2.1 {
            preference 25;
        }
    }
}
```

If you are done configuring the devices, enter **commit** from configuration mode on both devices.

Verification

Confirm that the configuration is working properly.

- Checking the Routing Tables on page 30
- Pinging the Remote Addresses on page 31
- Making Sure That the Backup Route Becomes the Active Route on page 31

Checking the Routing Tables

Purpose Make sure that the static routes appear in the routing tables of Device B and Device D.

Action

```
user@B> show route protocol static
inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.47.0/24    *[Static/5] 02:02:03
                    > to 172.16.1.2 via ge-1/2/0.0
                    [Static/25] 01:58:21
                    > to 192.168.2.2 via fe-1/2/1.2
```

```
user@D> show route protocol static
inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 02:02:12
                    > to 172.16.1.1 via ge-1/2/0.1
                    [Static/25] 01:58:31
                    > to 192.168.2.1 via fe-1/2/1.3
```

Meaning The asterisks (*) in the routing tables show the active routes. The backup routes are listed next.

Pinging the Remote Addresses

Purpose Verify that the static routes are working.

From Device B, ping one of the loopback interface addresses on Device D.

From Device D, ping one of the loopback interface addresses on Device B.

Action

```
user@B> ping 192.168.47.5
PING 192.168.47.5 (192.168.47.5): 56 data bytes
64 bytes from 192.168.47.5: icmp_seq=0 ttl=64 time=156.126 ms
64 bytes from 192.168.47.5: icmp_seq=1 ttl=64 time=120.393 ms
64 bytes from 192.168.47.5: icmp_seq=2 ttl=64 time=175.361 ms
```

```
user@D> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.315 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=31.819 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.268 ms
```

Making Sure That the Backup Route Becomes the Active Route

Purpose If the primary route becomes unusable, make sure that the backup secondary route becomes active.

Action

1. Disable the active route by deactivating the ge-1/2/0.0 interface on Device B.

   user@B# deactivate interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
   user@B# commit

2. Check Device B's routing table.

   ```
   user@B> show route protocol static
   inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
   + = Active Route, - = Last Active, * = Both
   ```



```
192.168.47.0/24     *[Static/25] 02:06:24
                     > to 192.168.2.2 via fe-1/2/1.2
```

Meaning The backup route has become the active route.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Configuring Static Routing on Logical Systems on page 19

Example: Controlling Static Routes in Routing and Forwarding Tables

Understanding Static Route Control in Routing and Forwarding Tables

You can control the importation of static routes into the routing and forwarding tables in a number of ways. Primary ways include assigning one or more of the following attributes to the route:

- **retain**—Keeps the route in the forwarding table after the routing process shuts down or the device reboots.
- **no-readvertise**—Prevents the route from being readvertised to other routing protocols.
- **passive**—Rejects traffic destined for the route.

This topic includes the following sections:

Route Retention

By default, static routes are not retained in the forwarding table when the routing process shuts down. When the routing process starts up again, any routes configured as static routes must be added to the forwarding table again. To avoid this latency, routes can be flagged as **retain**, so that they are kept in the forwarding table even after the routing process shuts down. Retention ensures that the routes are always in the forwarding table, even immediately after a system reboot.

Readvertisement Prevention

Static routes are eligible for readvertisement by other routing protocols by default. In a stub area where you might not want to readvertise these static routes under any circumstances, you can flag the static routes as **no-readvertise**.

Forced Rejection of Passive Route Traffic

Generally, only active routes are included in the routing and forwarding tables. If a static route's next-hop address is unreachable, the route is marked **passive**, and it is not included in the routing or forwarding tables. To force a route to be included in the routing tables regardless of next-hop reachability, you can flag the route as **passive**. If a route is flagged **passive** and its next-hop address is unreachable, the route is included in the routing table, and all traffic destined for the route is rejected.

Example: Preventing a Static Route from Being Readvertised

This example shows how to prevent a static route from being readvertised into OSPF, thereby preventing the route from appearing in the routing and forwarding tables.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

This example shows how to configure a routing policy that readvertises static routes into OSPF, with the exception of one static route that is not readvertised because it is tagged with the **no-readvertise** statement.

Figure 5 on page 33 shows the sample network.

Figure 5: Customer Routes Connected to a Service Provider

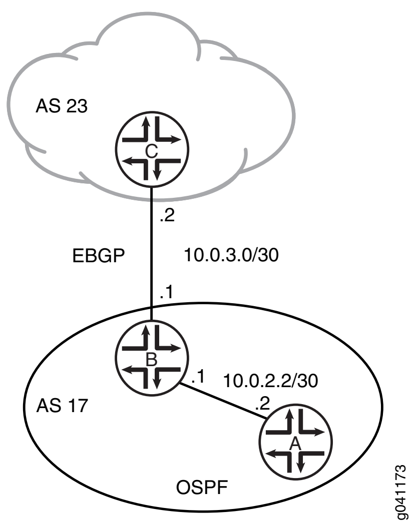

Configuration

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 A

```
set interfaces fe-1/2/0 unit 4 description A->B
set interfaces fe-1/2/0 unit 4 family inet address 10.0.2.2/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.4
```

Device B

```
set interfaces fe-1/2/0 unit 3 description B->A
set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces fe-1/2/1 unit 6 description B->C
set interfaces fe-1/2/1 unit 6 family inet address 10.0.3.1/30
set protocols bgp group ext type external
set protocols bgp group ext peer-as 23
set protocols bgp group ext neighbor 10.0.3.2
set protocols ospf export send-static
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set policy-options policy-statement send-static from protocol static
set policy-options policy-statement send-static then accept
set routing-options static route 0.0.0.0/0 next-hop 10.0.3.2
set routing-options static route 192.168.0.0/24 next-hop 10.0.3.2
set routing-options static route 192.168.0.0/24 no-readvertise
set routing-options autonomous-system 17
```

Device C

```
set interfaces fe-1/2/0 unit 7 description B->C
set interfaces fe-1/2/0 unit 7 family inet address 10.0.3.2/30
set interfaces lo0 unit 5 family inet address 192.168.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext peer-as 17
set protocols bgp group ext neighbor 10.0.3.1
set routing-options autonomous-system 23
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Configure the interface to Device B.

   ```
   [edit interfaces fe-1/2/0 unit 4]
   user@A# set description A->B
   user@A# set family inet address 10.0.2.2/30
   ```

2. Configure OSPF to form an OSPF peer relationship with Device B.

   ```
   [edit protocols ospf area 0.0.0.0]
   user@A# set interface fe-1/2/0.4
   ```

Step-by-Step Procedure

To configure Device B:

1. Configure the interfaces to Device A and Device C.

   ```
   [edit interfaces]
   user@B# set fe-1/2/0 unit 3 description B->A
   user@B# set fe-1/2/0 unit 3 family inet address 10.0.2.1/30
   user@B# set fe-1/2/1 unit 6 description B->C
   user@B# set fe-1/2/1 unit 6 family inet address 10.0.3.1/30
   ```

2. Configure one or more static routes and the autonomous system (AS) number.

```
[edit routing-options]
user@B# set static route 0.0.0.0/0 next-hop 10.0.3.2
user@B# set static route 192.168.0.0/24 next-hop 10.0.3.2
user@B# set autonomous-system 17
```

3. Configure the routing policy.

   This policy exports static routes from the routing table into OSPF.

   ```
   [edit policy-options policy-statement send-static]
   user@B# set from protocol static
   user@B# set then accept
   ```

4. Include the **no-readvertise** statement to prevent the 192.168.0.0/24 route from being exported into OSPF.

   ```
   [edit routing-options]
   user@B# set static route 192.168.0.0/24 no-readvertise
   ```

5. Configure the routing protocols.

   The BGP configuration forms an external BGP (EBGP) peer relationship with Device C.

   The OSPF configuration forms an OSPF peer relationship with Device A and applies the **send-static** routing policy.

   ```
   [edit protocols]
   user@B# set bgp group ext type external
   user@B# set bgp group ext peer-as 23
   user@B# set bgp group ext neighbor 10.0.3.2
   user@B# set ospf export send-static
   user@B# set ospf area 0.0.0.0 interface fe-1/2/0.3
   ```

Step-by-Step Procedure

To configure Device C:

1. Create the interface to Device B, and configure the loopback interface.

   ```
   [edit interfaces ]
   user@C# set fe-1/2/0 unit 7 description B->C
   user@C# set fe-1/2/0 unit 7 family inet address 10.0.3.2/30
   user@C# set lo0 unit 5 family inet address 192.168.0.1/32
   ```

2. Configure the EBGP peering session with Device B.

   ```
   [edit protocols bgp group ext]
   user@C# set type external
   user@C# set peer-as 17
   user@C# set neighbor 10.0.3.1
   ```

3. Configure the AS number.

   ```
   [edit routing-options]
   user@C# set autonomous-system 23
   ```

Results

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

Device A

```
user@A# show interfaces
fe-1/2/0 {
   unit 4 {
      description A->B;
      family inet {
         address 10.0.2.2/30;
      }
   }
}

user@A# show protocols
ospf {
   area 0.0.0.0 {
      interface fe-1/2/0.4;
   }
}
```

Device B

```
user@B# show interfaces
interfaces {
   fe-1/2/0 {
      unit 3 {
         description B->A;
         family inet {
            address 10.0.2.1/30;
         }
      }
   }
   fe-1/2/1 {
      unit 6 {
         description B->C;
         family inet {
            address 10.0.3.1/30;
         }
      }
   }
}

user@B# show policy-options
policy-statement send-static {
   from protocol static;
   then accept;
}

user@B# show protocols
bgp {
   group ext {
      type external;
      peer-as 23;
      neighbor 10.0.3.2;
   }
}
ospf {
   export send-static;
   area 0.0.0.0 {
      interface fe-1/2/0.3;
   }
```



```
}

user@B# show routing-options
static {
   route 0.0.0.0/0 next-hop 10.0.3.2;
   route 192.168.0.0/24 {
      next-hop 10.0.3.2;
      no-readvertise;
   }
}
autonomous-system 17;
```

Device C

```
user@C# show interfaces
fe-1/2/0 {
   unit 7 {
      description B->C;
      family inet {
         address 10.0.3.2/30;
      }
   }
}
lo0 {
   unit 5 {
      family inet {
         address 192.168.0.1/32;
      }
   }
}

user@C# show protocols
bgp {
   group ext {
      type external;
      peer-as 17;
      neighbor 10.0.3.1;
   }
}

user@C# show routing-options
autonomous-system 23;
```

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

Verification

Confirm that the configuration is working properly.

Checking the Routing Table

Purpose Make sure that the **no-readvertise** statement is working.

Action

1. On Device A, run the **show route protocol ospf** command to make sure that the 192.168.0.0/24 route does not appear in Device A's routing table.

   ```
   user@A> show route protocols ospf
   inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
   + = Active Route, - = Last Active, * = Both
   ```



```
0.0.0.0/0          *[OSPF/150] 00:03:15, metric 0, tag 0
                    > to 10.0.2.1 via fe-1/2/0.4
224.0.0.5/32       *[OSPF/10] 00:04:07, metric 1
                      MultiRecv
```

2. On Device B, deactivate the **no-readvertise** statement.

   user@B# deactivate routing-options static route 192.168.0.0/24 no-readvertise

3. On Device A, rerun the **show route protocol ospf** command to make sure that the 192.168.0.0/24 route appears in Device A's routing table.

```
user@A> show route protocols ospf
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:04:24, metric 0, tag 0
                    > to 10.0.2.1 via fe-1/2/0.4
192.168.0.0/24     *[OSPF/150] 00:00:15, metric 0, tag 0
                    > to 10.0.2.1 via fe-1/2/0.4
224.0.0.5/32       *[OSPF/10] 00:05:16, metric 1
                      MultiRecv
```

Meaning The **no-readvertise** statement is working as expected.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Configuring Static Routing on Logical Systems on page 19
- Example: Configuring Static Route Preferences and Qualified Next Hops on page 26

Examples: Creating a Routing Table and Populating It with Routes

Understanding Junos OS Routing Tables

Junos OS automatically creates and maintains several routing tables. Each routing table is used for a specific purpose. In addition to these automatically created routing tables, you can create your own routing tables.

Each routing table populates a portion of the forwarding table. Thus, the forwarding table is partitioned based on routing tables. This allows for specific forwarding behavior for each routing table. For example, for VPNs, each VPN-based routing table has its own VPN-specific partition in the forwarding table.

It is common for the routing software to maintain unicast routes and multicast routes in different routing tables. You also might have policy considerations that would lead you to create separate routing tables to manage the propagation of routing information.

Creating routing tables is optional. If you do not create any, Junos OS uses its default routing tables, which are as follows:

- **inet.0**—For IP version 4 (IPv4) unicast routes. This table stores interface local and direct routes, static routes, and dynamically learned routes.
- **inet.1**—For the IPv4 multicast forwarding cache. This table stores the IPv4 (S,G) group entries that are dynamically created as a result of join state information.
- **inet.2**—For subsequent address family indicator (SAFI) 2 routes, when multiprotocol BGP (MBGP) is enabled. This table stores unicast routes that are used for multicast reverse-path-forwarding (RPF) lookup. The routes in this table can be used by the Distance Vector Multicast Routing Protocol (DVMRP), which requires a specific RPF table. In contrast, Protocol Independent Multicast (PIM) does not need this table because it can perform RPF checks against the inet.0 table. You can import routes from inet.0 into inet.2 using routing information base (RIB) groups, or install routes directly into inet.2 from a multicast routing protocol.
- **inet.3**—For IPv4 MPLS. This table stores the egress address of an MPLS label-swiched path (LSP), the LSP name, and the outgoing interface name. This routing table is used only when the local device is the ingress node to an LSP.
- **inet6.0**—For IP version 6 (IPv6) unicast routes. This table stores interface local and direct routes, static routes, and dynamically learned routes.
- ***instance-name*.inet.0**—If you configure a routing instance, Junos OS creates the default unicast routing table ***instance-name*.inet.0**.
- ***instance-name*.inet.2**—If you configure **routing-instances *instance-name* protocols bgp family inet multicast** in a routing instance of type VRF, Junos OS creates the ***instance-name*.inet.2** table.

  Another way to create the ***instance-name*.inet.2** table is to use the **rib-group** statement. See "Example: Exporting Specific Routes from One Routing Table Into Another Routing Table" on page 46.

  NOTE: **Importing inet-vpn multicast routes from the bgp.l3vpn.2 table into the *instance-name*.inet.2 table does not create the *instance-name*.inet.2 table. The import operation works only if the *instance-name*.inet.2 table already exists.**

- ***instance-name*.inetflow.0**—If you configure a flow route, Junos OS creates the flow routing table ***instance-name*.inetflow.0**.
- **bgp.l2vpn.0**—For Layer 2 VPN routes learned from BGP. This table stores routes learned from other provider edge (PE) routers. The Layer 2 routing information is copied into Layer 2 VPN routing and forwarding instances (VRFs) based on target communities.
- **bgp.l3vpn.0**—For Layer 3 VPN routes learned from BGP. This table stores routes learned from other PE routers. Routes in this table are copied into a Layer 3 VRF when there is a matching route table.

- **mpls.0**—For MPLS label switching operations. This table is used when the local device is a transit router.
- **iso.0**—For IS-IS routes. When you are using IS-IS to support IP routing, this table contains only the local device's network entity title (NET).
- **juniper_private**—For Junos OS to communicate internally between the Routing Engine and PIC hardware.

Example: Creating Routing Tables

This example shows how to create a custom routing table.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Creating routing tables is optional. You might have policy considerations that would lead you to create separate routing tables to manage the propagation of routing information. This capability is rarely used, but it is demonstrated here for completeness.

If you do not create any routing tables, Junos OS uses its default routing tables.

NOTE: If you want to add static, aggregate, generated, or martian routes only to the default IPv4 unicast routing table (**inet.0**), you do not have to create any routing tables because, by default, these routes are added to **inet.0**. You can add these routes by including the **static**, **aggregate**, **generate**, and **martians** statements.

To explicitly create a routing table, include the **rib** statement and child statements under the **rib** statement.

The routing table name, ***routing-table-name***, includes the protocol family, optionally followed by a period and a number. The protocol family can be **inet** for the IPv4 family, **inet6** for the IPv6 family, or **iso** for the International Standards Organization (ISO) protocol family. The number represents the routing instance. The first instance is 0.

This example shows how to configure a custom IPv4 routing table called inet.14. The example also shows how to populate the routing table with a single static route.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Configuration

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.

```
set routing-options rib inet.14 static route 10.2.0.0/16 discard
```

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 create a routing table:

1. Configure the routing table.

   ```
   [edit routing-options]
   user@host# set rib inet.14 static route 10.2.0.0/16 discard
   ```

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

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

Results

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

```
user@host# show routing-options
rib inet.14 {
   static {
      route 10.2.0.0/16 discard;
   }
}
```

Verification

Confirm that the configuration is working properly.

- Checking the Routing Table on page 41

Checking the Routing Table

Purpose

Make sure that the static route appears in the custom routing table.

Action

```
user@host> show route table inet.14
inet.14: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.2.0.0/16        *[Static/5] 00:00:09
                      Discard
```

Meaning The static route is in the custom routing table.

Example: Importing Direct and Static Routes Into a Routing Instance

This example shows how to populate the routing table that is created when you configure a virtual router.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

You can install routes into more than one routing table. For example, you might want a simple configuration that allows you to install a static route into the default routing table **inet.0**, as well as a second routing table **vpna.inet.0**. Instead of configuring the same static route for each routing table, you can use routing table groups to insert the route into multiple tables. To create a routing table group, include the **rib-groups** statement.

This example shows how to export static routes, direct routes, and local routes from the default IPv4 unicast routing table (**inet.0**) and import them into the IPv4 unicast routing table of a virtual router called vpna (**vpna.inet.0**).

NOTE: **To explicitly create a routing table, include the rib statement. In this case, you do not need to use the rib statement because when you configure a routing instance, Junos OS automatically creates the routing table *instance-name*.inet.0.**

In this example, Device A and Device B are directly connected to each other. Device A also has a virtual router configured called vpna. Device A's **inet.0** routing table has direct and local routes (also known as interface routes). These routes are imported into vpna's **inet.0** routing table (**vpna.inet.0**). Device A also has a static route configured. This static route is also imported into **vpna.inet.0**.

Configuration

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 A

```
set interfaces fe-1/2/0 unit 4 description A->B
set interfaces fe-1/2/0 unit 4 family inet address 10.0.2.2/30
set interfaces lo0 unit 1 family inet address 1.1.1.1/32
set routing-instances vpna instance-type virtual-router
set routing-options interface-routes rib-group inet group1
set routing-options static rib-group group1
set routing-options static route 192.168.1.0/24 discard
set routing-options rib-groups group1 import-rib inet.0
set routing-options rib-groups group1 import-rib vpna.inet.0
```

Device B

```
set interfaces fe-1/2/0 unit 3 description B->A
set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces lo0 unit 2 family inet address 2.2.2.2/32
```

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 this example:

1. Configure the routing instance.

   ```
   [edit routing-instances]
   user@A# set vpna instance-type virtual-router
   ```

2. Configure the interfaces.

   ```
   [edit interfaces]
   user@A# set fe-1/2/0 unit 4 description A->B
   user@A# set fe-1/2/0 unit 4 family inet address 10.0.2.2/30
   user@A# set lo0 unit 1 family inet address 1.1.1.1/32
   ```

3. Configure one or more static routes.

   ```
   [edit routing-options]
   user@A# set static route 192.168.1.0/24 discard
   ```

4. Include the direct and local routes in a routing table group called group1.

   The **interface-routes** statement specifies the direct and local routes to match against.

   ```
   [edit routing-options]
   user@A# set interface-routes rib-group inet group1
   ```

5. Include all static routes in the group1 routing table group.

   The **static** statement specifies the protocol (static) to match against.

   ```
   [edit routing-options]
   user@A# set static rib-group group1
   ```

6. Configure the primary routing table for group1.

The primary routing table determines the address family of the routing table group. To configure an IPv4 group, specify **inet.0** as the primary table. To configure an IPv6 group, specify **inet6.0** as the primary routing table.

```
[edit routing-options]
user@A# set rib-groups group1 import-rib inet.0
```

7. Configure the secondary routing table for group1.

```
[edit routing-options]
user@A# set rib-groups group1 import-rib vpna.inet.0
```

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

```
[edit routing-options]
user@A# commit
```

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 B:

1. Configure the interfaces.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 3 description B->A
user@B# set fe-1/2/0 unit 3 family inet address 10.0.2.1/30

user@B# set lo0 unit 2 family inet address 2.2.2.2/32
```

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

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

Results

Confirm your configuration by issuing the **show interfaces**, **show routing-instances**, 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@A# show interfaces
fe-1/2/0 {
  unit 4 {
    description A->B;
    family inet {
      address 10.0.2.2/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 1.1.1.1/32;
    }
  }
```



```
}

user@A# show routing-instances
vpna {
   instance-type virtual-router;
}

user@A# show routing-options
interface-routes {
   rib-group inet group1;
}
static {
   rib-group group1;
   route 192.168.1.0/24 discard;
}
rib-groups {
   group1 {
      import-rib [ inet.0 vpna.inet.0 ];
   }
}

user@B# show interfaces
fe-1/2/0 {
   unit 3 {
      description B->A;
      family inet {
         address 10.0.2.1/30;
      }
   }
}
lo0 {
   unit 2 {
      family inet {
         address 2.2.2.2/32;
      }
   }
}
```

Verification

Confirm that the configuration is working properly.

Checking the Routing Tables

Purpose Make sure that the expected routes appear in the routing tables.

Action

```
user@A> show route table inet.0

inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32         *[Direct/0] 02:51:24
                    > via lo0.1
10.0.2.0/30        *[Direct/0] 03:19:06
                    > via fe-1/2/0.4
10.0.2.2/32        *[Local/0] 03:19:07
                      Local via fe-1/2/0.4
192.168.1.0/24     *[Static/5] 00:44:21
                      Discard
```

```
user@A> show route table vpna.inet.0

vpna.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32         *[Direct/0] 02:35:39
                    > via lo0.1
10.0.2.0/30        *[Direct/0] 02:35:39
                    > via fe-1/2/0.4
10.0.2.2/32        *[Local/0] 02:35:39
                      Local via fe-1/2/0.4
192.168.1.0/24     *[Static/5] 00:44:28
                      Discard
```

Meaning The static route and the interface routes appear in both routing tables.

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table

This example shows how to duplicate specific routes from one routing table into another routing table within the same routing instance.

Requirements

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

Overview

This example uses the **auto-export** statement and the **rib-group** statement to accomplish the goal of exporting specific routes from one routing table to another.

Consider the following points:

- When **auto-export** is configured in a routing instance, the **vrf-import** and **vrf-export** policies are examined. Based on the route target and community information in the policies, the **auto-export** function performs route leaking among the local routing instance inet.0 tables.

- You can use the **rib-group** statement if it is necessary to import routes into tables other than *instance*.inet.0. If a RIB group is used, the RIB group's **export-rib** and **import-policy** statements are not used. Only the **import-rib** statement is used. To use a RIB group with **auto-export**, the routing instance should specify explicit **vrf-import** and **vrf-export** policies. The **vrf-import** and **vrf-export** policies can be extended to contain additional terms to filter routes as needed for the RIB group.

In this example, access-internal routes are added into the vpna.inet.0 routing table. The access-internal routes are also duplicated into the vpna.inet.2 routing table.

Configuration

- Configuring Specific Route Export Between Routing Tables on page 48

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.

```
set interfaces fe-1/3/1 vlan-tagging
set interfaces fe-1/3/1 unit 0 vlan-id 512
set interfaces fe-1/3/1 unit 0 family inet address 10.168.100.3/24
set interfaces lo0 unit 0 family inet address 192.168.3.3/32
set routing-options rib-groups rib-group-vpna-access-internal import-rib vpna.inet.2
set routing-options autonomous-system 63000
set policy-options policy-statement vpna-export term a from protocol bgp
set policy-options policy-statement vpna-export term a then community add vpna-comm
set policy-options policy-statement vpna-export term a then accept
set policy-options policy-statement vpna-export term b from protocol access-internal
set policy-options policy-statement vpna-export term b then accept
set policy-options policy-statement vpna-export term c then reject
set policy-options policy-statement vpna-import term a from protocol bgp
set policy-options policy-statement vpna-import term a from community vpna-comm
set policy-options policy-statement vpna-import term a then accept
set policy-options policy-statement vpna-import term b from instance vpna
set policy-options policy-statement vpna-import term b from protocol access-internal
set policy-options policy-statement vpna-import term b then accept
set policy-options policy-statement vpna-import term c then reject
set policy-options community vpna-comm members target:63000:100
set routing-instances vpna instance-type vrf
set routing-instances vpna interface fe-1/3/1.1
set routing-instances vpna route-distinguisher 100:1
set routing-instances vpna vrf-import vpna-import
set routing-instances vpna vrf-export vpna-export
set routing-instances vpna routing-options auto-export family inet unicast rib-group
  rib-group-vpna-access-internal
set routing-instances vpna protocols bgp group bgp-vpna type external
set routing-instances vpna protocols bgp group bgp-vpna family inet multicast
set routing-instances vpna protocols bgp group bgp-vpna peer-as 100
set routing-instances vpna protocols bgp group bgp-vpna neighbor 10.0.0.10
```


Configuring Specific Route Export Between Routing Tables

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the device:

1. Configure the interfaces.

   ```
   [edit interfaces fe-1/3/1]
   user@host# set vlan-tagging
   user@host# set unit 0 vlan-id 512
   user@host# set unit 0 family inet address 10.168.100.3/24
   [edit interfaces lo0 unit 0]
   user@host# set family inet address 192.168.3.3/32
   ```

2. Configure the routing policy that specifies particular routes for import into vpna.inet.0 and export from vpna.inet.0.

   ```
   [edit policy-options policy-statement vpna-export]
   user@host# set term a from protocol bgp
   user@host# set term a then community add vpna-comm
   user@host# set term a then accept
   user@host# set term b from protocol access-internal
   user@host# set term b then accept
   user@host# set term c then reject
   [edit policy-options policy-statement vpna-import]
   user@host# set term a from protocol bgp
   user@host# set term a from community vpna-comm
   user@host# set term a then accept
   user@host# set term b from instance vpna
   user@host# set term b from protocol access-internal
   user@host# set term b then accept
   user@host# set term c then reject
   [edit policy-options]
   user@host# set community vpna-comm members target:63000:100
   ```

3. Configure the routing instance.

   ```
   [edit routing-instances vpna]
   user@host# set instance-type vrf
   user@host# set interface fe-1/3/1.1
   user@host# set route-distinguisher 100:1
   user@host# set vrf-import vpna-import
   user@host# set vrf-export vpna-export
   ```

   The **vrf-import** and **vrf-export** statements are used to apply the **vpna-import** and **vpna-export** routing policies configured in 2.

4. Configure the RIB group, and import routes into the **vpna.inet.2** routing table.

   ```
   [edit routing-options]
   user@host# set rib-groups rib-group-vpna-access-internal import-rib vpna.inet.2
   ```

5. Configure the **auto-export** statement to enable the routes to be exported from one routing table into another.

```
[edit routing-options]
user@host# set auto-export family inet unicast rib-group
  rib-group-vpna-access-internal
```

6. Configure BGP.

```
[edit routing-instances vpna protocols bgp group bgp-vpna]
user@host# set type external
user@host# set family inet multicast
user@host# set peer-as 100
user@host# set neighbor 100.0.0.10
```

7. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@host# set autonomous-system 63000
```

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

```
user@host# show interfaces
fe-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 512;
    family inet {
      address 10.168.100.3/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.3.3/32;
    }
  }
}
```

```
user@host# show policy-options
policy-statement vpna-export {
  term a {
    from {
      protocol bgp;
    }
    then {
      community add vpna-comm;
      accept;
    }
  }
  term b {
    from protocol access-internal;
    then accept;
  }
  term c {
    then reject;
```

```
    }
}
policy-statement vpna-import {
    term a {
        from {
            protocol bgp;
            community vpna-comm;
        }
        then accept;
    }
    term b {
        from {
            instance vpna;
            protocol access-internal;
        }
        then accept;
    }
    term c {
        then reject;
    }
}
community vpna-comm members target:63000:100;
```

```
user@host# show routing-options
rib-groups {
    rib-group-vpna-access-internal {
        import-rib vpna.inet.2;
    }
}
autonomous-system 63000;
```

```
user@host# show routing-instances
vpna {
    instance-type vrf;
    interface fe-1/3/1.1;
    route-distinguisher 100:1;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
        auto-export {
            family inet {
                unicast {
                    rib-group rib-group-vpna-access-internal;
                }
            }
        }
    }
    protocols {
        bgp {
            group bgp-vpna {
                type external;
                family inet {
                    multicast;
                }
                peer-as 100;
                neighbor 100.0.0.10;
            }
```

```
        }
      }
    }
```

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

Verification

Confirm that the configuration is working properly by running the **show table route vpna.inet.0** and **show route table vpna.inet.2** commands.

Related Documentation

- Examples: Configuring Static Routes on page 9

Example: Configuring Point-to-Multipoint LSPs with Static Routes

Understanding Point-to-Multipoint LSPs

A point-to-multipoint MPLS label-switched path (LSP) is an LDP-signaled or RSVP-signaled LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the inbound (ingress) router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in Figure 6 on page 52. Device PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Device PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Device P1 replicates the packet and forwards it to Routers PE2 and PE3. Device P2 sends the packet to Device PE4.

Figure 6: Point-to-Multipoint LSPs

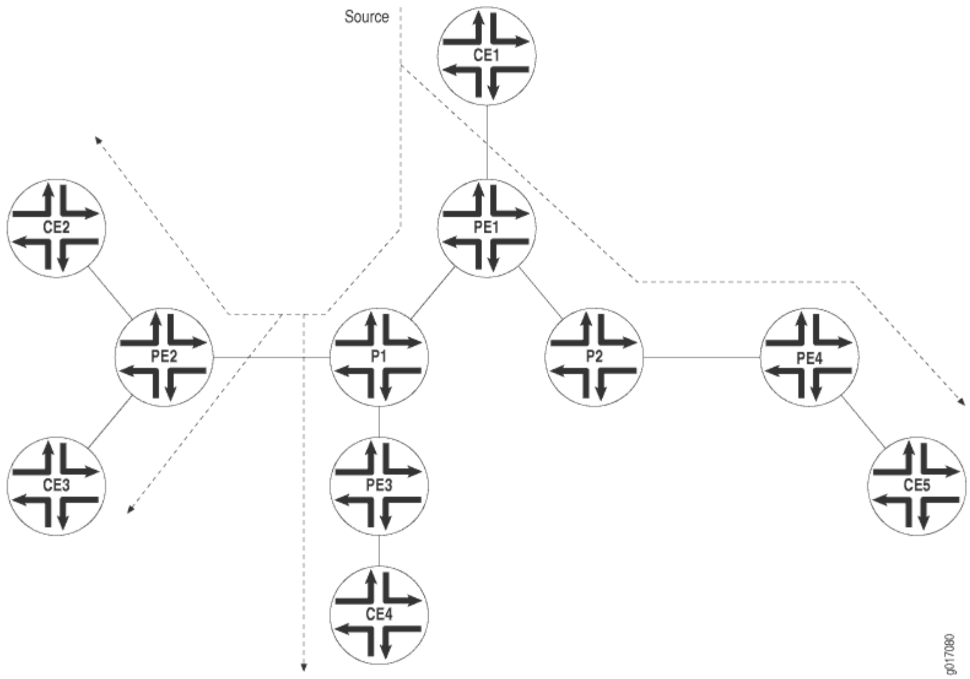

Following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP allows you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an outbound (egress) router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any primary paths fail, traffic can be quickly switched to the bypass.
- You can configure subpaths either statically or dynamically.
- You can enable graceful restart on point-to-multipoint LSPs.

Point-to-Multipoint LSP Configuration Overview

To set up a point-to-multipoint LSP:

1. Configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers.
2. Specify a pathname on the primary LSP and this same path name on each branch LSP.

NOTE: By default, the branch LSPs are dynamically signaled by means of Constrained Shortest Path First (CSPF) and require no configuration. You can alternatively configure the branch LSPs as static paths.

Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP

This example shows how to configure a collection of paths to create an RSVP-signaled point-to-multipoint label-switched path (LSP).

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, multiple routing devices serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. On the provider edge (PE), Device PE1 is the ingress node. The branches go from PE1 to PE2, PE1 to PE3, and PE1 to PE4. Static unicast routes on the ingress node (PE1) point to the egress nodes.

This example also demonstrates static routes with a next hop that is a point-to-multipoint LSP, using the p2mp-lsp-next-hop statement. This is useful when implementing filter-based forwarding.

NOTE: Another option is to use the lsp-next-hop statement to configure a regular point-to-point LSP to be the next hop. Though not shown in this example, you can optionally assign an independent preference and metric to the next hop.

Topology Diagram

Figure 7 on page 54 shows the topology used in this example.

Figure 7: RSVP-Signaled Point-to-Multipoint LSP

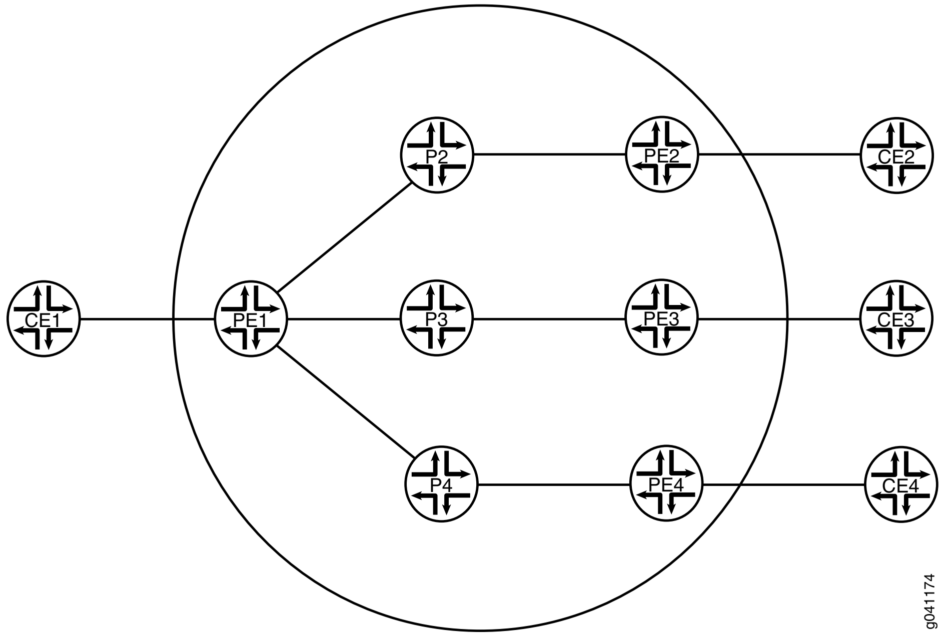

Configuration

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 PE1

```
set interfaces ge-2/0/2 unit 0 description PE1-to-CE1
set interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set interfaces fe-2/0/10 unit 1 description PE1-to-P2
set interfaces fe-2/0/10 unit 1 family inet address 2.2.2.1/24
set interfaces fe-2/0/10 unit 1 family mpls
set interfaces fe-2/0/9 unit 8 description PE1-to-P3
set interfaces fe-2/0/9 unit 8 family inet address 6.6.6.1/24
set interfaces fe-2/0/9 unit 8 family mpls
set interfaces fe-2/0/8 unit 9 description PE1-to-P4
set interfaces fe-2/0/8 unit 9 family inet address 3.3.3.1/24
set interfaces fe-2/0/8 unit 9 family mpls
set interfaces lo0 unit 1 family inet address 100.10.10.10/32
set protocols rsvp interface fe-2/0/10.1
set protocols rsvp interface fe-2/0/9.8
set protocols rsvp interface fe-2/0/8.9
set protocols rsvp interface lo0.1
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path PE1-PE2 to 100.50.50.50
set protocols mpls label-switched-path PE1-PE2 link-protection
set protocols mpls label-switched-path PE1-PE2 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE3 to 100.70.70.70
set protocols mpls label-switched-path PE1-PE3 link-protection
set protocols mpls label-switched-path PE1-PE3 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE4 to 100.40.40.40
set protocols mpls label-switched-path PE1-PE4 link-protection
```

```
set protocols mpls label-switched-path PE1-PE4 p2mp p2mp1
set protocols mpls interface fe-2/0/10.1
set protocols mpls interface fe-2/0/9.8
set protocols mpls interface fe-2/0/8.9
set protocols mpls interface lo0.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface fe-2/0/10.1
set protocols ospf area 0.0.0.0 interface fe-2/0/9.8
set protocols ospf area 0.0.0.0 interface fe-2/0/8.9
set protocols ospf area 0.0.0.0 interface lo0.1
set routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set routing-options router-id 100.10.10.10
```

Device CE1

```
set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30
set interfaces ge-1/3/2 unit 0 description CE1-to-PE1
set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10
```

Device CE2

```
set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30
set interfaces ge-1/3/3 unit 0 description CE2-to-PE2
set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10
```

Device CE3

```
set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30
set interfaces ge-2/0/1 unit 0 description CE3-to-PE3
set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10
```

Device CE4

```
set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30
set interfaces ge-3/1/3 unit 0 description CE4-to-PE4
set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9
```

Configuring the Ingress Label-Switched Router (LSR) (Device PE1)

Step-by-Step Procedure

To configure Device PE1:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit interfaces]
user@PE1# set ge-2/0/2 unit 0 description PE1-to-CE1
user@PE1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
user@PE1# set fe-2/0/10 unit 1 description PE1-to-P2
user@PE1# set fe-2/0/10 unit 1 family inet address 2.2.2.1/24
user@PE1# set fe-2/0/10 unit 1 family mpls
user@PE1# set fe-2/0/9 unit 8 description PE1-to-P3
user@PE1# set fe-2/0/9 unit 8 family inet address 6.6.6.1/24
user@PE1# set fe-2/0/9 unit 8 family mpls
user@PE1# set fe-2/0/8 unit 9 description PE1-to-P4
user@PE1# set fe-2/0/8 unit 9 family inet address 3.3.3.1/24
user@PE1# set fe-2/0/8 unit 9 family mpls
user@PE1# set lo0 unit 1 family inet address 100.10.10.10/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

   [edit protocols]
   user@PE1# **set rsvp interface fe-2/0/10.1**
   user@PE1# **set rsvp interface fe-2/0/9.8**
   user@PE1# **set rsvp interface fe-2/0/8.9**
   user@PE1# **set rsvp interface lo0.1**
   user@PE1# **set mpls interface fe-2/0/10.1**
   user@PE1# **set mpls interface fe-2/0/9.8**
   user@PE1# **set mpls interface fe-2/0/8.9**
   user@PE1# **set mpls interface lo0.1**
   user@PE1# **set ospf area 0.0.0.0 interface ge-2/0/2.0**
   user@PE1# **set ospf area 0.0.0.0 interface fe-2/0/10.1**
   user@PE1# **set ospf area 0.0.0.0 interface fe-2/0/9.8**
   user@PE1# **set ospf area 0.0.0.0 interface fe-2/0/8.9**
   user@PE1# **set ospf area 0.0.0.0 interface lo0.1**

3. Configure the MPLS point-to-multipoint LSPs.

   [edit protocols]
   user@PE1# **set mpls label-switched-path PE1-PE2 to 100.50.50.50**
   user@PE1# **set mpls label-switched-path PE1-PE2 p2mp p2mp1**
   user@PE1# **set mpls label-switched-path PE1-PE3 to 100.70.70.70**
   user@PE1# **set mpls label-switched-path PE1-PE3 p2mp p2mp1**
   user@PE1# **set mpls label-switched-path PE1-PE4 to 100.40.40.40**
   user@PE1# **set mpls label-switched-path PE1-PE4 p2mp p2mp1**

4. (Optional) Enable link protection on the LSPs.

   Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

   [edit protocols]
   user@PE1# **set mpls label-switched-path PE1-PE2 link-protection**
   user@PE1# **set mpls label-switched-path PE1-PE3 link-protection**
   user@PE1# **set mpls label-switched-path PE1-PE4 link-protection**

5. Enable MPLS to perform traffic engineering for OSPF.

   [edit protocols]
   user@PE1# **set mpls traffic-engineering bgp-igp**

   This causes the ingress routes to be installed in the inet.0 routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

6. Enable traffic engineering for OSPF.

   [edit protocols]
   user@PE1# **set ospf traffic-engineering**

   This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

7. Configure the router ID.

   [edit routing-options]
   user@PE1# **set router-id 100.10.10.10**

8. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```
[edit routing-options]
user@PE1# set static route 5.5.5.0/24p2mp-lsp-next-hop p2mp1
user@PE1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
```

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

```
[edit]
user@PE1# commit
```

Configuring the Transit and Egress LSRs (Devices P2, P3, P4, PE2, PE3, and PE4)

Step-by-Step Procedure

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@P2# set interfaces fe-2/0/10 unit 2 description P2-to-PE1
user@P2# set interfaces fe-2/0/10 unit 2 family inet address 2.2.2.2/24
user@P2# set interfaces fe-2/0/10 unit 2 family mpls
user@P2# set interfaces fe-2/0/9 unit 10 description P2-to-PE2
user@P2# set interfaces fe-2/0/9 unit 10 family inet address 5.5.5.1/24
user@P2# set interfaces fe-2/0/9 unit 10 family mpls
user@P2# set interfaces lo0 unit 2 family inet address 100.20.20.20/32
user@PE2# set interfaces ge-2/0/3 unit 0 description PE2-to-CE2
user@PE2# set interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@PE2# set interfaces fe-2/0/10 unit 5 description PE2-to-P2
user@PE2# set interfaces fe-2/0/10 unit 5 family inet address 5.5.5.2/24
user@PE2# set interfaces fe-2/0/10 unit 5 family mpls
user@PE2# set interfaces lo0 unit 5 family inet address 100.50.50.50/32
user@P3# set interfaces fe-2/0/10 unit 6 description P3-to-PE1
user@P3# set interfaces fe-2/0/10 unit 6 family inet address 6.6.6.2/24
user@P3# set interfaces fe-2/0/10 unit 6 family mpls
user@P3# set interfaces fe-2/0/9 unit 11 description P3-to-PE3
user@P3# set interfaces fe-2/0/9 unit 11 family inet address 7.7.7.1/24
user@P3# set interfaces fe-2/0/9 unit 11 family mpls
user@P3# set interfaces lo0 unit 6 family inet address 100.60.60.60/32
user@PE3# set interfaces ge-2/0/1 unit 0 description PE3-to-CE3
user@PE3# set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@PE3# set interfaces fe-2/0/10 unit 7 description PE3-to-P3
user@PE3# set interfaces fe-2/0/10 unit 7 family inet address 7.7.7.2/24
user@PE3# set interfaces fe-2/0/10 unit 7 family mpls
user@PE3# set interfaces lo0 unit 7 family inet address 100.70.70.70/32
user@P4# set interfaces fe-2/0/10 unit 3 description P4-to-PE1
user@P4# set interfaces fe-2/0/10 unit 3 family inet address 3.3.3.2/24
user@P4# set interfaces fe-2/0/10 unit 3 family mpls
user@P4# set interfaces fe-2/0/9 unit 12 description P4-to-PE4
user@P4# set interfaces fe-2/0/9 unit 12 family inet address 4.4.4.1/24
user@P4# set interfaces fe-2/0/9 unit 12 family mpls
user@P4# set interfaces lo0 unit 3 family inet address 100.30.30.30/32
user@PE4# set interfaces ge-2/0/0 unit 0 description PE4-to-CE4
user@PE4# set interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@PE4# set interfaces fe-2/0/10 unit 4 description PE4-to-P4
```

```
user@PE4# set interfaces fe-2/0/10 unit 4 family inet address 4.4.4.2/24
user@PE4# set interfaces fe-2/0/10 unit 4 family mpls
user@PE4# set interfaces lo0 unit 4 family inet address 100.40.40.40/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@P2# set protocols rsvp interface fe-2/0/10.2
user@P2# set protocols rsvp interface fe-2/0/9.10
user@P2# set protocols rsvp interface lo0.2
user@P2# set protocols mpls interface fe-2/0/10.2
user@P2# set protocols mpls interface fe-2/0/9.10
user@P2# set protocols mpls interface lo0.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/9.10
user@P2# set protocols ospf area 0.0.0.0 interface lo0.2
user@PE2# set protocols rsvp interface fe-2/0/10.5
user@PE2# set protocols rsvp interface lo0.5
user@PE2# set protocols mpls interface fe-2/0/10.5
user@PE2# set protocols mpls interface lo0.5
user@PE2# set protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@PE2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.5
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.5
user@P3# set protocols rsvp interface fe-2/0/10.6
user@P3# set protocols rsvp interface fe-2/0/9.11
user@P3# set protocols rsvp interface lo0.6
user@P3# set protocols mpls interface fe-2/0/10.6
user@P3# set protocols mpls interface fe-2/0/9.11
user@P3# set protocols mpls interface lo0.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/9.11
user@P3# set protocols ospf area 0.0.0.0 interface lo0.6
user@PE3# set protocols rsvp interface fe-2/0/10.7
user@PE3# set protocols rsvp interface lo0.7
user@PE3# set protocols mpls interface fe-2/0/10.7
user@PE3# set protocols mpls interface lo0.7
user@PE3# set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@PE3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.7
user@PE3# set protocols ospf area 0.0.0.0 interface lo0.7
user@P4# set protocols rsvp interface fe-2/0/10.3
user@P4# set protocols rsvp interface fe-2/0/9.12
user@P4# set protocols rsvp interface lo0.3
user@P4# set protocols mpls interface fe-2/0/10.3
user@P4# set protocols mpls interface fe-2/0/9.12
user@P4# set protocols mpls interface lo0.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/9.12
user@P4# set protocols ospf area 0.0.0.0 interface lo0.3
user@PE4# set protocols rsvp interface fe-2/0/10.4
user@PE4# set protocols rsvp interface lo0.4
user@PE4# set protocols mpls interface fe-2/0/10.4
user@PE4# set protocols mpls interface lo0.4
user@PE4# set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@PE4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.4
user@PE4# set protocols ospf area 0.0.0.0 interface lo0.4
```


3. Enable traffic engineering for OSPF.

```
[edit]
user@P2# set protocols ospf traffic-engineering
user@P2# set protocols ospf traffic-engineering
user@P3# set protocols ospf traffic-engineering
user@PE2# set protocols ospf traffic-engineering
user@PE3# set protocols ospf traffic-engineering
user@PE4# set protocols ospf traffic-engineering
```

   This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```
[edit]
user@P2# set routing-options router-id 100.20.20.20
user@P3# set routing-options router-id 100.60.60.60
user@P4# set routing-options router-id 100.30.30.30
user@PE2# set routing-options router-id 100.50.50.50
user@PE3# set routing-options router-id 100.70.70.70
user@PE4# set routing-options router-id 100.40.40.40
```

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

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

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

Device PE1

```
user@PE1# show interfaces
ge-2/0/2 {
   unit 0 {
      description R1-to-CE1;
      family inet {
         address 10.0.244.10/30;
      }
   }
}
fe-2/0/10 {
   unit 1 {
      description PE1-to-P2;
      family inet {
         address 2.2.2.1/24;
      }
      family mpls;
   }
}
fe-2/0/9 {
   unit 8 {
      description PE1-to-P2;
      family inet {
         address 6.6.6.1/24;
      }
      family mpls;
```



```
    }
}
fe-2/0/8 {
    unit 9 {
        description PE1-to-P3;
        family inet {
            address 3.3.3.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 100.10.10.10/32;
        }
    }
}

user@PE1# show protocols
rsvp {
    interface fe-2/0/10.1;
    interface fe-2/0/9.8;
    interface fe-2/0/8.9;
    interface lo0.1;
}
mpls {
    traffic-engineering bgp-igp;
    label-switched-path PE1-to-PE2 {
        to 100.50.50.50;
        link-protection;
        p2mp p2mp1;
    }
    label-switched-path PE1-to-PE3 {
        to 100.70.70.70;
        link-protection;
        p2mp p2mp1;
    }
    label-switched-path PE1-to-PE4 {
        to 100.40.40.40;
        link-protection;
        p2mp p2mp1;
    }
    interface fe-2/0/10.1;
    interface fe-2/0/9.8;
    interface fe-2/0/8.9;
    interface lo0.1;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/2.0;
        interface fe-2/0/10.1;
        interface fe-2/0/9.8;
        interface fe-2/0/8.9;
        interface lo0.1;
```



```
   }
}
user@PE1# show routing-options
static {
   route 5.5.5.0/24 {
      p2mp-lsp-next-hop p2mp1;
   }
   route 7.7.7.0/24 {
      p2mp-lsp-next-hop p2mp1;
   }
   route 4.4.4.0/24 {
      p2mp-lsp-next-hop p2mp1;
   }
}
router-id 100.10.10.10;
```

Device P2

```
user@P2# show interfaces
fe-2/0/10 {
   unit 2 {
      description P2-to-PE1;
      family inet {
         address 2.2.2.2/24;
      }
      family mpls;
   }
fe-2/0/9 {
   unit 10 {
      description P2-to-PE2;
      family inet {
         address 5.5.5.1/24;
      }
      family mpls;
   }
}
lo0 {
   unit 2 {
      family inet {
         address 100.20.20.20/32;
      }
   }
}

user@P2# show protocols
rsvp {
   interface fe-2/0/10.2;
   interface fe-2/0/9.10;
   interface lo0.2;
}
mpls {
   interface fe-2/0/10.2;
   interface fe-2/0/9.10;
   interface lo0.2;
}
ospf {
   traffic-engineering;
```

```
  area 0.0.0.0 {
    interface fe-2/0/10.2;
    interface fe-2/0/9.10;
    interface lo0.2;
  }
}
```

user@P2# **show routing-options**

```
router-id 100.20.20.20;
```

Device P3

user@P3# **show interfaces**

```
fe-2/0/10 {
  unit 6 {
    description P3-to-PE1;
    family inet {
      address 6.6.6.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 11 {
    description P3-to-PE3;
    family inet {
      address 7.7.7.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 6 {
    family inet {
      address 100.60.60.60/32;
    }
  }
}
```

user@P3# **show protocols**

```
rsvp {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
mpls {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.6;
    interface fe-2/0/9.11;
    interface lo0.6;
  }
}
```



```
user@P2# show routing-options
router-id 100.60.60.60;
```

Device P4

```
user@P4# show interfaces
fe-2/0/10 {
   unit 3 {
      description P4-to-PE1;
      family inet {
         address 3.3.3.2/24;
      }
      family mpls;
   }
}
fe-2/0/9 {
   unit 12 {
      description P4-to-PE4;
      family inet {
         address 4.4.4.1/24;
      }
      family mpls;
   }
}
lo0 {
   unit 3 {
      family inet {
         address 100.30.30.30/32;
      }
   }
}

user@P4# show protocols
rsvp {
   interface fe-2/0/10.3;
   interface fe-2/0/9.12;
   interface lo0.3;
}
mpls {
   interface fe-2/0/10.3;
   interface fe-2/0/9.12;
   interface lo0.3;
}
ospf {
   traffic-engineering;
   area 0.0.0.0 {
      interface fe-2/0/10.3;
      interface fe-2/0/9.12;
      interface lo0.3;
   }
}

user@P3# show routing-options
router-id 100.30.30.30;
```

Device PE2

```
user@PE2# show interfaces
   ge-2/0/3 {
      unit 0 {
         description PE2-to-CE2;
```

```
            family inet {
                address 10.0.224.10/30;
            }
        }
    }
    fe-2/0/10 {
        unit 5 {
            description PE2-to-P2;
            family inet {
                address 5.5.5.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 5 {
            family inet {
                address 100.50.50.50/32;
            }
        }
    }
}

user@PE2# show protocols
rsvp {
    interface fe-2/0/10.5;
    interface lo0.5;
}
mpls {
    interface fe-2/0/10.5;
    interface lo0.5;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/3.0;
        interface fe-2/0/10.5;
        interface lo0.5;
    }
}

user@PE2# show routing-options
router-id 100.50.50.50;
```

Device PE3

```
user@PE3# show interfaces
    ge-2/0/1 {
        unit 0 {
            description PE3-to-CE3;
            family inet {
                address 10.0.134.10/30;
            }
        }
    }
    fe-2/0/10 {
        unit 7 {
            description PE3-to-P3;
```

```
            family inet {
                address 7.7.7.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 7 {
            family inet {
                address 100.70.70.70/32;
            }
        }
    }
}
```

```
user@PE3# show protocols
rsvp {
    interface fe-2/0/10.7;
    interface lo0.7;
}
mpls {
    interface fe-2/0/10.7;
    interface lo0.7;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/1.0;
        interface fe-2/0/10.7;
        interface lo0.7;
    }
}
```

```
user@PE3# show routing-options
router-id 100.70.70.70;
```

Device PE4

```
user@PE4# show interfaces
    ge-2/0/0 {
        unit 0 {
            description PE4-to-CE4;
            family inet {
                address 10.0.104.9/30;
            }
        }
    }
    fe-2/0/10 {
        unit 4 {
            description PE4-to-P4;
            family inet {
                address 4.4.4.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 4 {
```

```
          family inet {
            address 100.40.40.40/32;
          }
        }
      }
    }

    user@PE4# show protocols
    rsvp {
      interface fe-2/0/10.4;
      interface lo0.4;
    }
    mpls {
      interface fe-2/0/10.4;
      interface lo0.4;
    }
    ospf {
      traffic-engineering;
      area 0.0.0.0 {
        interface ge-2/0/0.0;
        interface fe-2/0/10.4;
        interface lo0.4;
      }
    }

    user@PE4# show routing-options
    router-id 100.40.40.40;
```

Configuring Device CE1

Step-by-Step Procedure

To configure Device CE1:

1. Configure an interface to Device PE1.

   ```
   [edit interfaces]
   user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
   user@CE1# set ge-1/3/2 unit 0 description CE1-to-PE1
   ```

2. Configure static routes from Device CE1 to the three other customer networks, with Device PE1 as the next hop.

   ```
   [edit routing-options]
   user@CE1# set static route 10.0.104.8/30 next-hop 10.0.244.10
   user@CE1# set static route 10.0.134.8/30 next-hop 10.0.244.10
   user@CE1# set static route 10.0.224.8/30 next-hop 10.0.244.10
   ```

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

   ```
   [edit]
   user@CE1# commit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** 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
ge-1/3/2 {
```



```
    unit 0 {
        family inet {
            address 10.0.244.9/30;
            description CE1-to-PE1;
        }
    }
}
```

```
user@CE1# show routing-options
static {
    route 10.0.104.8/30 next-hop 10.0.244.10;
    route 10.0.134.8/30 next-hop 10.0.244.10;
    route 10.0.224.8/30 next-hop 10.0.244.10;
}
```

Configuring Device CE2

Step-by-Step Procedure

To configure Device CE2:

1. Configure an interface to Device PE2.

   ```
   [edit interfaces]
   user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
   user@CE2# set ge-1/3/3 unit 0 description CE2-to-PE2
   ```

2. Configure a static route from Device CE2 to CE1, with Device PE2 as the next hop.

   ```
   [edit routing-options]
   user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
   ```

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

   ```
   [edit]
   user@CE2# commit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** 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@CE2# show interfaces
ge-1/3/3 {
    unit 0 {
        family inet {
            address 10.0.224.9/30;
            description CE2-to-PE2;
        }
    }
}
```

```
user@CE2# show routing-options
static {
    route 10.0.244.8/30 next-hop 10.0.224.10;
}
```

Configuring Device CE3

Step-by-Step Procedure

To configure Device CE3:

1. Configure an interface to Device PE3.

```
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-PE3
```

2. Configure a static route from Device CE3 to CE1, with Device PE3 as the next hop.

```
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
```

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

```
[edit]
user@CE3# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** 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@CE3# show interfaces
ge-2/0/1 {
  unit 0 {
    family inet {
      address 10.0.134.9/30;
      description CE3-to-PE3;
    }
  }
}
```

```
user@CE3# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.134.10;
}
```

Configuring Device CE4

Step-by-Step Procedure

To configure Device CE4:

1. Configure an interface to Device PE4.

```
[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-PE4
```

2. Configure a static route from Device CE4 to CE1, with Device PE4 as the next hop.

```
[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9
```

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

```
[edit]
user@CE4# commit
```


Results From configuration mode, confirm your configuration by entering the **show interfaces** 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@CE4# show interfaces
ge-3/1/3 {
  unit 0 {
    family inet {
      address 10.0.104.10/30;
      description CE4-to-PE4;
    }
  }
}
```

```
user@CE4# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.104.9;
}
```

Verification

Confirm that the configuration is working properly.

- Verifying Connectivity on page 69
- Verifying the State of the Point-to-Multipoint LSP on page 70
- Checking the Forwarding Table on page 70

Verifying Connectivity

Purpose Make sure that the devices can ping each other.

Action Run the **ping** command from CE1 to the interface on CE2 connecting to PE2.

```
user@CE1> ping 10.0.224.9
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to PE3.

```
user@CE1> ping 10.0.134.9
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to PE4.

```
user@CE1> ping 10.0.104.10
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

Verifying the State of the Point-to-Multipoint LSP

Purpose Make sure that the ingress, transit, and egress LSRs are in the Up state.

Action Run the **show mpls lsp p2mp** command on all of the LSRs. Only the ingress LSR is shown here.

```
user@PE1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To              From            State Rt P     ActivePath       LSPname
100.40.40.40    100.10.10.10    Up     0 *                      PE1-PE4
100.70.70.70    100.10.10.10    Up     0 *                      PE1-PE3
100.50.50.50    100.10.10.10    Up     0 *                      PE1-PE2
Total 3 displayed, Up 3, Down 0
...
```

Checking the Forwarding Table

Purpose Make sure that the routes are set up as expected by running the **show route forwarding-table** command. Only the routes to the remote customer networks are shown here.

Action

```
user@PE1> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
...
10.0.104.8/30      user     0 3.3.3.2            ucst  1006     6 fe-2/0/8.9
10.0.134.8/30      user     0 6.6.6.2            ucst  1010     6 fe-2/0/9.8
10.0.224.8/30      user     0 2.2.2.2            ucst  1008     6 fe-2/0/10.1
...
```

Related Documentation

- *(OBSOLETE) RSVP Signaling Protocol*
- *Example: Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems*

Example: Configuring Static Routes for CLNS

Understanding Static Routes for CLNS

The Connectionless Network Service (CLNS) is an ISO Layer 3 protocol that uses network service access point (NSAP) reachability information instead of IPv4 or IPv6 prefixes.

You can configure static routes to exchange CLNS routes within a CLNS island. A *CLNS island* is typically an IS-IS level 1 area that is part of a single IGP routing domain. An island can contain more than one area. CLNS islands can be connected by VPNs.

Example: Configuring Static Routes for CLNS

This example shows how to configure static routes for CLNS.

Requirements

Before you begin, configure the network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

In this example, you configure static routes for CLNS. In the absence of an interior gateway protocol (IGP) on a certain link, a routing device might need to be configured with static routes for CLNS prefixes to be reachable by way of that link. This might be useful, for example, at an autonomous system (AS) boundary.

When you configure static routes for CLNS, consider the following tasks:

- Specify the **iso.0** routing table option to configure a primary instance CLNS static route.
- Specify the **instance-name.iso.0** routing table option to configure a CLNS static route for a particular routing instance.
- Specify the **route** ***nsap-prefix*** statement to configure the destination for the CLNS static route.
- Specify the **next-hop (*interface-name* | *iso-net*)** statement to configure the next hop, specified as an ISO network entity title (NET) or interface name.
- Include the **qualified-next-hop (*interface-name* | *iso-net*)** statement to configure a secondary backup next hop, specified as an ISO network entity title or interface name.

Configuration

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.

```
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
   47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
set routing-options rib iso.0 static iso-route
   47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.eee0/152
   qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference
   20
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.eee0/152
   qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

Step-by-Step Procedure

To configure static routes for CLNS:

1. Configure the routes.

```
[edit routing-options rib iso.0 static]
user@host# set iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
   47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
user@host# set iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152
   next-hop t1-0/2/2.0
user@host# set iso-route 47.0005.80ff.f800.0000.eee0/152 qualified-next-hop
   47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference 20
user@host# set iso-route 47.0005.80ff.f800.0000.eee0/152 qualified-next-hop
   47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

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

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

Results

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


```
user@host# show routing-options
rib iso.0 {
  static {
    iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
      47.0005.80ff.f800.0000.0108.0001.1921.6800.4212;
    iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0;
    iso-route 47.0005.80ff.f800.0000.eee0/152 {
      qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 {
        preference 20;
        metric 10;
      }
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

Checking the Routing Table

Purpose Make sure that the expected routes appear in the routing table.

Action

```
user@host> show route table iso.0

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

47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152
                   *[Static/5] 00:00:25
                    > via t1-0/2/2.0
47.0005.80ff.f800.0000.eee0/84
                   *[Static/20] 00:04:01, metric 10, metric2 10
                    > to #75 0.12.0.34.0.56 via fe-0/0/1.0
47.0005.80ff.f800.0000.ffff.ffff/104
                   *[Static/5] 00:04:01, metric2 0
                    > via t1-0/2/2.0
```

Meaning The static routes appear in the routing table.

CHAPTER 3

Aggregate and Generated Routes

Example: Summarizing Static Routes Through Route Aggregation

Understanding Aggregate Routes

Route aggregation allows you to combine groups of routes with common addresses into a single entry in the routing table. This decreases the size of the routing table as well as the number of route advertisements sent by the routing device.

An aggregate route becomes active when it has one or more *contributing routes*. A contributing route is an active route that is a more specific match for the aggregate destination. For example, for the aggregate destination **128.100.0.0/16**, routes to **128.100.192.0/19** and **128.100.67.0/24** are contributing routes, but routes to **128.0.0.0./8** and **128.0.0.0/16** are not.

A route can contribute only to a single aggregate route. However, an active aggregate route can recursively contribute to a less specific matching aggregate route. For example, an aggregate route to the destination **128.100.0.0/16** can contribute to an aggregate route to **128.96.0.0/13**.

When an aggregate route becomes active, it is installed in the routing table with the following information:

- Reject next hop—If a more-specific packet does not match a more-specific route, the packet is rejected and an ICMP unreachable message is sent to the packet's originator.
- Metric value as configured with the **aggregate** statement.
- Preference value that results from the policy filter on the primary contributor, if a filter is specified.
- AS path as configured in the **aggregate** statement, if any. Otherwise, the path is computed by aggregating the paths of all contributing routes.
- Community as configured in the **aggregate** statement, if any is specified.

NOTE: You can configure only one aggregate route for each destination prefix.

To configure aggregate routes in the default routing table (**inet.0**), include the **aggregate** statement:

```
aggregate {
  defaults {
    ... aggregate-options ...
  }
  route destination-prefix {
    policy policy-name;
    ... aggregate-options ...
  }
}
```

To configure aggregate routes in one of the other routing tables, or to explicitly configure aggregate routes in the default routing table (**inet.0**), include the **aggregate** statement:

```
rib routing-table-name {
  aggregate {
    defaults {
      ... aggregate-options ...
    }
    route destination-prefix {
      policy policy-name;
      ... aggregate-options ...
    }
  }
}
```

NOTE: You cannot configure aggregate routes for the IPv4 multicast routing table (inet.1) nor the IPv6 multicast routing table (inet6.1).

The **aggregate** statement consists of two parts:

- **defaults**—(Optional) Here you specify global aggregate route options. These are treated as global defaults and apply to all the aggregate routes you configure in the **aggregate** statement.
- **route**—Here you configure individual aggregate routes. In this part of the **aggregate** statement, you optionally can configure aggregate route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the **aggregate** statement.

When you configure an individual aggregate route in the **route** part of the **aggregate** statement, specify the destination of the route (in **route *destination-prefix***) in one of the following ways:

- ***network/mask-length***, where ***network*** is the network portion of the IP address and ***mask-length*** is the destination prefix length.
- **default** if this is the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.

After you have configured aggregate routes, you can have a protocol advertise the routes by configuring a policy that is then exported by a routing protocol.

You can associate a routing policy when configuring an aggregate route's destination prefix in the **routes** part of the **aggregate** statement. Doing so provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route and, if the contributor is accepted, the policy can modify the default preferences.

The following algorithm is used to compare two aggregate contributing routes in order to determine which one is the primary or preferred contributor:

1. Compare the protocol's **preferences** of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.
2. Compare the protocol's **preferences2** of the contributing routes. The lower preference2 value is better. If only one route has **preferences2**, then this route is preferred.
3. The preference values are the same. Proceed with a numerical comparison of the prefix values.
   a. The primary contributor is the numerically smallest prefix value.
   b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.
4. At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.

To associate a routing policy with an aggregate route, include the **policy** statement when configuring the route:

```
aggregate (defaults | route) {
  policy policy-name;
}
```

In the **defaults** and **route** parts of the **aggregate** statement, you can specify ***aggregate-options***, which define additional information about aggregate routes that is included with the route when it is installed in the routing table. All aggregate options are optional. Aggregate options that you specify in the **defaults** part of the **aggregate** statement are treated as global defaults and apply to all the aggregate routes you configure in the **aggregate** statement. Aggregate options that you specify in the **route** part of the **aggregate** statement override any global aggregate options and apply to that destination only.

To configure aggregate route options, include one or more of them in the **defaults** or **route** part of the **aggregate** statement:

```
[edit]
routing-options {
  aggregate {
    (defaults | route) {
      (active | passive);
      as-path <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator
        as-number in-address>;
      community [ community-ids ];
      discard;
      (brief | full);
      (metric | metric2 | metric3 | metric4) metric <type type>;
      (preference | preference2 | color | color2) preference <type type>;
      tag string;
    }
  }
}
```

Configuring a Metric Value for Aggregate Routes

You can specify up to four metric values, starting with **metric** (for the first metric value) and continuing with **metric2**, **metric3**, and **metric4** by including one or more of the following statements:

```
aggregate (defaults | route) {
  (metric | metric2 | metric3 | metric4) metric <type type>;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

In the **type** option, you can specify the type of route.

Configuring a Preference Value for Aggregate Routes

By default, aggregate routes have a preference value of 130. If the routing table contains a dynamic route to a destination that has a better (lower) preference value than this, the dynamic route is chosen as the active route and is installed in the forwarding table.

To modify the default preference value, specify a primary preference value (**preference**). You also can specify secondary preference value (**preference2**); and colors, which are even finer-grained preference values (**color** and **color2**). To do this, include one or more of the following statements:

```
aggregate (defaults | route) {
  (preference | preference2 | color | color2) preference <type type>;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

The preference value can be a number in the range from 0 through 4,294,967,295 ($2^{32} - 1$) with a lower number indicating a more preferred route. For more information about preference values, see *Route Preferences Overview*.

In the **type** option, you can specify the type of route.

Configuring the Next Hop for Aggregate Routes

By default, when aggregate routes are installed in the routing table, the next hop is configured as a reject route. That is, the packet is rejected and an ICMP unreachable message is sent to the packet's originator.

When you configure an individual route in the **route** part of the **aggregate** statement, or when you configure the defaults for aggregate routes, you can specify a discard next hop. This means that if a more specific packet does not match a more specific route, the packet is rejected and a reject route for this destination is installed in the routing table, but ICMP unreachable messages are not sent.

Being able to discard next hops allows you to originate a summary route, which can be advertised through dynamic routing protocols, and allows you to discard received traffic that does not match a more specific route than the summary route. To discard next hops, include the **discard** option:

```
discard;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Associating BGP Communities with Aggregate Routes

By default, no BGP community information is associated with aggregate routes. To associate community information with the routes, include the **community** option:

```
aggregate (defaults | route) {
  community [ community-ids ];
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement. ***community-value*** is the community identifier and can be a number in the range from 0 through 65,535.

community-ids is one or more community identifiers for either communities or extended communities.

The format for community identifiers is:

as-number:*community-value*

as-number is the AS number and can be a value in the range from 1 through 65,534.

You also can specify ***community-ids*** for communities as one of the following well-known community names, which are defined in RFC 1997:

- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.

- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

You can explicitly exclude BGP community information with an aggregate route using the **none** option. Include **none** when configuring an individual route in the **route** portion of the **aggregate** statement to override a **community** option specified in the **defaults** portion of the statement.

NOTE: Extended community attributes are not supported at the **[edit routing-options]** hierarchy level. You must configure extended communities at the **[edit policy-options]** hierarchy level. For information about configuring extended communities information, see the "Configuring the Extended Communities Attribute" section in the *Routing Policy Feature Guide for Routing Devices*. For information about configuring 4-byte AS numbers and extended communities, see *Using 4-Byte Autonomous System Numbers in BGP Networks*.

Associating AS Paths with Aggregate Routes

By default, the AS path for aggregate routes is built from the component routes. To manually specify the AS path and associate AS path information with the routes, include the **as-path** option:

```
aggregate (defaults | route) {
  as-path <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator
    as-number in-address>;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

as-path is the AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([ ]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path.

NOTE: In Junos OS Release 9.1 and later, the numeric AS range is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. For the AS number, you can configure a value from 1 through 4,294,967,295. All releases of Junos OS support 2-byte AS numbers. The 2-byte AS number range is 1 through 65,535 (this is a subset of the 4-byte range).

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>*.*<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.

You also can specify the AS path using the BGP origin attribute, which indicates the origin of the AS path information:

- **egp**—Path information originated in another AS.
- **igp**—Path information originated within the local AS.
- **incomplete**—Path information was learned by some other means.

To attach the BGP **ATOMIC_AGGREGATE** path attribute to the aggregate route, specify the **atomic-aggregate** option. This path attribute indicates that the local system selected a less specific route rather than a more specific route.

To attach the BGP **AGGREGATOR** path attribute to the aggregate route, specify the **aggregator** option. When using this option, you must specify the last AS number that formed the aggregate route (encoded as two octets), followed by the IP address of the BGP system that formed the aggregate route.

Including AS Numbers in Aggregate Route Paths

By default, all AS numbers from all contributing paths are included in the aggregate route's path. To include only the longest common leading sequences from the contributing AS paths, include the **brief** option when configuring the route. If doing this results in AS numbers being omitted from the aggregate route, the BGP **ATOMIC_ATTRIBUTE** path attribute is included with the aggregate route.

```
aggregate (defaults | route) {
  brief;
}
```

To explicitly have all AS numbers from all contributing paths be included in the aggregate route's path, include the **full** option when configuring routes. Include this option when configuring an individual route in the **route** portion of the **aggregate** statement to override a **retain** option specified in the **defaults** portion of the statement.

```
aggregate (defaults | route) {
  full;
}
```


For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Configuring an OSPF Tag String for Aggregate Routes

By default, no OSPF tag strings are associated with aggregate routes. You can specify an OSPF tag string by including the **tag** option:

```
aggregate (defaults | route) {
  tag string;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Controlling Retention of Inactive Aggregate Routes in the Routing and Forwarding Tables

Static routes are only removed from the routing table if the next hop becomes unreachable, which happens if there are no contributing routes. To have an aggregate route remain continually installed in the routing and forwarding tables, include the **passive** option when configuring the route:

```
aggregate (defaults | route) {
  passive;
}
```

Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with **reject** next hops when they are inactive.

To explicitly remove aggregate routes when they become inactive, include the **active** option when configuring routes. Include this option when configuring an individual route in the **route** portion of the **aggregate** statement to override a **passive** option specified in the **defaults** portion of the statement.

```
aggregate (defaults | route) {
  active;
}
```

Example: Summarizing Routes Through Route Aggregation

This example shows how to summarize routes by configuring aggregate routes.

Requirements

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

Overview

In this example, Device R1 is connected to customer networks 10.200.1.0/24 and 10.200.2.0/24. For demonstration purposes, these routes are represented in this example as loopback interfaces on Device R1.

Device R2 has static routes configured to reach Device R1's customer networks. Device R2 also has a routing policy configured to advertise all static routes to its neighbors in autonomous system (AS) 65001.

Device R3 is in AS 65001 and receives the static routes from Device R2. When Device R3 sends information about these routes to Device ISP, the information is summarized as a single aggregate route. The aggregate route is 10.200.0.0/16.

Device ISP injects a default route into AS 65001, and Device R3 advertises the default route.

This example shows the configuration for all of the devices and the step-by-step configuration on Device R3.

Figure 8 on page 83 shows the sample network.

Figure 8: Aggregate Route Advertised to an ISP

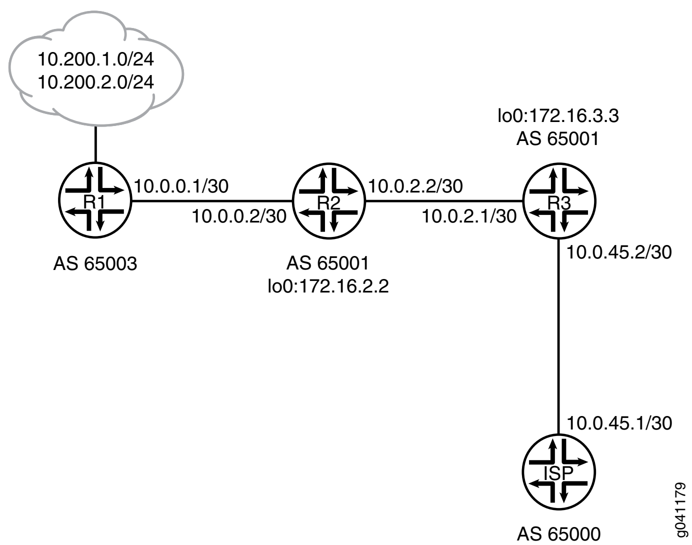

Configuration

CLI Quick Configuration

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

Device R1

```
set interfaces ge-1/2/0 unit 2 description R1->R2
set interfaces ge-1/2/0 unit 2 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 10.200.1.1/32
set interfaces lo0 unit 1 family inet address 10.200.2.2/32
```

```
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65001
set protocols bgp group ext neighbor 10.0.0.2
set protocols ospf area 0.0.0.0 interface ge-1/2/0.2
set routing-options autonomous-system 65003
```

Device R2

```
set interfaces ge-1/2/0 unit 1 description R2->R1
set interfaces ge-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 4 description R2->R3
set interfaces ge-1/2/1 unit 4 family inet address 10.0.2.2/30
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols bgp group int type internal
set protocols bgp group int local-address 172.16.2.2
set protocols bgp group int export send-customer-routes
set protocols bgp group int neighbor 172.16.3.3
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65003
set protocols bgp group ext neighbor 10.0.0.1
set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
set protocols ospf area 0.0.0.0 interface ge-1/2/1.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set policy-options policy-statement send-customer-routes from protocol static
set policy-options policy-statement send-customer-routes then accept
set routing-options static route 10.200.1.0/24 next-hop 10.0.0.1
set routing-options static route 10.200.2.0/24 next-hop 10.0.0.1
set routing-options autonomous-system 65001
```

Device R3

```
set interfaces ge-1/2/0 unit 3 description R3->R2
set interfaces ge-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces ge-1/2/1 unit 6 description R3->ISP
set interfaces ge-1/2/1 unit 6 family inet address 10.0.45.2/30
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-aggregate
set protocols bgp group ext peer-as 65000
set protocols bgp group ext neighbor 10.0.45.1
set protocols bgp group int type internal
set protocols bgp group int local-address 172.16.3.3
set protocols bgp group int neighbor 172.16.2.2
set protocols ospf export send-default
set protocols ospf area 0.0.0.0 interface ge-1/2/0.3
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement send-aggregate term 1 from protocol aggregate
set policy-options policy-statement send-aggregate term 1 then accept
set policy-options policy-statement send-aggregate term suppress-specific-routes from
  route-filter 10.200.0.0/16 longer
set policy-options policy-statement send-aggregate term suppress-specific-routes then
  reject
set policy-options policy-statement send-default from route-filter 0.0.0.0/0 exact
set policy-options policy-statement send-default then accept
set routing-options aggregate route 10.200.0.0/16
set routing-options autonomous-system 65001
```

Device ISP

```
set interfaces ge-1/2/0 unit 7 family inet address 10.0.45.1/30
set protocols bgp group ext type external
```



```
set protocols bgp group ext export advertise-default
set protocols bgp group ext peer-as 65001
set protocols bgp group ext neighbor 10.0.45.2
set policy-options policy-statement advertise-default term 1 from route-filter 0.0.0.0/0
   exact
set policy-options policy-statement advertise-default term 1 then accept
set routing-options static route 0.0.0.0/0 discard
set routing-options autonomous-system 65000
```

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 R3:

1. Configure the interfaces.

```
[edit interfaces]
user@R3# set ge-1/2/0 unit 3 description R3->R2
user@R3# set ge-1/2/0 unit 3 family inet address 10.0.2.1/30

user@R3# set ge-1/2/1 unit 6 description R3->ISP
user@R3# set ge-1/2/1 unit 6 family inet address 10.0.45.2/30

user@R3# set lo0 unit 3 family inet address 172.16.3.3/32
```

2. Configure the AS number.

```
[edit routing-options]
user@R3# set autonomous-system 65001
```

3. Configure the BGP session with the ISP device.

```
[edit protocols bgp group ext]
user@R3# set type external
user@R3# set peer-as 65000
user@R3# set neighbor 10.0.45.1
```

4. Configure the BGP session with Device R2.

```
[edit protocols bgp group int]
user@R3# set type internal
user@R3# set local-address 172.16.3.3
user@R3# set neighbor 172.16.2.2
```

5. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface ge-1/2/0.3
user@R3# set interface lo0.3 passive
```

6. Configure the aggregate route.

```
[edit routing-options]
user@R3# set aggregate route 10.200.0.0/16
```

7. Configure the routing policy to advertise the aggregate route.

   The first term in this policy advertises the aggregate route. The second term prevents more specific routes from being advertised.

   ```
   [edit policy-options policy-statement send-aggregate]
   user@R3# set term 1 from protocol aggregate
   user@R3# set term 1 then accept
   user@R3# set term suppress-specific-routes from route-filter 10.200.0.0/16 longer
   user@R3# set term suppress-specific-routes then reject
   ```

8. Apply the aggregate route policy to the external BGP session with Device ISP.

   ```
   [edit protocols bgp group ext]
   user@R3# set export send-aggregate
   ```

9. Configure the routing policy to advertise the default route from Device ISP.

   ```
   [edit policy-options policy-statement send-default]
   user@R3# set from route-filter 0.0.0.0/0 exact
   user@R3# set then accept
   ```

10. Apply the default routing policy to OSPF.

    ```
    [edit protocols ospf]
    user@R3# set export send-default
    ```

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

    ```
    [edit]
    user@R3# commit
    ```

Results

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

```
user@R3# show interfaces
ge-1/2/0 {
   unit 3 {
      description R3->R2;
      family inet {
         address 10.0.2.1/30;
      }
   }
}
ge-1/2/1 {
   unit 6 {
      description R3->ISP;
      family inet {
         address 10.0.45.2/30;
      }
   }
}
lo0 {
   unit 3 {
      family inet {
```



```
      address 172.16.3.3/32;
    }
  }
}
user@R3# show protocols
bgp {
  group ext {
    type external;
    export send-aggregate;
    peer-as 65000;
    neighbor 10.0.45.1;
  }
  group int {
    type internal;
    local-address 172.16.3.3;
    neighbor 172.16.2.2;
  }
}
ospf {
  export send-default;
  area 0.0.0.0 {
    interface ge-1/2/0.3;
    interface lo0.3 {
      passive;
    }
  }
}
user@R3# show policy-options
policy-statement send-aggregate {
  term 1 {
    from protocol aggregate;
    then accept;
  }
  term suppress-specific-routes {
    from {
      route-filter 10.200.0.0/16 longer;
    }
    then reject;
  }
}
policy-statement send-default {
  from {
    route-filter 0.0.0.0/0 exact;
  }
  then accept;
}
user@R3# show routing-options
aggregate {
  route 10.200.0.0/16;
}
autonomous-system 65001;
```

Verification

Confirm that the configuration is working properly.

- Verifying That Device R3 Has the Expected Routes on page 88
- Verifying That Device R3 Advertises the Aggregate Route to Device ISP on page 88

Verifying That Device R3 Has the Expected Routes

Purpose Make sure that Device R3 has the specific static routes.

Action

```
user@R3>show route terse protocol bgp

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 0.0.0.0/0          B 170        100            >10.0.45.1      65000 I
* 10.200.1.0/24      B 170        100            >10.0.2.2       I
* 10.200.2.0/24      B 170        100            >10.0.2.2       I
```

Meaning The output shows that Device R3 has the specific routes to the 10.200.1.0/24 and 10.200.2.0/24 networks.

Verifying That Device R3 Advertises the Aggregate Route to Device ISP

Purpose Make sure that Device R3 does not send the specific static routes and only sends the summarized aggregate route.

Action

```
user@R3>show route advertising-protocol bgp 10.0.45.1
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.200.0.0/16           Self                                    I
```

Meaning The output shows that Device R3 sends only the summarized route to Device ISP.

Related Documentation

- Example: Conditionally Generating Static Routes on page 88

Example: Conditionally Generating Static Routes

- Understanding Conditionally Generated Routes on page 88
- Example: Configuring a Conditional Default Route Policy on page 90

Understanding Conditionally Generated Routes

Generated routes are used as the *route of last resort*. A packet is forwarded to the route of last resort when the routing tables have no information about how to reach that packet's destination. One use of route generation is to generate a default route to use if the routing table contains a route from a peer on a neighboring backbone.

A generated route becomes active when it has one or more *contributing routes*. A contributing route is an active route that is a more specific match for the generated destination. For example, for the destination 128.100.0.0/16, routes to 128.100.192.0/19 and 128.100.67.0/24 are contributing routes, but routes to 128.0.0.0./8, 128.0.0.0/16, and 128.100.0.0/16 are not.

A route can contribute only to a single generated route. However, an active generated route can recursively contribute to a less specific matching generated route. For example, a generated route to the destination 128.100.0.0/16 can contribute to a generated route to 128.96.0.0/13.

By default, when generated routes are installed in the routing table, the next hop is chosen from the primary contributing route.

NOTE: You can configure only one generated route for each destination prefix.

To configure generated routes in the default routing table (**inet.0**), include the **generate** statement:

```
generate {
  (defaults | route destination ) {
    (active | passive);
    as-path <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator
      as-number in-address>;
    community [ community-ids ];
    discard;
    (brief | full);
    (metric | metric2 | metric3 | metric4) metric <type type>;
    policy policy-name;
    (preference | preference2 | color | color2) preference <type type>;
    tag string;
  }
}
```

NOTE: You cannot configure generated routes for the IPv4 multicast routing table (**inet.1**) or the IPv6 multicast routing table (**inet6.1**).

The **generate** statement consists of two parts:

- **defaults**—Here you specify global generated route options. These are treated as global defaults and apply to all the generated routes you configure in the **generate** statement. This part of the **generate** statement is optional.
- **route**—Here you configure individual generated routes. In this part of the **generate** statement, you optionally can configure generated route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the **generate** statement.

Example: Configuring a Conditional Default Route Policy

This example shows how to configure a conditional default route on one routing device and redistribute the default route into OSPF.

- Requirements on page 90
- Overview on page 90
- Configuration on page 91
- Verification on page 95

Requirements

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

Overview

In this example, OSPF area 0 contains three routing devices. Device R3 has a BGP session with an external peer, for example, an Internet service provider (ISP).

To propagate a static route into BGP, this example includes the **discard** statement when defining the route. The ISP injects a default static route into BGP, which provides the customer network with a default static route to reach external networks. The static route has a discard next hop. This means that if a packet does not match a more specific route, the packet is rejected and a reject route for this destination is installed in the routing table, but Internet Control Message Protocol (ICMP) unreachable messages are not sent. The discard next hop allows you to originate a summary route, which can be advertised through dynamic routing protocols.

Device R3 exports the default route into OSPF. The route policy on Device R3 is conditional such that if the connection to the ISP goes down, the default route is no longer exported into OSPF because it is no longer active in the routing table. This policy prevents packets from being silently dropped without notification (also known as black holing).

This example shows the configuration for all of the devices and the step-by-step configuration on Device R3.

Figure 9 on page 91 shows the sample network.

Figure 9: OSPF with a Conditional Default Route to an ISP

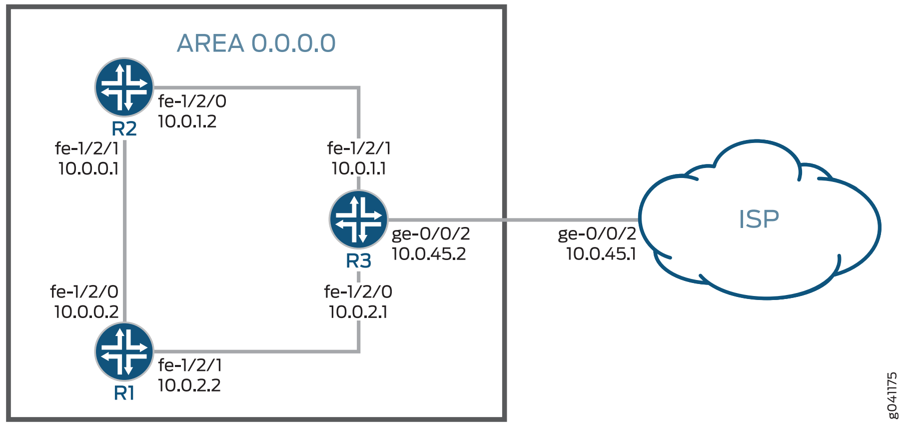

Configuration

CLI Quick Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 0 description R1->R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.1.2/30
set interfaces fe-1/2/1 unit 2 description R1->R2
set interfaces fe-1/2/1 unit 2 family inet address 10.0.0.1/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
```

Device R2

```
set interfaces fe-1/2/0 unit 1 description R2->R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 4 description R2->R3
set interfaces fe-1/2/1 unit 4 family inet address 10.0.2.2/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
```

Device R3

```
set interfaces fe-1/2/0 unit 3 description R3->R2
set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces fe-1/2/1 unit 5 description R3->R1
set interfaces fe-1/2/1 unit 5 family inet address 10.0.1.1/30
set interfaces ge-0/0/2 unit 0 description R3->ISP
set interfaces ge-0/0/2 unit 0 family inet address 10.0.45.2/30
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65000
set protocols bgp group ext neighbor 10.0.45.1
set protocols ospf export gendefault
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set policy-options policy-statement gendefault term upstreamroutes from protocol bgp
set policy-options policy-statement gendefault term upstreamroutes from as-path
  upstream
set policy-options policy-statement gendefault term upstreamroutes from route-filter
  0.0.0.0/0 upto /16
```

```
set policy-options policy-statement gendefault term upstreamroutes then next-hop
  10.0.45.1
set policy-options policy-statement gendefault term upstreamroutes then accept
set policy-options policy-statement gendefault term end then reject
set policy-options as-path upstream "^65000 "
set routing-options generate route 0.0.0.0/0 policy gendefault
set routing-options autonomous-system 65001
```

Device ISP

```
set interfaces ge-0/0/2 unit 0 family inet address 10.0.45.1/30
set protocols bgp group ext type external
set protocols bgp group ext export advertise-default
set protocols bgp group ext peer-as 65001
set protocols bgp group ext neighbor 10.0.45.2
set policy-options policy-statement advertise-default term 1 from route-filter 0.0.0.0/0
  exact
set policy-options policy-statement advertise-default term 1 then accept
set routing-options static route 0.0.0.0/0 discard
set routing-options autonomous-system 65000
```

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 R3:

1. Configure the device interfaces.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 3 description R3->R2
user@R3# set fe-1/2/0 unit 3 family inet address 10.0.2.1/30

user@R3# set fe-1/2/1 unit 5 description R3->R1
user@R3# set fe-1/2/1 unit 5 family inet address 10.0.1.1/30

user@R3# set ge-0/0/2 unit 0 description R3->ISP
user@R3# set ge-0/0/2 unit 0 family inet address 10.0.45.2/30
```

2. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 65001
```

3. Configure the BGP session with the ISP device.

```
[edit protocols bgp group ext]
user@R3# set type external
user@R3# set peer-as 65000
user@R3# set neighbor 10.0.45.1
```

4. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface fe-1/2/1.4
user@R3# set interface fe-1/2/0.3
```


5. Configure the routing policy.

```
[edit policy-options policy-statement gendefault]
user@R3# set term upstreamroutes from protocol bgp
user@R3# set term upstreamroutes from as-path upstream
user@R3# set term upstreamroutes from route-filter 0.0.0.0/0 upto /16
user@R3# set term upstreamroutes then next-hop 10.0.45.1
user@R3# set term upstreamroutes then accept

user@R3# set term end then reject

[edit policy-options]
user@R3# set as-path upstream "^65000 "
```

6. Configure the generated route, associating the routing policy with the generated route.

```
[edit routing-options]
user@R3# set generate route 0.0.0.0/0 policy gendefault
```

7. Apply the export policy to OSPF.

```
[edit protocols ospf]
user@R3# set export gendefault
```

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

```
[edit]
user@R3# commit
```

Results

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

```
user@R3# show
interfaces {
  fe-1/2/0 {
    unit 3 {
      description R3->R2;
      family inet {
        address 10.0.2.1/30;
      }
    }
  }
  fe-1/2/1 {
    unit 5 {
      description R3->R1;
      family inet {
        address 10.0.1.1/30;
      }
    }
  }
  ge-1/2/0 {
    unit 0 {
```

```
            description R3->ISP;
            family inet {
                address 10.0.45.2/30;
            }
        }
    }
}
protocols {
    bgp {
        group ext {
            type external;
            peer-as 65000;
            neighbor 10.0.45.1;
        }
    }
    ospf {
        export gendefault;
        area 0.0.0.0 {
            interface fe-1/2/1.4;
            interface fe-1/2/0.3;
        }
    }
}
policy-options {
    policy-statement gendefault {
        term upstreamroutes {
            from {
                protocol bgp;
                as-path upstream;
                route-filter 0.0.0.0/0 upto /16;
            }
            then {
                next-hop 10.0.45.1;
                accept;
            }
        }
        term end {
            then reject;
        }
    }
    as-path upstream "^65000 ";
}
routing-options {
    generate {
        route 0.0.0.0/0 policy gendefault;
    }
    autonomous-system 65001;
}
```


Verification

Confirm that the configuration is working properly.

- Verifying That the Route to the ISP Is Working on page 95
- Verifying That the Static Route Is Redistributed on page 95
- Testing the Policy Condition on page 96

Verifying That the Route to the ISP Is Working

Purpose Make sure connectivity is established between Device R3 and the ISP's router.

Action

```
user@R3> ping 10.0.45.1
PING 10.0.45.1 (10.0.45.1): 56 data bytes
64 bytes from 10.0.45.1: icmp_seq=0 ttl=64 time=1.185 ms
64 bytes from 10.0.45.1: icmp_seq=1 ttl=64 time=1.199 ms
64 bytes from 10.0.45.1: icmp_seq=2 ttl=64 time=1.186 ms
```

Meaning The **ping** command confirms reachability.

Verifying That the Static Route Is Redistributed

Purpose Make sure that the BGP policy is redistributing the static route into Device R3's routing table. Also make sure that the OSPF policy is redistributing the static route into the routing tables of Device R1 and Device R2.

Action

user@R3> **show route protocol bgp**

```
inet.0: 9 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:00:25, localpref 100
                      AS path: 65000 I
                    > to 10.0.45.1 via ge-0/0/2.6
```

user@R1> **show route protocol ospf**

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

0.0.0.0/0          *[OSPF/150] 00:03:58, metric 0, tag 0
                    > to 10.0.1.1 via fe-1/2/0.0
10.0.2.0/30        *[OSPF/10] 03:37:45, metric 2
                      to 10.0.1.1 via fe-1/2/0.0
                    > to 10.0.0.2 via fe-1/2/1.2
224.0.0.5/32       *[OSPF/10] 03:38:41, metric 1
                      MultiRecv
```

user@R2> **show route protocol ospf**

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

0.0.0.0/0          *[OSPF/150] 00:04:04, metric 0, tag 0
                    > to 10.0.2.1 via fe-1/2/1.4
10.0.1.0/30        *[OSPF/10] 03:37:46, metric 2
                      to 10.0.0.1 via fe-1/2/0.1
                    > to 10.0.2.1 via fe-1/2/1.4
224.0.0.5/32       *[OSPF/10] 03:38:47, metric 1
                      MultiRecv
```

Meaning

The routing tables contain the default 0.0.0.0/0 route. If Device R1 and Device R2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Device R3 for further processing. If Device R3 receives packets destined for networks not specified in its routing table, those packets will be sent to the ISP for further processing.

Testing the Policy Condition

Purpose

Deactivate the interface to make sure that the route is removed from the routing tables if the external network becomes unreachable.

Action

```
user@R3> deactivate interfaces ge-0/0/2 unit 0 family inet address 10.0.45.2/30
user@R3> commit
```

```
user@R1> show route protocol ospf

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

10.0.2.0/30        *[OSPF/10] 03:41:48, metric 2
                      to 10.0.1.1 via fe-1/2/0.0
                    > to 10.0.0.2 via fe-1/2/1.2
224.0.0.5/32       *[OSPF/10] 03:42:44, metric 1
                      MultiRecv
```

```
user@R2> show route protocol ospf
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/30        *[OSPF/10] 03:42:10, metric 2
                      to 10.0.0.1 via fe-1/2/0.1
                    > to 10.0.2.1 via fe-1/2/1.4
224.0.0.5/32       *[OSPF/10] 03:43:11, metric 1
                      MultiRecv
```

Meaning The routing tables on Device R1 and Device R2 do not contain the default 0.0.0.0/0 route. This verifies that the default route is no longer present in the OSPF domain. To reactivate the ge-0/0/2.6 interface, issue the **activate interfaces ge-0/0/2 unit 0 family inet address 10.0.45.2/30** configuration mode command.

Related Documentation

- Example: Summarizing Static Routes Through Route Aggregation on page 75

CHAPTER 4

Bidirectional Forwarding Detection for Static Routes

Examples: Configuring BFD for Static Routes

Understanding BFD for Static Routes

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes. In Junos OS Release 8.2 and later, BFD also supports multihop static routes.

To enable failure detection, include the bfd-liveness-detection statement in the static route configuration.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is not supported for any other protocol.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route *destination-prefix*]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

NOTE: **If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.**

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

NOTE: SRX Series devices do not support distributed BFD.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.

NOTE: You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.

NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

NOTE: If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Junos OS also supports BFD over multihop static routes. For example, you can configure BFD over a Layer 3 path to provide path integrity over that path. You can limit the number of hops by specifying the time to live (TTL).

To configure BFD over multihop static routes, include the following statements:

```
static route destination-prefix {
  bfd-liveness-detection {
    local-address ip-address;
    minimum-receive-ttl number;
  }
}
```

To specify the source address for the multihop static route and to enable multihop BFD support, include the **local-address** statement.

To specify the number of hops, include the **minimum-receive-ttl** statement. You must configure this statement for a multihop BFD session. You can configure a value in the range from 1 through 255. It is optional for a single-hop BFD session. If you configure the **minimum-receive-ttl** statement for a single-hop session, the value must be 255.

On M Series and T Series platforms only, starting in Junos OS Release 12.3, multihop BFD runs on the CPU in the FPC, DPC, or MPC. This is referred to as *distributed BFD*. Previously, multihop BFD ran from the Routing Engine.

Example: Configuring BFD for Static Routes

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 10 on page 104 shows the sample network.

Figure 10: Customer Routes Connected to a Service Provider

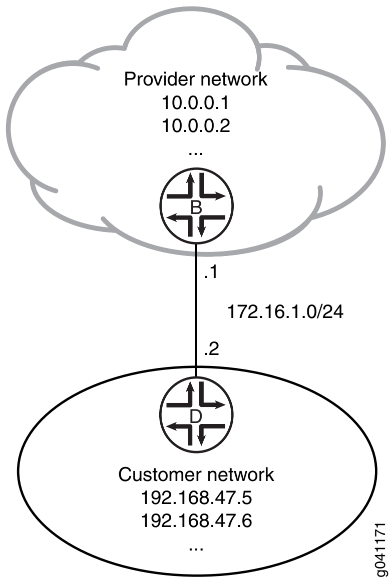

Configuration

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 B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```


Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

   ```
   [edit interfaces]
   user@B# set ge-1/2/0 unit 0 description B->D
   user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
   user@B# set lo0 unit 57 family inet address 10.0.0.1/32
   user@B# set lo0 unit 57 family inet address 10.0.0.2/32
   ```

2. On Device B, create a static route and set the next-hop address.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
   ```

3. On Device B, configure BFD for the static route.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
     1000
   ```

4. On Device B, configure tracing operations for BFD.

   ```
   [edit protocols]
   user@B# set bfd traceoptions file bfd-trace
   user@B# set bfd traceoptions flag all
   ```

5. If you are done configuring Device B, commit the configuration.

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

6. On Device D, configure the interfaces.

   ```
   [edit interfaces]
   user@D# set ge-1/2/0 unit 1 description D->B
   user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
   user@D# set lo0 unit 2 family inet address 192.168.47.5/32
   user@D# set lo0 unit 2 family inet address 192.168.47.6/32
   ```

7. On Device D, create a static route and set the next-hop address.

   ```
   [edit routing-options]
   user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
   ```

8. On Device D, configure BFD for the static route.

   ```
   [edit routing-options]
   user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
   ```

9. On Device D, configure tracing operations for BFD.

   ```
   [edit protocols]
   user@D# set bfd traceoptions file bfd-trace
   user@D# set bfd traceoptions flag all
   ```

10. If you are done configuring Device D, commit the configuration.

   ```
   [edit]
   ```



```
user@D# commit
```

Results

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

Device B

```
user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
      address 10.0.0.2/32;
    }
  }
}

user@D# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    next-hop 172.16.1.2;
    bfd-liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

Device D

```
user@D# show interfaces
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
```



```
      address 192.168.47.5/32;
      address 192.168.47.6/32;
    }
  }
}
```

```
user@D# show routing-options
static {
  route 0.0.0.0/0 {
    next-hop 172.16.1.1;
    bfd-liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

- Verifying That BFD Sessions Are Up on page 107
- Viewing Detailed BFD Events on page 108

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.1.2               Up        lt-1/2/0.0     3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 172
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

```
user@D> show bfd session extensive
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.1.1               Up        lt-1/2/0.1     3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:35
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 170
 Min async interval 1.000, min slow interval 1.000
```

```
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 1, remote discriminator 2
 Echo mode disabled/inactive

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning The **TX interval 1.000, RX interval 1.000** output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55    Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
 78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55    IfIndex (3) len 4: 0
Nov 23 14:26:55    Protocol (1) len 1: BFD
Nov 23 14:26:55    Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
 6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: ppmd_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55    IfIndex (3) len 4: 0
Nov 23 14:26:55    Protocol (1) len 1: BFD
Nov 23 14:26:55    Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
 78 6d 69 74 20 74
```

Meaning BFD messages are being written to the trace file.

Example: Enabling BFD on Qualified Next Hops in Static Routes

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See Figure 11 on page 109.

Figure 11: BFD Enabled on Qualified Next Hops

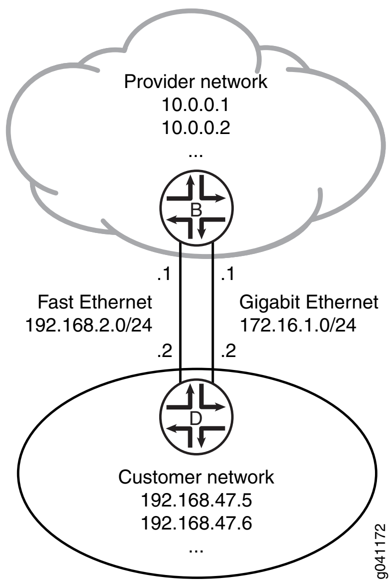

Configuration

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 B

```
set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
  bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
  bfd-liveness-detection minimum-interval 60
```

Device D

```
set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
```

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

   ```
   [edit interfaces fe-0/1/0]
   user@B# set unit 2 description secondary-B->D
   user@B# set unit 2 family inet address 192.168.2.1/24

   [edit interfaces ge-1/2/0]
   user@B# set unit 0 description B->D
   user@B# set unit 0 family inet address 172.16.1.1/24
   ```

2. On Device B, configure the static route with two next hops, both with BFD enabled.

   ```
   [edit routing-options static route 192.168.47.0/24]
   user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
     60
   user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
     60
   ```

3. On Device D, configure the interfaces.

   ```
   [edit interfaces fe-0/1/0]
   user@D# set unit 3 description secondary-D->B
   user@D# set unit 3 family inet address 192.168.2.2/24

   [edit interfaces ge-1/2/0]
   user@D# set unit 1 description D->B
   user@D# set unit 1 family inet address 172.16.1.2/24
   ```

4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

   In this case, BFD is enabled on the route, not on the next hops.

   ```
   [edit routing-options static route 0.0.0.0/0]
   user@D# set qualified-next-hop 192.168.2.1
   user@D# set qualified-next-hop 172.16.1.1
   user@D# set bfd-liveness-detection minimum-interval 60
   ```

Results

Confirm your configuration by issuing the **show interfaces** 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@B# show interfaces
fe-0/1/0 {
  unit 2 {
```



```
        description secondary-B->D;
        family inet {
            address 192.168.2.1/24;
        }
    }
}
ge-1/2/0 {
    unit 0 {
        description B->D;
        family inet {
            address 172.16.1.1/24;
        }
    }
}
```

```
user@B# show routing-options
static {
    route 192.168.47.0/24 {
        qualified-next-hop 192.168.2.2 {
            bfd-liveness-detection {
                minimum-interval 60;
            }
        }
        qualified-next-hop 172.16.1.2 {
            bfd-liveness-detection {
                minimum-interval 60;
            }
        }
    }
}
```

```
user@D# show interfaces
fe-0/1/0 {
    unit 3 {
        description secondary-D->B;
        family inet {
            address 192.168.2.2/24;
        }
    }
}
ge-1/2/0 {
    unit 1 {
        description D->B;
        family inet {
            address 172.16.1.2/24;
        }
    }
}
```

```
user@D# show routing-options
static {
    route 0.0.0.0/0 {
        qualified-next-hop 192.168.2.1;
        qualified-next-hop 172.16.1.1;
        bfd-liveness-detection {
            minimum-interval 60;
        }
```

```
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

- Checking the Routing Tables on page 112
- Verifying the BFD Sessions on page 112
- Removing BFD from Device D on page 113
- Removing BFD from One Next Hop on page 113

Checking the Routing Tables

Purpose Make sure that the static route appears in the routing table on Device B with two possible next hops.

Action

```
user@B> show route 192.168.47.0 extensive
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}
        *Static Preference: 5
                Next hop type: Router
                Address: 0x9334010
                Next-hop reference count: 1
                Next hop: 172.16.1.2 via ge-1/2/0.0
                Next hop: 192.168.2.2 via fe-0/1/0.2, selected
                State: <Active Int Ext>
                Age: 9
                Task: RT
                Announcement bits (1): 3-KRT
                AS path: I
```

Meaning Both next hops are listed. The next hop 192.168.2.2 is the selected route.

Verifying the BFD Sessions

Purpose Make sure that the BFD sessions are up.

Action

```
user@B> show bfd session

                                          Detect   Transmit
Address           State     Interface     Time     Interval  Multiplier
172.16.1.2        Up        ge-1/2/0.0    0.720     0.240        3
192.168.2.2       Up        fe-0/1/0.2    0.720     0.240        3

2 sessions, 2 clients
Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps
```

Meaning The output shows that the BFD sessions are up.

Removing BFD from Device D

Purpose Demonstrate what happens when the BFD session is down for both next hops.

Action

1. Deactivate BFD on Device D.

   [edit routing-options static route 0.0.0.0/0]
   user@D# **deactivate bfd-liveness-detection**
   user@D# **commit**

2. Rerun the **show bfd session** command on Device B.

   ```
   user@B> show bfd session

                                                   Detect   Transmit
   Address                 State     Interface      Time     Interval  Multiplier
   172.16.1.2               Down       ge-1/2/0.0     3.000     1.000        3
   192.168.2.2              Down       fe-0/1/0.2     3.000     1.000        3

   2 sessions, 2 clients
   Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps
   ```

3. Rerun the **show route 192.168.47.0** command on Device B.

   ```
   user@B> show route 192.168.47.0
   ```

Meaning As expected, when the BFD sessions are down, the static route is removed from the routing table.

Removing BFD from One Next Hop

Purpose Demonstrate what happens when only one next hop has BFD enabled.

Action

1. If it is not already deactivated, deactivate BFD on Device D.

   [edit routing-options static route 0.0.0.0/0]
   user@D# **deactivate bfd-liveness-detection**
   user@D# **commit**

2. Deactivate BFD on one of the next hops on Device B.

   [edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
   user@B# **deactivate bfd-liveness-detection**
   user@B# **commit**

3. Rerun the **show bfd session** command on Device B.

   ```
   user@B> show bfd session

                                        Detect   Transmit
   Address      State     Interface      Time     Interval  Multiplier
   192.168.2.2   Down      fe-0/1/0.2     3.000     1.000        3
   ```

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

   ```
   user@B> show route 192.168.47.0 extensive

   inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
   192.168.47.0/24 (1 entry, 1 announced)
   TSI:
   ```

```
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
        *Static Preference: 5
                Next hop type: Router, Next hop index: 624
                Address: 0x92f0178
                Next-hop reference count: 3
                Next hop: 172.16.1.2 via ge-1/2/0.0, selected
                State: <Active Int Ext>
                Age: 2:36
                Task: RT
                Announcement bits (1): 3-KRT
                AS path: I
```

Meaning As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

Related Documentation

- Example: Configuring BFD Authentication for Static Routes on page 114
- *Example: Configuring BFD for OSPF*
- *Example: Configuring BFD for BGP*
- *Example: Configuring BFD for IS-IS*
- *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*

Example: Configuring BFD Authentication for Static Routes

Understanding BFD Authentication for Static Routes

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.

NOTE: **We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.**

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.

NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for Static Routes

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

Requirements

Junos OS Release 9.6 or later (Canda and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Overview

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.

3. Configure the related security authentication keychain. This must be configured on the main router.

> TIP: We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.
>
> [edit]
> user@host> **set routing-options static route ipv4 bfd-liveness-detection authentication loose-check**

Figure 12 on page 117 shows the sample network.

Figure 12: Customer Routes Connected to a Service Provider

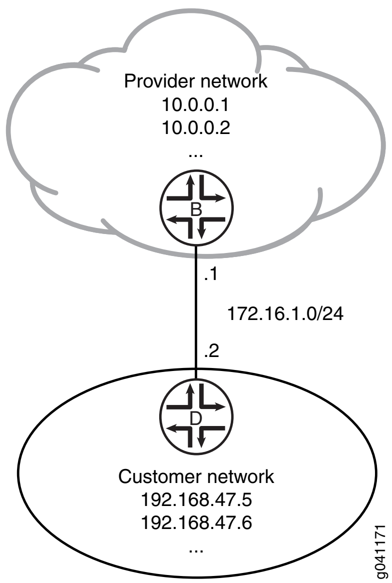

Configuration

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 B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
  key-chain bfd-kc4
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1
```


set security authentication-key-chains key-chain bfd-kc4 key 5 secret "9JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2aDkP5Ft0IQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time "2011-1-1.12:00:00 -0800"

Device D

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret "9JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2aDkP5Ft0IQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time "2011-1-1.12:00:00 -0800"

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

   [edit interfaces]
   user@B# **set ge-1/2/0 unit 0 description B->D**
   user@B# **set ge-1/2/0 unit 0 family inet address 172.16.1.1/24**

   user@B# **set lo0 unit 57 family inet address 10.0.0.1/32**
   user@B# **set lo0 unit 57 family inet address 10.0.0.2/32**

2. On Device B, create a static route and set the next-hop address.

   [edit routing-options]
   user@B# **set static route 192.168.47.0/24 next-hop 172.16.1.2**

3. On Device B, configure BFD for the static route.

   [edit routing-options]
   user@B# **set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000**

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

   [edit routing-options]
   user@B# **set static route 192.168.47.0/24 bfd-liveness-detection authentication algorithm keyed-sha-1**

NOTE: Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

5. On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

   This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
     key-chain bfd-kc4
   ```

6. On Device B, specify the unique security authentication information for BFD sessions:

   - The matching keychain name as specified in Step 5.
   - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
   - The secret data used to allow access to the session.
   - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

   ```
   [edit security authentication-key-chains key-chain bfd-kc4]
   user@B# set key 5 secret
     "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
     DkP5Ft0IQFclev7N"
   user@B# set key 5 start-time "2011-1-1.12:00:00 -0800"
   ```

7. If you are done configuring Device B, commit the configuration.

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

8. Repeat the configuration on Device D.

   The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

Results

Confirm your configuration by issuing the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device B

```
user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
```

```
            address 172.16.1.1/24;
          }
        }
      }
      lo0 {
        unit 57 {
          family inet {
            address 10.0.0.1/32;
            address 10.0.0.2/32;
          }
        }
      }
```

```
user@B# show routing-options
static {
  route 192.168.47.0/24 {
    next-hop 172.16.1.2;
    bfd-liveness-detection {
      minimum-interval 1000;
      authentication {
        key-chain bfd-kc4;
        algorithm keyed-sha-1;
      }
    }
  }
}
```

```
user@B# show security
authentication-key-chains {
  key-chain bfd-kc4 {
    key 5 {
      secret
        "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
        DkP5Ft0IQFclev7N"; ## SECRET-DATA
      start-time "2011-1-1.12:00:00 -0800";
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

- Verifying That BFD Sessions Are Up on page 120
- Viewing Details About the BFD Session on page 121
- Viewing Extensive BFD Session Information on page 121

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up.

Action From operational mode, enter the **show bfd session** command.

```
user@B> show bfd session
                                                  Detect   Transmit
Address                   State     Interface     Time     Interval  Multiplier
```



```
172.16.1.2               Up        ge-1/2/0.0     3.000     1.000        3

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning The command output shows that the BFD session is up.

Viewing Details About the BFD Session

Purpose View details about the BFD sessions and make sure that authentication is configured.

Action From operational mode, enter the show bfd session **detail** command.

```
user@B> show bfd session detail
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.1.2               Up        ge-1/2/0.0     3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000, Authenticate
 Session up time 00:53:58
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive BFD Session Information

Purpose View more detailed information about the BFD sessions.

Action From operational mode, enter the show bfd session **extensive** command.

```
user@B> show bfd session extensive
Address                  State     Interface      Time     Interval  Multiplier
172.16.1.2               Up        ge-1/2/0.0     3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000, Authenticate
        keychain bfd-kc4, algo keyed-sha-1, mode strict
 Session up time 01:39:45
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 3, remote discriminator 4
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

Related Documentation

- Examples: Configuring BFD for Static Routes on page 99

CHAPTER 5

Packet Forwarding Behavior

Example: Enabling Indirect Next Hops on the Packet Forwarding Engine

Understanding Indirect Next Hops

Junos OS supports the concept of an indirect next hop for all routing protocols that support indirectly connected next hops, also known as third-party next hops.

Because routing protocols such as internal BGP (IBGP) can send routing information about indirectly connected routes, Junos OS relies on routes from intra-AS routing protocols (OSPF, IS-IS, RIP, and static) to resolve the best directly connected next hop. The Routing Engine performs route resolution to determine the best directly connected next hop and installs the route to the Packet Forwarding Engine.

By default, Junos OS does not maintain the route for indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. As a result, when a rerouting event occurs, potentially thousands of route to forwarding next-hop bindings must be updated, which increases the route convergence time. Figure 13 on page 123 illustrates the route to forwarding next-hop bindings with indirect next hop disabled.

Figure 13: Route to Forwarding Next-Hop Bindings

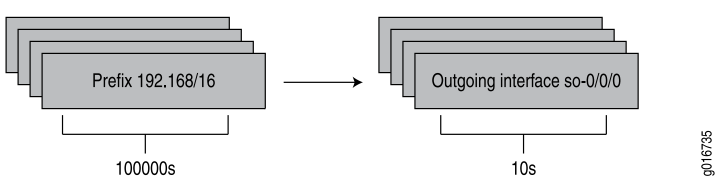

You can enable Junos OS to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. As a result, fewer route to forwarding next-hop bindings need to be updated, which improves the route convergence time. Figure 14 on page 124 illustrates the route to forwarding next-hop bindings with indirect next hop enabled.

Figure 14: Route to Forwarding Indirect Next-Hop Bindings

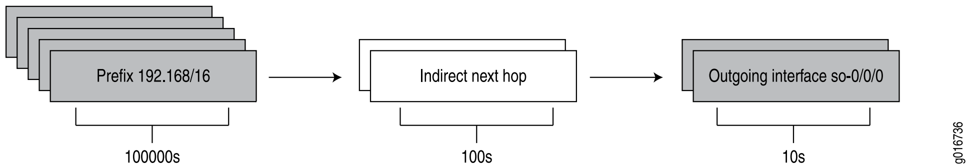

Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

This example shows how to use indirect next hops to promote faster network convergence (for example, in BGP networks) by decreasing the number of forwarding table changes required when a change in the network topology occurs.

Requirements

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

Overview

In this example, several devices are connected over unequal-cost paths. From Device R1 to Device R2, the path through Device R3 has a higher IGP metric than the path through Device R4. Device R1 has an internal BGP connection to Device R2. Device R0 injects multiple routes into the network, and Device R1 advertises those routes to Device R2. Because Device R2 is not directly connected to Device R1, Device R2's forwarding table contains indirect next hops. An interior gateway protocol, in this case OSPF, is running on the internal links among Devices R1, R2, R3, and R4. Each router is advertising its loopback interface IPv4 address.

On Device R2, the **indirect-next-hop** statement enables Junos OS to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. As a result, fewer route to forwarding next-hop bindings need to be updated, which improves the route convergence time if a path fails.

Figure 15 on page 125 shows the sample network.

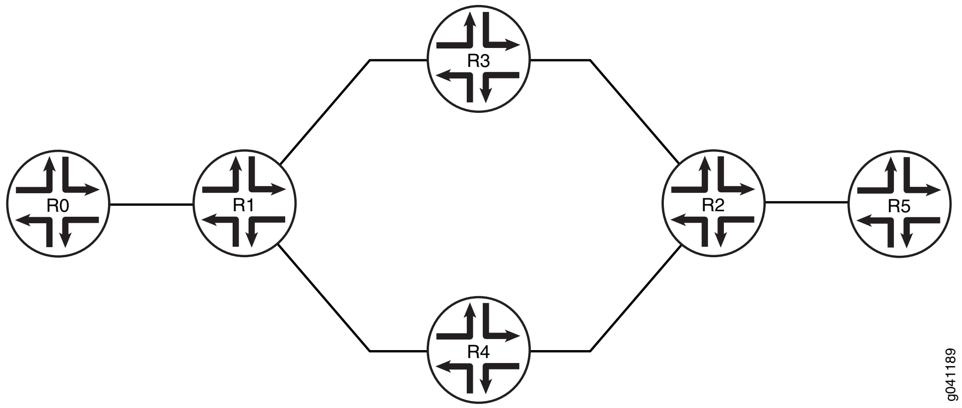

The "CLI Quick Configuration" on page 125 section shows the full configuration on all of the devices in Figure 15 on page 125. Otherwise, the example focuses on Device R0, Device R1, and Device R2.

Configuration

CLI Quick Configuration

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

Device R0

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 1.1.0.1/32
set interfaces lo0 unit 1 family inet address 1.1.0.2/32
set interfaces lo0 unit 1 family inet address 1.1.0.3/32
set interfaces lo0 unit 1 family inet address 1.1.0.4/32
set interfaces lo0 unit 1 family inet address 1.1.0.5/32
set interfaces lo0 unit 1 family inet address 1.1.0.6/32
set interfaces lo0 unit 1 family inet address 1.1.0.7/32
set interfaces lo0 unit 1 family inet address 1.1.0.8/32
set interfaces lo0 unit 1 family inet address 1.1.0.9/32
set routing-options static route 0.0.0.0/0 next-hop 10.0.0.2
```

Device R1

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces fe-1/2/2 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 2 family inet address 1.1.1.1/32
set protocols bgp export send-local
set protocols bgp export send-static
set protocols bgp group int type internal
set protocols bgp group int local-address 1.1.1.1
set protocols bgp group int neighbor 2.2.2.2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.5
set protocols ospf area 0.0.0.0 interface fe-1/2/2.9
set protocols ospf area 0.0.0.0 interface lo0.2
set policy-options policy-statement send-local from protocol local
set policy-options policy-statement send-local from protocol direct
set policy-options policy-statement send-local then accept
set policy-options policy-statement send-static from protocol static
set policy-options policy-statement send-static then accept
set routing-options static route 1.1.0.2/32 next-hop 10.0.0.1
```

```
set routing-options static route 1.1.0.1/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.3/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.4/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.5/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.6/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.7/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.8/32 next-hop 10.0.0.1
set routing-options static route 1.1.0.9/32 next-hop 10.0.0.1
set routing-options autonomous-system 65500
```

Device R2

```
set interfaces fe-1/2/0 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/2/1 unit 18 family inet address 10.0.0.18/30
set interfaces fe-1/2/2 unit 21 family inet
set interfaces lo0 unit 3 family inet address 2.2.2.2/32
set protocols bgp export send-local
set protocols bgp group int type internal
set protocols bgp group int local-address 2.2.2.2
set protocols bgp group int family inet unicast
set protocols bgp group int family inet-vpn unicast
set protocols bgp group int neighbor 1.1.1.1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.14
set protocols ospf area 0.0.0.0 interface fe-1/2/1.18
set protocols ospf area 0.0.0.0 interface lo0.3
set policy-options policy-statement send-local from protocol local
set policy-options policy-statement send-local from protocol direct
set policy-options policy-statement send-local then accept
set routing-options autonomous-system 65500
set routing-options forwarding-table indirect-next-hop
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces fe-1/2/1 unit 13 family inet address 10.0.0.13/30
set interfaces lo0 unit 4 family inet address 3.3.3.3/32
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6 metric 5000
set protocols ospf area 0.0.0.0 interface fe-1/2/1.13 metric 5000
set protocols ospf area 0.0.0.0 interface lo0.4
```

Device R4

```
set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-1/2/1 unit 17 family inet address 10.0.0.17/30
set interfaces lo0 unit 5 family inet address 4.4.4.4/32
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-1/2/1.17
set protocols ospf area 0.0.0.0 interface lo0.5
```

Device R5

```
set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30
set interfaces lo0 unit 6 family inet address 5.5.5.5/32
```

Configuring Device R0

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 R0:

1. Configure the interfaces, including multiple routes that can be injected into the network for demonstration purposes.

```
[edit interfaces]
user@R0# set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.1/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.2/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.3/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.4/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.5/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.6/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.7/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.8/32
user@R0# set interfaces lo0 unit 1 family inet address 1.1.0.9/32
```

2. Configure a static default route for network reachability.

```
[edit routing-options]
user@R0# set routing-options static route 0.0.0.0/0 next-hop 10.0.0.2
```

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

```
[edit]
user@R0# commit
```

Configuring Device R1

Step-by-Step Procedure

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

To configure Device R1:

1. Configure the interfaces, including multiple routes that can be injected into the network for demonstration purposes.

```
[edit interfaces]
user@R1# set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
user@R1# set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
user@R1# set interfaces fe-1/2/2 unit 9 family inet address 10.0.0.9/30

user@R1# set interfaces lo0 unit 2 family inet address 1.1.1.1/32
```

2. Configure BGP.

```
[edit routing-options]
user@R1# set protocols bgp export send-local
```

```
user@R1# set protocols bgp export send-static
user@R1# set protocols bgp group int type internal
user@R1# set protocols bgp group int local-address 1.1.1.1
user@R1# set protocols bgp group int neighbor 2.2.2.2
```

3. Configure OSPF.

```
[edit protocols]
user@R1# set protocols ospf area 0.0.0.0 interface fe-1/2/1.5
user@R1# set protocols ospf area 0.0.0.0 interface fe-1/2/2.9
user@R1# set protocols ospf area 0.0.0.0 interface lo0.2
```

4. Configure the routing policies.

```
user@R1# set policy-options policy-statement send-local from protocol local
user@R1# set policy-options policy-statement send-local from protocol direct
user@R1# set policy-options policy-statement send-local then accept

user@R1# set policy-options policy-statement send-static from protocol static
user@R1# set policy-options policy-statement send-static then accept
```

5. Configure a set of static routes to the set of interfaces configured on Device R0.

```
user@R1# set routing-options static route 1.1.0.2/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.1/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.3/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.4/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.5/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.6/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.7/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.8/32 next-hop 10.0.0.1
user@R1# set routing-options static route 1.1.0.9/32 next-hop 10.0.0.1
```

6. Configure the autonomous system (AS) identifier.

```
user@R1# set routing-options autonomous-system 65500
```

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

```
[edit]
user@R1# commit
```

Configuring Device R2

Step-by-Step Procedure

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

To configure Device R2:

1. Configure the interfaces, including multiple routes that can be injected into the network for demonstration purposes.

```
[edit interfaces]
user@R2# set interfaces fe-1/2/0 unit 14 family inet address 10.0.0.14/30
user@R2# set interfaces fe-1/2/1 unit 18 family inet address 10.0.0.18/30
user@R2# set interfaces fe-1/2/2 unit 21 family inet address 10.0.0.21/30;
```



```
user@R2# set interfaces lo0 unit 3 family inet address 2.2.2.2/32
```

2. Configure BGP.

```
[edit routing-options]
user@R2# set protocols bgp export send-local
user@R2# set protocols bgp group int type internal
user@R2# set protocols bgp group int local-address 2.2.2.2
user@R2# set protocols bgp group int family inet unicast
user@R2# set protocols bgp group int family inet-vpn unicast
user@R2# set protocols bgp group int neighbor 1.1.1.1
```

3. Configure OSPF.

```
[edit protocols]
user@R2# set protocols ospf area 0.0.0.0 interface fe-1/2/0.14
user@R2# set protocols ospf area 0.0.0.0 interface fe-1/2/1.18
user@R2# set protocols ospf area 0.0.0.0 interface lo0.3
```

4. Configure the routing policies.

```
user@R2# set policy-options policy-statement send-local from protocol local
user@R2# set policy-options policy-statement send-local from protocol direct
user@R2# set policy-options policy-statement send-local then accept
```

5. Configure the AS identifier.

```
user@R2# set routing-options autonomous-system 65500
```

6. Enable indirect next hops in the forwarding plane.

```
user@R2# set routing-options forwarding-table indirect-next-hop
```

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

```
[edit]
user@R2# commit
```

Results

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

Device R0

```
user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 1.1.0.1/32;
      address 1.1.0.2/32;
      address 1.1.0.3/32;
      address 1.1.0.4/32;
      address 1.1.0.5/32;
```

```
            address 1.1.0.6/32;
            address 1.1.0.7/32;
            address 1.1.0.8/32;
            address 1.1.0.9/32;
        }
    }
}

user@R0# show routing-options
static {
    route 0.0.0.0/0 next-hop 10.0.0.2;
}
```

Device R1

```
user@R1# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            address 10.0.0.2/30;
        }
    }
}
fe-1/2/1 {
    unit 5 {
        family inet {
            address 10.0.0.5/30;
        }
    }
}
fe-1/2/2 {
    unit 9 {
        family inet {
            address 10.0.0.9/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 1.1.1.1/32;
        }
    }
}

user@R1# show protocols
bgp {
    export [ send-local send-static ];
    group int {
        type internal;
        local-address 1.1.1.1;
        neighbor 2.2.2.2;
    }
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/1.5;
        interface fe-1/2/2.9;
```

```
        interface lo0.2;
    }
}
```

```
user@R1# show policy-options
policy-statement send-local {
    from protocol [ local direct ];
    then accept;
}
policy-statement send-static {
    from protocol static;
    then accept;
}
```

```
user@R1# show routing-options
static {
    route 1.1.0.2/32 next-hop 10.0.0.1;
    route 1.1.0.1/32 next-hop 10.0.0.1;
    route 1.1.0.3/32 next-hop 10.0.0.1;
    route 1.1.0.4/32 next-hop 10.0.0.1;
    route 1.1.0.5/32 next-hop 10.0.0.1;
    route 1.1.0.6/32 next-hop 10.0.0.1;
    route 1.1.0.7/32 next-hop 10.0.0.1;
    route 1.1.0.8/32 next-hop 10.0.0.1;
    route 1.1.0.9/32 next-hop 10.0.0.1;
}
autonomous-system 65500;
```

Device R2

```
user@R2# show interfaces
fe-1/2/0 {
    unit 14 {
        family inet {
            address 10.0.0.14/30;
        }
    }
}
fe-1/2/1 {
    unit 18 {
        family inet {
            address 10.0.0.18/30;
        }
    }
}
fe-1/2/2 {
    unit 21 {
        family inet {
            address 10.0.0.21/30
        }
    }
}
lo0 {
    unit 3 {
        family inet {
            address 2.2.2.2/32;
        }
    }
}
```



```
user@R2# show protocols
bgp {
  export send-local;
  group int {
    type internal;
    local-address 2.2.2.2;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    neighbor 1.1.1.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.14;
    interface fe-1/2/1.18;
    interface lo0.3;
  }
}

user@R2# show policy-options
policy-statement send-local {
  from protocol [ local direct ];
  then accept;
}

user@R2# show routing-options
autonomous-system 65500;
forwarding-table {
  indirect-next-hop;
}
```

Configure Device R3, Device R4, and Device R5, as shown in "CLI Quick Configuration" on page 125.

Verification

Confirm that the configuration is working properly.

Verifying That the Routes Have the Expected Indirect-Next-Hop Flag

Purpose Make sure that Device R2 is configured to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table.

Action

```
user@R2> show krt indirect-next-hop
Indirect Nexthop:
Index: 262143 Protocol next-hop address: 1.1.1.1
  RIB Table: inet.0
  Policy Version: 0                     References: 4
  Locks: 2                              0x9290488
  Flags: 0x1
  Ref RIB Table: unknown
        Next hop: 10.0.0.17 via fe-1/2/1.18
Indirect Nexthop:
Index: 262142 Protocol next-hop address: 10.0.0.1
  RIB Table: inet.0
  Policy Version: 0                     References: 9
  Locks: 2                              0x9290570
  Flags: 0x1
  Ref RIB Table: unknown
        Next hop: 10.0.0.17 via fe-1/2/1.18
```

Meaning

The **0x1** flag in the output indicates that Device R2 is configured to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. When the **indirect-next-hop** statement is deleted or deactivated, this flag changes to **0x0**.

NOTE: The **show krt indirect-next-hop** command is hidden and is therefore undocumented. The **show krt indirect-next-hop** command is shown here because this is the only command that verifies the indirect next-hop feature. The best verification method is, of course, monitoring network performance during reconvergence after a path failure.

Related Documentation

- Example: Controlling Static Routes in Routing and Forwarding Tables on page 32

Example: Configuring Unicast Reverse-Path-Forwarding Check

Understanding Unicast Reverse Path Forwarding

IP spoofing can occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router or switch determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router or switch forwards the packet to the destination address. If it is not from a valid path, the router or switch discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.

NOTE: **Reverse path forwarding is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.**

Example: Configuring Unicast Reverse-Path-Forwarding Check

Unicast reverse path forwarding (RPF) helps protect against DoS and DDoS attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast RPF to filter incoming traffic.

- Requirements on page 134
- Overview on page 134
- Configuration on page 135
- Verification on page 140

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a DoS attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and DDoS attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding table entry for its source address. If the device uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the device forwards the packet. If the device does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the device discards the packet.

In this example, Device B has unicast RPF configured. Device A is using OSPF to advertise a prefix for the link that connects to Device D. OSPF is enabled on the links between Device B and Device C and the links between Device A and Device C, but not on the links between Device A and Device B. Therefore, Device B learns about the route to Device D through Device C.

This example also includes a fail filter. When a packet fails the unicast RPF check, the fail filter is evaluated to determine if the packet should be accepted anyway. The fail filter in this example allows Device B's interfaces to accept Dynamic Host Configuration

Protocol (DHCP) packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

Figure 16 on page 135 shows the sample network.

Figure 16: Unicast RPF Sample Topoolgy

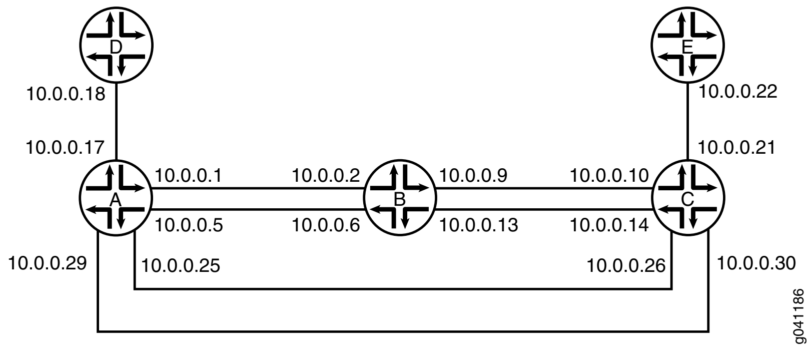

Configuration

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 A

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces fe-0/0/2 unit 5 family inet address 10.0.0.5/30
set interfaces fe-0/0/1 unit 17 family inet address 10.0.0.17/30
set interfaces fe-0/1/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/1/1 unit 29 family inet address 10.0.0.29/30
set protocols ospf export send-direct
set protocols ospf area 0.0.0.0 interface fe-0/1/1.25
set protocols ospf area 0.0.0.0 interface fe-1/1/1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct from route-filter 10.0.0.16/30 exact
set policy-options policy-statement send-direct then accept
```

Device B

```
set interfaces fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/1/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/1 unit 9 family inet address 10.0.0.9/30
set interfaces fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/0 unit 13 family inet address 10.0.0.13/30
set protocols ospf area 0.0.0.0 interface fe-0/1/1.9
set protocols ospf area 0.0.0.0 interface fe-0/1/0.13
set routing-options forwarding-table unicast-reverse-path active-paths
set firewall filter rpf-special-case-dhcp term allow-dhcp from source-address 0.0.0.0/32
set firewall filter rpf-special-case-dhcp term allow-dhcp then count rpf-dhcp-traffic
set firewall filter rpf-special-case-dhcp term allow-dhcp then accept
set firewall filter rpf-special-case-dhcp term default then log
set firewall filter rpf-special-case-dhcp term default then reject
```


Device C

```
set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-0/0/2 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/0/2 unit 21 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 30 family inet address 10.0.0.30/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-0/0/2.14
set protocols ospf area 0.0.0.0 interface fe-1/2/2.26
set protocols ospf area 0.0.0.0 interface fe-1/2/1.30
```

Device D

```
set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30
```

Device E

```
set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30
```

Configuring Device A

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 A:

1. Configure the interfaces.

```
[edit interfaces]
user@A# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@A# set fe-0/0/2 unit 5 family inet address 10.0.0.5/30

user@A# set fe-0/0/1 unit 17 family inet address 10.0.0.17/30

user@A# set fe-0/1/1 unit 25 family inet address 10.0.0.25/30

user@A# set fe-1/1/1 unit 29 family inet address 10.0.0.29/30
```

2. Configure OSPF.

```
[edit protocols ospf]
user@A# set export send-direct
user@A# set area 0.0.0.0 interface fe-0/1/1.25
user@A# set area 0.0.0.0 interface fe-1/1/1.29
```

3. Configure the routing policy.

```
[edit policy-options policy-statement send-direct]
user@A# set from protocol direct
user@A# set from route-filter 10.0.0.16/30 exact
user@A# set then accept
```

4. If you are done configuring Device A, commit the configuration.

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


Configuring Device B

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 B:

1. Configure the interfaces.

   [edit interfaces]
   user@B# **set fe-1/2/0 unit 2 family inet address 10.0.0.2/30**

   user@B# **set fe-1/1/1 unit 6 family inet address 10.0.0.6/30**

   user@B# **set fe-0/1/1 unit 9 family inet address 10.0.0.9/30**

   user@B# **set fe-0/1/0 unit 13 family inet address 10.0.0.13/30**

2. Configure OSPF.

   [edit protocols ospf area 0.0.0.0]
   user@B# **set interface fe-0/1/1.9**
   user@B# **set interface fe-0/1/0.13**

3. Configure unicast RPF, and apply the optional fail filter.

   [edit interfaces]
   user@B# **set fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp**

   user@B# **set fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp**

   user@B# **set fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp**

   user@B# **set fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp**

4. (Optional) Configure the fail filter that gets evaluated if a packet fails the RPF check.

   [edit firewall filter rpf-special-case-dhcp]
   user@B# **set term allow-dhcp from source-address 0.0.0.0/32**
   user@B# **set term allow-dhcp then count rpf-dhcp-traffic**
   user@B# **set term allow-dhcp then accept**
   user@B# **set term default then log**
   user@B# **set term default then reject**

5. (Optional) Configure only active paths to be considered in the RPF check.

   This is the default behavior.

   [edit routing-options forwarding-table]
   user@B# set unicast-reverse-path **active-paths**

6. If you are done configuring Device B, commit the configuration.

   [edit]
   user@B# **commit**

Results

Confirm your configuration by issuing the **show firewall**, **show interfaces**, **show protocols**, **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.

Device A

```
user@A# show interfaces
fe-1/2/0 {
   unit 1 {
      family inet {
         address 10.0.0.1/30;
      }
   }
}
fe-0/0/2 {
   unit 5 {
      family inet {
         address 10.0.0.5/30;
      }
   }
}
fe-0/0/1 {
   unit 17 {
      family inet {
         address 10.0.0.17/30;
      }
   }
}
fe-0/1/1 {
   unit 25 {
      family inet {
         address 10.0.0.25/30;
      }
   }
}
fe-1/1/1 {
   unit 29 {
      family inet {
         address 10.0.0.29/30;
      }
   }
}

user@A# show protocols
ospf {
   export send-direct;
   area 0.0.0.0 {
      interface fe-0/1/1.25;
      interface fe-1/1/1.29;
   }
}

user@A# show policy-options
policy-statement send-direct {
   from {
```



```
        protocol direct;
        route-filter 10.0.0.16/30 exact;
    }
    then accept;
}
```

Device B

```
user@B# show firewall
filter rpf-special-case-dhcp {
    term allow-dhcp {
        from {
            source-address {
                0.0.0.0/32;
            }
        }
        then {
            count rpf-dhcp-traffic;
            accept;
        }
    }
    term default {
        then {
            log;
            reject;
        }
    }
}
user@B# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.2/30;
        }
    }
}
fe-1/1/1 {
    unit 6 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.6/30;
        }
    }
}
fe-0/1/1 {
    unit 9 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.9/30;
        }
    }
}
fe-0/1/0 {
    unit 13 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.13/30;
```

```
    }
  }
}

user@B# show protocols
ospf {
  area 0.0.0.0 {
    interface fe-0/1/1.9;
    interface fe-0/1/0.13;
  }
}

user@B# show routing-options
forwarding-table {
  unicast-reverse-path active-paths;
}
```

Enter the configurations on Device C, Device D, and Device E, as shown in "CLI Quick Configuration" on page 135.

Verification

Confirm that the configuration is working properly.

Confirm That Unicast RPF Is Enabled

Purpose Make sure that the interfaces on Device B have unicast RPF enabled.

Action

```
user@B> show interfaces fe-0/1/0.13 extensive
 Logical interface fe-0/1/0.13 (Index 73) (SNMP ifIndex 553) (Generation 208)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             999390
     Output bytes  :            1230122
     Input  packets:              12563
     Output packets:              12613
    Local statistics:
     Input  bytes  :             998994
     Output bytes  :            1230122
     Input  packets:              12563
     Output packets:              12613
    Transit statistics:
     Input  bytes  :                396                  0 bps
     Output bytes  :                  0                  0 bps
     Input  packets:                  0                  0 pps
     Output packets:                  0                  0 pps
    Protocol inet, MTU: 1500, Generation: 289, Route table: 22
      Flags: Sendbcast-pkt-to-re, uRPF
      RPF Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.12/30, Local: 10.0.0.13, Broadcast: 10.0.0.15,
Generation: 241
```

Meaning The **uRPF** flag confirms that unicast RPF is enabled on this interface.

Confirm That the Source Addresses Are Blocked

Purpose Use the **ping** command to make sure that Device B blocks traffic from unexpected source addresses.

Action From Device A, ping Device B’s interfaces, using 10.0.0.17 as the source address.

```
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.6 (10.0.0.6): 56 data bytes
^C
--- 10.0.0.6 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

Meaning As expected, the ping operation fails.

Confirm That the Source Addresses Are Unblocked

Purpose Use the **ping** command to make sure that Device B does not block traffic when the RPF check is deactivated.

Action

1. Deactivate the RPF check on one of the interfaces.
2. Rerun the ping operation.

```
user@B> deactivate interfaces fe-1/1/1.6 family inet rpf-check
```

```
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=63 time=1.263 ms
```

```
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.263/1.289/1.316/0.027 ms
```

Meaning As expected, the ping operation succeeds.

Related Documentation

- Example: Enabling Indirect Next Hops on the Packet Forwarding Engine on page 123

Understanding the Default Routing Table Groups for Interface Routes on Packet Transport Routers

On PTX Series Packet Transport Routers, the default interface-route routing table groups differ from that of other Junos OS routing devices.

The PTX Series routers are MPLS transit platforms that do IP forwarding, typically using interior gateway protocol (IGP) routes. Interface routes are directly connected and local routes.

PTX Series routers are unlike other Junos OS routing devices in that they force an indirect next-hop resolution. PTX Series routers need the indirect next hop be resolved to create the chained composite next hop. This can cause routes to be hidden when the next-hop type is unusable.

To prevent routes from being hidden, PTX Series platforms automatically copy the routes in inet.0 into inet.2 and inet.3, and the routes in inet6.0 into inet6.2 and inet6.3.

The default interface routing table configuration on the PTX Series routers is as follows:

```
user@host# show routing-options | display inheritance defaults
##
## 'interface-routes' was inherited from group 'junos-defaults'
##
interface-routes {
  ##
  ## 'rib-group' was inherited from group 'junos-defaults'
  ##
  rib-group {
    ##
    ## 'junos-ifrg-inet0-to-inet2-and-inet3' was inherited from group 'junos-defaults'
    ##
    inet junos-ifrg-inet0-to-inet2-and-inet3;
    ##
    ## 'junos-ifrg-inet60-to-inet62-and-inet63' was inherited from group 'junos-defaults'
    ##
    inet6 junos-ifrg-inet60-to-inet62-and-inet63;
  }
}
rib-groups {
  ##
  ## 'junos-ifrg-inet0-to-inet2-and-inet3' was inherited from group 'junos-defaults'
  ##
  junos-ifrg-inet0-to-inet2-and-inet3 {
```



```
        ##
        ## 'inet.0' was inherited from group 'junos-defaults'
        ## 'inet.2' was inherited from group 'junos-defaults'
        ## 'inet.3' was inherited from group 'junos-defaults'
        ##
        import-rib [ inet.0 inet.2 inet.3 ];
    }
    ##
    ## 'junos-ifrg-inet60-to-inet62-and-inet63' was inherited from group 'junos-defaults'
    ##
    junos-ifrg-inet60-to-inet62-and-inet63 {
        ##
        ## 'inet6.0' was inherited from group 'junos-defaults'
        ## 'inet6.2' was inherited from group 'junos-defaults'
        ## 'inet6.3' was inherited from group 'junos-defaults'
        ##
        import-rib [ inet6.0 inet6.2 inet6.3 ];
    }
}
```

Related Documentation

- *Chained Composite Next Hops for Transit Devices*
- *Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Routers*

CHAPTER 6

Martian Addresses

Example: Configuring Martian Addresses

Understanding Martian Addresses

Martian addresses are host or network addresses about which all routing information is ignored. When received by the routing device, these routes are ignored. They commonly are sent by improperly configured systems on the network and have destination addresses that are obviously invalid.

In IPv6, the loopback address and the multicast resolve and discard routes are the default martian addresses.

In Junos OS Release 10.4R5 and later, the reserved IPv6 multicast address space (ff00::/8 and ff02::/16) is added to the list of martian addresses.

In Junos OS Release 9.6 and later, you can configure Class E addresses on interfaces. Class E addresses are treated like any other unicast address for the purpose of forwarding. To allow Class E addresses to be configured on interfaces, you must remove the Class E prefix from the list of martian addresses. To remove the Class E prefix from the list of martian addresses include the **martians 240/4 orlonger allow** statement at the **[edit routing-options]** hierarchy level.

To view the default and configured martian routes, run the **show route martians** command.

IPv4 Martian Addresses

user@host> **show route martians table inet.**

```
inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.1:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed

inet.2:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.3:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed
```

IPv6 Martian Addresses

user@host> **show route martians table inet6**

```
inet6.0:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.1:
             ::1/128 exact -- disallowed

inet6.2:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.3:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
```

Example: Configuring Martian Addresses

This example shows how to remove the Class E prefix from the list of martian addresses.

Requirements

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

Overview

In this example, Junos OS defaults are modified to allow the 240.0.0.0/4 address block. This block of addresses is known as the experimental Class E addresses. In Junos OS Release 9.6 and later, you can configure Class E addresses on interfaces and use them for forwarding traffic. However, to do this, you must first allow routing on this address block.

This example also shows how to modify the martian addresses in the IPv6 routing table, **inet6.0**.

Configuration

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.

```
set routing-options rib inet.1 martians 240.0.0.0/4 orlonger allow
set routing-options rib inet6.0 martians fd00::/8 orlonger
set routing-options rib inet.3 martians 240.0.0.0/4 orlonger allow
set routing-options rib inet.2 martians 240.0.0.0/4 orlonger allow
set routing-options martians 240.0.0.0/4 orlonger allow
```

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 martian routes:

1. Allow Class C addresses in the default unicast routing table.

   ```
   [edit routing-options]
   user@host# set martians 240.0.0.0/4 orlonger allow
   ```

2. Allow Class C addresses in the routing table that is used for the IPv4 multicast forwarding cache.

   ```
   [edit routing-options]
   user@host# set rib inet.1 martians 240.0.0.0/4 orlonger allow
   ```

3. Allow Class C addresses in the routing table that is used for multicast reverse path forwarding (RPF) lookup.

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



```
user@host# set rib inet.2 martians 240.0.0.0/4 orlonger allow
```

4. Allow Class C addresses in the routing table that stores MPLS LSP information.

```
[edit routing-options]
user@host# set rib inet.3 martians 240.0.0.0/4 orlonger allow
```

5. Add a disallowed martian route to the IPv6 unicast routing table.

```
[edit routing-options]
user@host# set rib inet6.0 martians fd00::/8 orlonger
```

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

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

Results

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

```
user@host# show routing-options
rib inet.1 {
  martians {
    240.0.0.0/4 orlonger allow;
  }
}
rib inet6.0 {
  martians {
    fd00::/8 orlonger;
  }
}
rib inet.3 {
  martians {
    240.0.0.0/4 orlonger allow;
  }
}
rib inet.2 {
  martians {
    240.0.0.0/4 orlonger allow;
  }
}
martians {
  240.0.0.0/4 orlonger allow;
}
```

Verification

Confirm that the configuration is working properly.

- Verifying That the 240.0.0.0/4 Routes Are Now Accepted on page 149
- Verifying That the fd00::/8 Routes Are Now Rejected on page 149

Verifying That the 240.0.0.0/4 Routes Are Now Accepted

Purpose Make sure that the 240.0.0.0/4 route appears in the routing tables as allowed.

Action

```
user@host> show route martians table inet.
inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- allowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.1:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- allowed

inet.2:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- allowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.3:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- allowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed
```

Meaning The output shows that the 240.0.0.0/4 route is allowed.

Verifying That the fd00::/8 Routes Are Now Rejected

Purpose Make sure that the fd00::/8 route appears in the IPv6 unicast routing table as disallowed.

Action

```
user@host> show route martians table inet6.0
inet6.0:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
             fd00::/8 orlonger -- disallowed
```

Meaning The output shows that the fd00::/8 route is disallowed.

Related Documentation

- Examples: Creating a Routing Table and Populating It with Routes on page 38
- Example: Controlling Static Routes in Routing and Forwarding Tables on page 32

CHAPTER 7

Debugging and Trace Operations

Example: Tracing Global Routing Protocol Operations

Understanding Global Routing Protocol Tracing Operations

Global routing protocol tracing operations track all general routing operations and record them in a log file. To set protocol-specific tracing operations and to modify the global tracing operations for an individual protocol, configure tracing for that protocol.

Using the **traceoptions** statement, you can specify the following global routing protocol tracing flags:

- **all**—All tracing operations
- **condition-manager**—Condition manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-synchronization**—Nonstop routing synchronization events
- **parse**—Configuration parsing
- **policy**—Policy operations and actions
- **regex-parse**—Regular expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

NOTE: Use the **all** flag with caution. This flag might cause the CPU to become very busy.

Example: Tracing Global Routing Protocol Operations

This example shows how to list and view files that are created when you enable global routing trace operations.

Requirements

You must have the **view** privilege.

Overview

To configure global routing protocol tracing, include the **traceoptions** statement at the **[edit routing-options]** hierarchy level:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}
```

The flags in a **traceoptions flag** statement are identifiers. When you use the **set** command to configure a flag, any flags that might already be set are not modified. In the following example, setting the **timer** tracing flag has no effect on the already configured **task** flag. Use the **delete** command to delete a particular flag.

```
[edit routing-options traceoptions]
user@host# show
flag task;
user@host# set traceoptions flag timer
user@host# show
flag task;
flag timer;
user@host# delete traceoptions flag task
user@host# show
flag timer;
```

This example shows how to configure and view a trace file that tracks changes in the routing table. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.

TIP: To view a list of hierarchy levels that support tracing operations, enter the **help apropos traceoptions** command in configuration mode.

Configuration

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.

```
set routing-options traceoptions file routing-table-changes
set routing-options traceoptions file size 10m
set routing-options traceoptions file files 10
set routing-options traceoptions flag route
set routing-options static route 1.1.1.2/32 next-hop 10.0.45.6
```

Configuring Trace Operations

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 the trace operations:

1. Configure trace operations.

   ```
   [edit routing-options traceoptions]
   user@host# set file routing-table-changes
   user@host# set file size 10m
   user@host# set file files 10
   user@host# set flag route
   ```

2. Configure a static route to cause a change in the routing table.

   ```
   [edit routing-options static]
   user@host# set route 1.1.1.2/32 next-hop 10.0.45.6
   ```

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

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

Viewing the Trace File

Step-by-Step Procedure

To view the trace file:

1. In operational mode, list the log files on the system.

   ```
   user@host> file list /var/log
   /var/log:
   ...
   routing-table-changes
   ...
   ```

2. View the contents of the **routing-table-changes** file.

   ```
   user@host> file show /var/log/routing-table-changes
   Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
   Dec 15 11:09:29.496507
   Dec 15 11:09:29.496507 Tracing flags enabled: route
   Dec 15 11:09:29.496507
   ```



```
Dec 15 11:09:29.533203 inet_routerid_notify: Router ID: 192.168.4.1
Dec 15 11:09:29.533334 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.533381 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.533420 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.534915 inet_routerid_notify: Router ID: 192.168.4.1
Dec 15 11:09:29.542934 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.549253 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.556878 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.582990 rt_static_reinit: examined 3 static nexthops, 0
unreferenced
Dec 15 11:09:29.589920
Dec 15 11:09:29.589920 task_reconfigure reinitializing done
...
```

3. Filter the output of the log file.

```
user@host> file show /var/log/routing-table-changes | match 1.1.1.2
Dec 15 11:15:30.780314 ADD      1.1.1.2/32          nhid 0 gw 10.0.45.6
   Static   pref 5/0 metric  at-0/2/0.0 <ctive Int Ext>
Dec 15 11:15:30.782276 KRT Request: send len 216 v104 seq 0 ADD route/user
af 2 table 0 infot 0 addr 1.1.1.2 nhop-type unicast nhindex 663
```

4. View the tracing operations in real time by running the **monitor start** command with an optional **match** condition.

```
user@host> monitor start routing-table-changes | match 1.1.1.2
Aug 10 19:21:40.773467 BGP RECV         0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlri: 0.0.0.0/0 qualified bnp->ribact 0x0
l2afcb 0x0
```

5. Deactivate the static route.

```
user@host# deactivate routing-options static route 1.1.1.2/32
user@host# commit
```

```
*** routing-table-changes ***
Dec 15 11:42:59.355557 CHANGE   1.1.1.2/32          nhid 663 gw 10.0.45.6
     Static   pref 5/0 metric  at-0/2/0.0 <Delete Int Ext>
Dec 15 11:42:59.426887 KRT Request: send len 216 v104 seq 0 DELETE route/user
 af 2 table 0 infot 0 addr 1.1.1.2 nhop-type discard filtidx 0
Dec 15 11:42:59.427366 RELEASE  1.1.1.2/32          nhid 663 gw 10.0.45.6
     Static   pref 5/0 metric  at-0/2/0.0 <Release Delete Int Ext>
```

6. Halt the **monitor** command by pressing Enter and typing **monitor stop**.

```
[Enter]
user@host> monitor stop
```

7. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

   When configuration is deactivated, it appears in the configuration with the **inactive** tag.

```
[edit routing-options]
user@host# deactivate traceoptions
user@host# commit
```

```
[edit routing-options]
user@host# show

inactive: traceoptions {
```



```
        file routing-table-changes size 10m files 10;
        flag route;
    }
    static {
        inactive: route 1.1.1.2/32 next-hop 10.0.45.6;
    }
```

8. To reactivate trace operations, use the **activate** configuration-mode statement.

   [edit routing-options]
   user@host# **activate traceoptions**
   user@host# **commit**

Results

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

```
user@host# show routing-options
traceoptions {
  file routing-table-changes size 10m files 10;
  flag route;
}
static {
  route 1.1.1.2/32 next-hop 10.0.45.6;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Trace Log File Is Operating

Purpose Make sure that events are being written to the log file.

Action

```
user@host> show log routing-table-changes
Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
```

Related Documentation

- *Example: Configuring BGP Trace Operations*

CHAPTER 8

Configuration Statements

access (Static Access Routes)

Syntax

```
access {
  route ip-prefix</prefix-length> {
    metric route-cost;
    next-hop next-hop;
    preference route-distance;
    qualified-next-hop next-hop;
    tag tag-number
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure access routes.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

access-internal (Static Access-Internal Routes)

Syntax

```
access-internal {
  route ip-prefix</prefix-length> {
    next-hop next-hop;
    qualified-next-hop next-hop
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure parameters for internal access routes.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management*
- *Configuring Dynamic Access-Internal Routes for PPP Subscriber Management*

active

Syntax (active | passive);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with **reject** next hops when they are inactive.

- **active**—Remove a route from the routing and forwarding tables when it becomes inactive.
- **passive**—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.

Include the **active** statement when configuring an individual route in the **route** portion of the **static** statement to override a **passive** option specified in the **defaults** portion of the statement.

Default **active**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88

aggregate (Routing)

Syntax

```
aggregate {
  defaults {
    ... aggregate-options ...
  }
  route destination-prefix {
    policy policy-name;
    ... aggregate-options ...
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure aggregate routes.

Options

aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***aggregate-options***. Each option is explained separately.

- **(active | passive);**
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number ip-address*>**;
- **(brief | full)**;
- **community [*community-ids*]**;
- **discard**;
- **(metric | metric2 | metric3 | metric4) *value* <type *type*>**;
- **(preference | preference2 | color | color2) *preference* <type *type*>**;
- **tag *string***;

defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Summarizing Routes Through Route Aggregation on page 82

as-path (Routing Options)

Syntax as-path <*as-path*> <aggregator *as-number ip-address*> <atomic-aggregate> <origin (egp | igp | incomplete)>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.

In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.

Default No AS path information is associated with static routes.

Options **aggregator**—(Optional) Attach the BGP **aggregator** path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for ***as-number***, followed by the IP address of the BGP system that formed the aggregate route for ***ip-address***.

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([ ]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for ***as-path***. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*

auto-export

Syntax

```
auto-export {
  disable;
  family inet {
    disable;
    flow {
      disable;
      rib-group rib-group;
    }
    multicast {
      disable;
      rib-group rib-group;
    }
    unicast {
      disable;
      rib-group rib-group;
    }
  }
  family inet6 {
    disable;
    multicast {
      disable;
      rib-group rib-group;
    }
    unicast {
      disable;
      rib-group rib-group;
    }
  }
  family iso {
    disable;
    unicast {
      disable;
      rib-group rib-group;
    }
  }
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Export routes between routing instances.

This statement enables you to leak routes between VPN routing and forwarding (VRF) instances that are locally configured on a provider edge (PE) router. Auto export is always applied on the local PE router, because it applies to only local prefix leaking by evaluating the export policy of each VRF and determining which route targets can be leaked. The standard VRF import and export policies affect remote PE prefix leaking.

You can use this statement as an alternative to using the VRF import and export policies.

Options **(disable | enable)**—Disable or enable auto-export.

Default: Enable

family—Address family.

inet—IP version 4 (IPv4) address family.

multicast—Multicast routing information.

unicast—Unicast routing information.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

autonomous-system

Syntax

autonomous-system *autonomous-system* <asdot-notation> <loops *number*> {
  independent-domain <no-attrset>;
}

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
asdot-notation option introduced in Junos OS Release 9.3.
asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.
no-attrset option introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify the routing device's AS number.

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options

autonomous-system—AS number. Use a number assigned to you by the NIC.

Range: 1 through 4,294,967,295 (2^{32} – 1) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops ***number***—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1

NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the **loops** statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the **loops** statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring Internal BGP Peering*

bfd

Syntax

```
bfd {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic.

Default If you do not include this statement, no BFD tracing operations are performed.

Options

disable—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the **/var/log** directory . We recommend that you place global routing protocol tracing output in the **routing-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the BFD protocol tracing options:

- **adjacency**—Trace adjacency messages.
- **all**—Trace all options for BFD.
- **error**—Trace all errors.
- **event**—Trace all events.
- **issu**—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

match *regular-expression*—(Optional) Regular expression for lines to be logged.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the trace file again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace–control—To add this statement to the configuration.

Related Documentation

- Example: Configuring BFD for Static Routes on page 103

bfd-liveness-detection (Routing Options Static Route)

Syntax

bfd-liveness-detection {
  authentication {
    algorithm *algorithm-name*;
    key-chain *key-chain-name*;
    loose-check;
  }
  detection-time {
    threshold *milliseconds*;
  }
  holddown-interval *milliseconds*;
  local-address *ip-address*;
  minimum-interval *milliseconds*;
  minimum-receive-interval *milliseconds*;
  minimum-receive-ttl *number*;
  multiplier *number*;
  neighbor *address*;
  no-adaptation;
  transmit-interval {
    minimum-interval *milliseconds*;
    threshold *milliseconds*;
  }
  version (1 | automatic);
}

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-options static route *destination-prefix*],

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

Release Information

Statement introduced before Junos OS Release 7.4.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2.
local-address statement introduced in Junos OS Release 8.2.
minimum-receive-ttl statement introduced in Junos OS Release 8.2.
Support for logical routers introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5.
no-adaptation statement introduced in Junos OS Release 9.0.
Support for IPv6 static routes introduced in Junos OS Release 9.1.
authentication algorithm, **authentication key-chain**, and **authentication loose-check** statements introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure bidirectional failure detection timers and authentication criteria for static routes.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route** ***destination-prefix*** **bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

minimum-receive-ttl *number*—Configure the time to live (TTL) for the multihop BFD session.

Range: 1 through 255

Default: 255

multiplier *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

neighbor *address*—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring BFD for Static Routes on page 103
- Example: Configuring BFD Authentication for Static Routes on page 116

brief

Syntax (brief | full);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.

- **brief**—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP **ATOMIC_ATTRIBUTE** path attribute is included with the aggregate route.
- **full**—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the **route** portion of the **generate** statement to override a **retain** option specified in the **defaults** portion of the statement.

Default **full**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- aggregate on page 160
- generate on page 189

color

See preference

community (Routing Options)

Syntax community ([*community-ids*] | no-advertise | no-export | no-export-subconfed | none);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Associate BGP community information with a static, aggregate, or generated route.

Default No BGP community information is associated with static routes.

Options ***community-ids***—One or more community identifiers. The ***community-ids*** format varies according to the type of attribute that you use.

The BGP community attribute format is ***as-number:community-value***:

- ***as-number***—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value.
- ***community-value***—Identifier of the community member. It can be a number from 0 through 65,535.

For more information about BGP community attributes, see the "Configuring the Extended Communities Attribute" section in the *Routing Policy Feature Guide for Routing Devices*.

For specifying the BGP community attribute only, you also can specify ***community-ids*** as one of the following well-known community names defined in RFC 1997:

- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

NOTE: Extended community attributes are not supported at the **[edit routing-options]** hierarchy level. You must configure extended communities at the **[edit policy-options]** hierarchy level. For information about configuring extended communities, see the *Routing Policy Feature Guide for Routing Devices*.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- **aggregate on page 160**
- **generate on page 189**
- **static on page 250**

confederation

Syntax confederation *confederation-autonomous-system* members [*autonomous-systems*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the routing device's confederation AS number.

If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more *confederations*. Each confederation is identified by its own AS number, which is called a *confederation AS number*. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.

The BGP path attributes **NEXT_HOP**, **LOCAL_PREF**, and **MULTI_EXIT_DISC**, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.

Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.

Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.

If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the **confederation** statement.

Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.

Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.

Options ***autonomous-systems***—AS numbers of the confederation members.

Range: 1 through 65,535

confederation-autonomous-system—Confederation AS number. Use one of the numbers assigned to you by the NIC.

Range: 1 through 65,535

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring BGP Confederations*

destination-networks

Syntax destination-networks *prefix*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*],
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*],
[edit routing-options dynamic-tunnels *tunnel-name*],
[edit routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.

Options ***prefix***—Destination prefix of the network.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring GRE Tunnels for Layer 3 VPNs*
- *Configuring Dynamic Tunnels*
- *Configuring RSVP Automatic Mesh*

disable (Routing Options)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options graceful-restart],
[edit logical-systems *logical-system-name* routing-options graceful-restart],
[edit routing-instances *routing-instance-name* routing-options graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Disable graceful restart.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS High Availability Library for Routing Devices*

discard

Syntax discard;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

To propagate static routes into the routing protocols, include the **discard** statement when you define the route, along with a routing policy.

NOTE: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as **route** ***address/mask-length*** **next-hop 127.0.0.1**) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.

Default When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88

- aggregate on page 160
- generate on page 189

dynamic-tunnels

Syntax

```
dynamic-tunnels tunnel-name {
  destination-networks prefix;
  gre;
  rsvp-te entry-name {
    destination-networks network-prefix;
    label-switched-path-template {
      default-template;
      template-name;
    }
  }
  source-address address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure a dynamic tunnel between two PE routers.

Options

tunnel-name—Name of the dynamic tunnel.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks*
- *Configuring GRE Tunnels for Layer 3 VPNs*
- *Configuring Dynamic Tunnels*

export (Routing Options)

Syntax export [*policy-name*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options forwarding-table],
[edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-instances *routing-instance-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Apply one or more policies to routes being exported from the routing table into the forwarding table.

In the **export** statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or a different **export** statement.

You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:

- Per-packet load balancing
- Class of service (CoS)

Options ***policy-name***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Load Balancing BGP Traffic*
- *Routing Policy Feature Guide for Routing Devices*
- *How a Routing Policy Is Evaluated*

export-rib

Syntax export-rib *routing-table-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the name of the routing table from which Junos OS should export routing information.

Options ***routing-table-name***—Routing table group name.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table on page 46
- **import-rib on page 194**
- **passive on page 220**

fate-sharing

Syntax

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Specify a backup path in case the primary path becomes unusable.

You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.

Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.

Options

cost ***value***—Cost assigned to the group.

Range: 1 through 65,535

Default: 1

from ***address***—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate **from** addresses in the group.

group ***group-name***—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.

to ***address***—(Optional) Address of egress router. For point-to-point link objects, you must specify both a **from** and a **to** address.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Ingress Router for MPLS-Signaled LSPs*
- *Junos OS MPLS Applications Library for Routing Devices*

filter

Syntax

```
filter {
  input filter-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the name of the routing table from which Junos OS should export routing information.

Options

input *filter-name*—Forwarding table filter name.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Forwarding Table Filters*
- *Applying Filters to Forwarding Tables*

firewall-install-disable

Syntax

```
firewall-install-disable;
```

Hierarchy Level

[edit routing-options flow],
[edit logical-systems *logical-system-name* routing-options flow],
[edit routing-instances *routing-instance-name* routing-options flow],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options flow]

Release Information

Statement introduced in Junos OS Releases 12.1X48 and 12.3.

Description

Disable installing flow-specification firewall filters in the firewall process (dfwd).

Default

If you omit the **firewall-install-disable** statement, the default behavior is **firewall-install-disable** mode.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Flow Routes*

flow

Syntax

```
flow {
  route name {
    match {
      match-conditions;
    }
    term-order (legacy | standard);
    then {
      actions;
    }
  }
  firewall-install-disable;
  term-order (legacy | standard);
  validation {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
term-order statement introduced in Junos OS Release 10.0
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a flow route.

Default

legacy

Options

actions—An action to take if conditions match.

match-conditions—Match packets to these conditions.

route *name*—Name of the flow route.

standard—Specify to use version 7 or later of the flow-specification algorithm.

term-order (legacy | standard)—Specify the version of the flow-specification algorithm.

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

then—Actions to take on matching packets.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Flow Routes*

forwarding-table

Syntax

```
forwarding-table {
  chained-composite-next-hop {
    ingress {
      l3vpn {
        extended-space;
      }
    }
  }
  export [ policy-name ];
  indexed-next-hop;
  (indirect-next-hop | no-indirect-next-hop);
  (indirect-next-hop-change-acknowledgements |
    no-indirect-next-hop-change-acknowledgements;)
  krt-nexthop-ack-timeout interval;
  unicast-reverse-path (active-paths | feasible-paths);
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Load Balancing BGP Traffic*

full

See brief

generate

Syntax

```
generate {
  defaults {
    generate-options;
  }
  route destination-prefix {
    policy policy-name;
    generate-options;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure generated routes, which are used as routes of last resort.

Options

defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the **generate** statement.

generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***generate-options***. Each option is explained separately.

- **(active | passive);**
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number in-address*>;**
- **(brief | full);**
- **community [*community-ids*];**
- **discard;**
- **(metric | metric2 | metric3 | metric4) *value* <type *type*>;**
- **(preference | preference2 | color | color2) *preference* <type *type*>;**
- **tag *string*;**

route *destination-prefix*—Configure a non-default generated route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.

- ***destination-prefix/prefix-length***—/***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Conditionally Generating Static Routes on page 88

graceful-restart (Enabling Globally)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-helper-recovery-time seconds;
  maximum-helper-restart-time seconds;
  notify-duration seconds;
  recovery-time seconds;
  restart-duration seconds;
  stale-routes-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.

For VPNs, the **graceful-restart** statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.

For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Default

Graceful restart is disabled by default.

Options

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

import (Routing Options)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution rib],
[edit logical-systems *logical-system-name* routing-options resolution rib],
[edit routing-instances *routing-instance-name* routing-options resolution rib],
[edit routing-options resolution rib]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify one or more import policies to use for route resolution.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

import-policy

Syntax import-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to routes imported into the routing table group. The **import-policy** statement complements the **import-rib** statement and cannot be used unless you first specify the routing tables to which routes are being imported.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options *policy-names*—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table on page 46
- export-rib on page 184
- passive on page 220

import-rib

Syntax import-rib [*routing-table-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be **inet.0** or **routing-instance-name.inet.0**. For IPv6 import routing tables, the primary routing table must be **inet6.0**.

In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options *routing-table-names*—Name of one or more routing tables.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table on page 46
- export-rib on page 184
- passive on page 220

independent-domain

Syntax independent-domain <no-attrset>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options autonomous-system *autonomous-system*],
[edit routing-instances *routing-instance-name* routing-options autonomous-system *autonomous-system*]

Release Information Statement introduced before Junos OS Release 7.4.
no-attrset option introduced in Junos OS Release 10.4.

Description Configure an independent AS domain.

The independent domain uses transitive path attribute 128 (attribute set) messages to tunnel the independent domain's BGP attributes through the internal BGP (IBGP) core.

This improves the transparency of Layer 3 VPN services for customer networks by preventing the IBGP routes that originate within an autonomous system (AS) in the customer network from being sent to a service provider's AS. Similarly, IBGP routes that originate within an AS in the service provider's network are prevented from being sent to a customer AS.

NOTE: **In Junos OS Release 10.3 and later, if BGP receives attribute 128 and you have not configured an independent domain in any routing instance, BGP treats the received attribute 128 as an unknown attribute.**

Options **no-attrset**—(Optional) Disables attribute set messages on the independent AS domain.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Layer 3 VPNs to Carry IBGP Traffic*
- autonomous-system on page 166

indexed-next-hop

Syntax indexed-next-hop;

Hierarchy Level [edit routing-options forwarding-table]

Release Information Statement introduced in Junos OS Release 12.1.

Description Allows Juniper Networks' devices to restore all next hops after a Graceful Routing Engine Switchover (GRES) or after a device restart. All next hops processed by the Junos OS software are converted to indexed next hops.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Indexed Next Hops*

indirect-next-hop

Syntax (indirect-next-hop | no-indirect-next-hop);

Hierarchy Level [edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.

NOTE:

- When virtual private LAN service (VPLS) is configured on the routing device, the **indirect-next-hop** statement is configurable at the **[edit routing-options forwarding-table]** hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.
- By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the **no-indirect-next-hop** statement.

Default Disabled.

Options **indirect-next-hop**—Enable indirectly connected next hops.

no-indirect-next-hop—Explicitly disable indirect next hops.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine on page 124

indirect-next-hop-change-acknowledgements

Syntax (indirect-next-hop-change-acknowledgements |
no-indirect-next-hop-change-acknowledgements);

Hierarchy Level [edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced in Junos OS Release 12.2.

Description Request that when an indirect next hop changes, an acknowledgement is sent.

Default Disabled by default in all platforms except PTX Series, where it is enabled by default.

Options **indirect-next-hop-change-acknowledgements**—Enable acknowledgements.

no-indirect-next-hop-change-acknowledgements—Explicitly disable acknowledgements.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine on page 124
- krt-nexthop-ack-timeout on page 204

input (Routing Options RIB)

Syntax input *filter-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options rib *routing-table-name* filter],
[edit routing-options rib *routing-table-name* filter]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the name of the input filter.

Options *filter-name*—Name of the input filter.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Forwarding Table Filters*
- *Applying Filters to Forwarding Tables*

install (Routing Options)

Syntax (install | no-install);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]
[edit routing-options static (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.

Options
install—Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the **route** portion of the **static** statement to override a **no-install** option specified in the **defaults** portion of the statement.

no-install—Do not install the route into the forwarding table, even if it is the route with the lowest preference.

Default: install

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring Static Routes on page 9
- **static on page 250**

instance-export

Syntax instance-export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Apply one or more policies to routes being exported from a routing instance.

Options ***policy-names***—Name of one or more export policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Routing Policy Feature Guide for Routing Devices*

instance-import

Syntax instance-import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Apply one or more policies to routes being imported into a routing instance.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Routing Policy Feature Guide for Routing Devices*

interface (Multicast Static Routes)

Syntax

```
interface interface-names {
  disable;
  maximum-bandwidth bps;
  no-qos-adjust;
  reverse-oif-mapping {
    no-qos-adjust;
  }
  subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable multicast traffic on an interface.

By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.

You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.

You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.

Static routes must be configured before you can enable multicast on an interface. Configuring the **interface** statement alone does not install any routes into the routing table. This feature relies on the static route configuration.

Options

interface-names—Name of one or more interfaces on which to enable multicast traffic.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Defining Interface Bandwidth Maximums*
- *Example: Configuring Multicast with Subscriber VLANs*

interface (Multicast Scoping)

Syntax interface [*interface-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the set of interfaces for multicast scoping.

Options ***interface-names***—Names of the interfaces on which to configure scoping. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify all.

NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- multicast on page 215

interface-routes

Syntax

```
interface-routes {
  family (inet | inet6) {
    export {
      lan;
      point-to-point;
    }
  }
  rib-group group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Description

Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0,** and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

Options **inet**—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Importing Direct and Static Routes Into a Routing Instance on page 42
- *Example: Configuring Multiple Routing Instances of OSPF*
- passive on page 220

krt-nexthop-ack-timeout

Syntax krt-nexthop-ack-timeout *interval*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced in Junos OS Release 12.2.

Description For indirect next-hop and multicast next-hop change acknowledgements, configure the time interval for which to wait for the next-hop acknowledgement. The routing protocol process (rpd) waits for the specified time period before changing the route to point to the new next hop.

If the acknowledgement is not received within the time period, it is assumed to have been received and the route is made to point to the new next hop.

Options *interval*—Kernel next-hop acknowledgement timeout interval.

Range: 1 through 100 seconds

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine on page 124
- indirect-next-hop-change-acknowledgements on page 198

lsp-next-hop (Static Routes)

Syntax

```
lsp-next-hop lsp-name {
  metric metric;
  preference preference;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*]
[edit routing-options static route *destination-prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify an LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.

NOTE: The preference and metric configured by means of the **lsp-next-hop** statement only apply to the LSP next hops. The LSP next-hop preference and metric override the route preference and metric (for that specific LSP next hop), similar to how the route preference overrides the default preference and metric (for that specific route).

Options

lsp-name—Name of the next-hop LSP.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on page 53

martians

Syntax

```
martians {
  destination-prefix match-type <allow>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure martian addresses.

Options

allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The **allow** option is the only supported action.

destination-prefix—Destination route you are configuring:

- ***destination-prefix/prefix-length—destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **default**—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address **0.0.0.0/0**.

match-type—Criteria that the destination must match:

- **exact**—Exactly match the route's mask length.
- **longer**—The route's mask length is greater than the specified mask length.
- **orlonger**—The route's mask length is equal to or greater than the specified mask length.
- **through *destination-prefix***—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.
- **upto *prefix-length***—The route's mask length falls between the two destination prefix lengths, inclusive.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Martian Addresses on page 146

maximum-paths

Syntax maximum-paths *path-limit* <log-interval *seconds*> <log-only | threshold *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a limit for the number of routes installed in a routing table based upon the route path.

NOTE: The **maximum-paths** statement is similar to the **maximum-prefixes** statement. The **maximum-prefixes** statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The **maximum-paths** limit applies the total number of routes (two). The **maximum-prefixes** limit applies to the total number of unique prefixes (one).

Options **log-interval** ***seconds***—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

path-limit—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.

Range: 1 through 4,294,967,295 ($2^{32} – 1$)

Default: No default

threshold ***value***—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the ***path-limit*** value that starts triggering the warnings.

Range: 1 through 100

NOTE: When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the ***path-limit*** value, then additional routes are rejected.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

maximum-prefixes

Syntax maximum-prefixes *prefix-limit* <log-interval *seconds*> <log-only | threshold *percentage*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a limit for the number of routes installed in a routing table based upon the route prefix.

Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.

NOTE: The **maximum-prefixes** statement is similar to the **maximum-paths** statement. The **maximum-prefixes** statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The **maximum-paths** limit applies the total number of routes (two). The **maximum-prefixes** limit applies to the total number of unique prefixes (one).

Options **log-interval *seconds***—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

prefix-limit—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.

Range: 1 through 4,294,967,295

Default: No default

threshold *value*—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the ***prefix-limit*** value that starts triggering the warnings.

Range: 1 through 100

NOTE: When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the ***prefix-limit*** value, then additional routes are rejected.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

med-igp-update-interval

Syntax

med-igp-update-interval *minutes*;

Hierarchy Level

[edit routing-options]

Release Information

Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp *offset* delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options

minutes—Interval to delay MED updates.

Range: 10 through 600

Default: 10 minutes

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates*
- *metric-out*

metric

Syntax metric *route-cost*;

Hierarchy Level [edit routing-options access route *ip-prefix*</*prefix-length*>]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the cost for an access route.

Options ***route-cost***—Specific cost you want to assign to the access route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring Static Routes on page 9

metric (Aggregate, Generated, or Static Route)

Syntax (metric | metric2 | metric3 | metric4) *metric* <type *type*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with **metric** (for the first metric value) and continuing with **metric2**, **metric3**, and **metric4**.

Options *metric*—Metric value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

type ***type***—(Optional) Type of route.

When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Summarizing Static Routes Through Route Aggregation on page 75
- Example: Conditionally Generating Static Routes on page 88
- **aggregate on page 160**
- **generate on page 189**
- **static on page 250**

metric (Qualified Next Hop on Static Route)

Syntax metric *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop],
[edit routing-options static route *destination-prefix* qualified-next-hop]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Metric value for a static route.

Options *metric*—Metric value.

Range: 0 through 4,294,967,295 (2^{32} – 1)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- qualified-next-hop on page 226
- static on page 250

multicast (Routing Options)

Syntax

```
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  interface interface-name {
    enable;
  }
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  ssm-groups {
    address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure generic multicast properties.

NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*
- *Examples: Configuring the Multicast Forwarding Cache*
- *Multicast Protocols Feature Guide for Routing Devices*
- (indirect-next-hop on page 197 | no-indirect-next-hop)

next-hop (Access)

Syntax next-hop *next-hop*;

Hierarchy Level [edit routing-options access route *ip-prefix</prefix-length>*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.

Options *next-hop*—Specific next-hop address you want to assign to the access route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring Static Routes on page 9

next-hop (Access Internal)

Syntax next-hop *next-hop*;

Hierarchy Level [edit routing-options access-internal route *ip-prefix</prefix-length>*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the next-hop address for an internal access route. Access routes are typically unnumbered interfaces.

Options *next-hop*—Specific next-hop address you want to assign to the internal access route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring Static Routes on page 9

no-delegate-processing

Syntax no-delegate-processing;

Hierarchy Level [edit routing-options ppm]

Release Information Statement introduced in Junos OS Release 10.1 for EX Series switches.

Description Disable distributed periodic packet management (PPM) processing and run all PPM processing on the Routing Engine.

PPM processing cannot be completely disabled on EX Series switches. You can only configure whether PPM processing is distributed between the access ports (EX3200 and EX4200 switches) or the line cards (EX8200 switches) and the Routing Engine or is handled just on the Routing Engine.

BEST PRACTICE: **Generally, you should only disable distributed PPM if Juniper Networks Customer Service advised you to do so. You should only disable distributed PPM if you have a compelling reason to disable it.**

Default Distributed PPM processing is enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*

nonstop-routing

Syntax nonstop-routing;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.

Default disabled

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Nonstop Active Routing*
- *Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)*

options (Routing Options)

Syntax

```
options {
  syslog (level level | upto level level);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.

Options

level *level*—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:

- **alert**—Conditions that should be corrected immediately, such as a corrupted system database.
- **critical**—Critical conditions, such as hard drive errors.
- **debug**—Software debugging messages.
- **emergency**—Panic or other conditions that cause the system to become unusable.
- **error**—Standard error conditions.
- **info**—Informational messages.
- **notice**—Conditions that are not error conditions, but might warrant special handling.
- **warning**—System warning messages.

upto level *level*—Log all messages up to a particular level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **syslog** in the *Junos OS Administration Library for Routing Devices*

p2mp-lsp-next-hop

Syntax

```
p2mp-lsp-next-hop {
  metric metric;
  preference preference;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*].
[edit routing-options static route *destination-prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify a point-to-multipoint LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Static LSPs*
- Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on page 53
- *Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems*

passive (Routing Options)

See active

policy (Aggregate and Generated Routes)

Syntax policy *policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Associate a routing policy when configuring an aggregate or generated route's destination prefix in the **routes** part of the **aggregate** or **generate** statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.

If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or *primary*, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.

The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:

1. Compare the protocol's **preference** of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.
2. Compare the protocol's **preference2** of the contributing routes. The lower **preference2** value is better. If only one route has **preference2**, then this route is preferred.
3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values.
    a. The primary contributor is the numerically smallest prefix value.
    b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options ***policy-name***—Name of a routing policy.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- **aggregate on page 160**
- **generate on page 189**

ppm

Syntax

```
ppm {
  no-delegate-processing;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description (M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).

After you disable PPM, PPM processing continues to run on the Routing Engine.

In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The **no-delegate-processing** statement disables the default behavior and restores the legacy behavior.

Default Distributed PPM processing is enabled for all protocols that use PPM.

Options **no-delegate-processing**—Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*
- *Configuring Distributed Periodic Packet Management*

preference (Routing Options)

Syntax (preference | preference2 | color | color2) *preference* <type *type*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value (**preference2**), as well as colors, which are even finer-grained preference values (**color** and **color2**).

If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.

Options *preference*—Preference value. A lower number indicates a more preferred route.

Range: 0 through 4,294,967,295 (2^{32} – 1)

Default: 5 (for static routes), 130 (for aggregate and generated routes)

type *type*—(Optional) Type of route.

Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- aggregate on page 160
- generate on page 189
- static on page 250

preference (Access)

Syntax preference *route-distance*;

Hierarchy Level [edit routing-options access route *ip-prefix</prefix-length>*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the distance for an access route.

Options ***route-distance***—Specific distance you want to assign to the access route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

prefix

Syntax prefix *destination-prefix*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the prefix for multicast scopes.

Options ***destination-prefix***—Address range for the multicast scope.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- *Example: Creating a Named Scope for Multicast Scoping*
- *multicast*

qualified-next-hop (Static Routes)

Syntax

```
qualified-next-hop (address | interface-name) {
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
        simple-password);
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  interface interface-name;
  metric metric;
  preference preference;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options rib inet6.0 static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit routing-options rib inet6.0 static route *destination-prefix*],
[edit routing-options static route *destination-prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options

address—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop interface for a static route, specify **qualified-next-hop** ***interface-name***, where ***interface-name*** is the name of the IPv4 or IPv6 unnumbered interface.

NOTE: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the unnumbered-address <*interface-name*> statement at the [edit interfaces <*interface-name*> unit <*logical-unit-number*> family <*family-name*>] hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Static Route Preferences and Qualified Next Hops on page 26
- Example: Enabling BFD on Qualified Next Hops in Static Routes on page 108

qualified-next-hop (Access)

Syntax qualified-next-hop *next-hop*;

Hierarchy Level [edit routing-options access route *ip-prefix*</*prefix-length*>]

Release Information
Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the qualified next-hop address for an access route.

Options *next-hop*—Specific qualified next-hop address you want to assign to the access route.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

qualified-next-hop (Access-Internal)

Syntax qualified-next-hop *next-hop*;

Hierarchy Level [edit routing-options access-internal route *ip-prefix</prefix-length>*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the qualified next-hop address for an internal access route.

Options *next-hop*—Specific qualified next-hop address you want to assign to the internal access route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

readadvertise

Syntax (readvertise | no-readvertise);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure whether static routes are eligible to be readvertised by routing protocols:

Default Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the **no-readvertise** statement.

Options **readvertise**—Readvertise static routes. Include the **readvertise** statement when configuring an individual route in the **route** portion of the **static** statement to override a **no-readvertise** option specified in the **defaults** portion of the statement.

no-readvertise—Mark a static route as being ineligible for readvertisement. Include the **no-readvertise** option when configuring the route.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Controlling Static Routes in Routing and Forwarding Tables on page 32
- static on page 250

resolution

Syntax

```
resolution {
  rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.

For example, you might want to direct **inet.2** route resolution to use topology routing tables **:red.inet.0** and **:blue.inet.0** for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.

You can specify up to two routing tables in the resolution-ribs statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure **resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]**, inet.0 is checked first and then inet.3 is checked.

NOTE: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the **show system processes** and the **show task memory** commands. For more information, see "Routing Protocol Process Memory FAQs" on page 380.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*
- *Example: Configuring Route Resolution on Route Reflectors*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*

resolution-ribs

Syntax

resolution-ribs [*routing-table-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution rib],
[edit logical-systems *logical-system-name* routing-options resolution rib],
[edit routing-instances *routing-instance-name* routing-options resolution rib],
[edit routing-options resolution rib]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Specify one or more routing tables to use for route resolution.

This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is **inet.3**, but you want to allow fallback resolution through **inet.0**. One example use case is overriding the **bgp.rtarget.0** (**family route-target**) routing table resolution from using only **inet.3** to using both **inet.3** and **inet.0**.

Options

routing-table-names—Name of one or more routing tables.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*

resolve

Syntax resolve;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the **inet.0** and **inet.3** routing tables.

Default Static routes can point only to a directly connected next hop.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- static on page 250

restart-duration

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols (isis | ospf | ospf3 | pim) graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3 | pim) graceful-restart],
[edit protocols (esis | isis | ospf | ospf3 | pim) graceful-restart],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3 | pim) graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the grace period for graceful restart globally.

Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.

Options *seconds*—Time for the graceful restart period.

Range:

The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:

- **[edit routing-options graceful-restart]** (global setting)—120 through 900
- ES-IS—30 through 300
- IS-IS—30 through 300
- OSPF/OSPFv3—1 through 3600
- PIM—30 through 300

Default:

The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:

- **[edit routing-options graceful-restart]** (global setting)—300
- ES-IS—180
- IS-IS—210
- OSPF/OSPFv3—180
- PIM—60

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Graceful Restart for VPNs*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

restart-duration (Routing Options)

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options graceful-restart],
[edit logical-systems *logical-system-name* routing-options graceful-restart],
[edit routing-instances *routing-instance-name* routing-options graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the restart timer for graceful restart.

Options *seconds*—Configure the time period for the restart to last.

Range: 120 through 900 seconds

Default: 300 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS High Availability Library for Routing Devices*

retain

Syntax (no-retain | retain);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:

Default Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

Options **no-retain**—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the **no-retain** statement. Include this statement when configuring an individual route in the **route** portion of the **static** statement to override a **retain** option specified in the **defaults** portion of the statement.

retain—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- static on page 250

rib (General)

Syntax

```
rib routing-table-name {
  aggregate {
    defaults {
      ... aggregate-options ...
    }
    route destination-prefix {
      policy policy-name;
      ... aggregate-options ...
    }
    generate {
      defaults {
        generate-options;
      }
      route destination-prefix {
        policy policy-name;
        generate-options;
      }
    }
    martians {
      destination-prefix match-type <allow>;
    }
  }
  static {
    defaults {
      static-options;
    }
    rib-group group-name;
    route destination-prefix {
      next-hop;
      static-options;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.

NOTE: The IPv4 multicast routing table (**inet.1**) and the IPv6 multicast routing table (**inet6.1**) are not supported for this statement.

Default If you do not specify a routing table name with the ***routing-table-name*** option, the software uses the default routing tables, which are **inet.0** for unicast routes and **inet.1** for the multicast cache.

Options ***routing-table-name***—Name of the routing table, in the following format: ***protocol*** [***.identifier***].

In a routing instance, the routing table name must include the routing instance name. For example, if the routing instance name is **link0**, the routing table name might be **link0.inet6.0**.

- ***protocol*** is the protocol family. It can be **inet6** for the IPv6 family, **inet** for the IPv4 family, **iso** for the ISO protocol family, or ***instance-name*.iso.0** for an ISO routing instance.
- ***identifier*** is a positive integer that specifies the instance of the routing table.

Default: inet.0

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Creating Routing Tables on page 40
- **passive on page 220**

rib (Route Resolution)

Syntax

```
rib routing-table-name {
  import [ policy-names ];
  resolution-ribs [ routing-table-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution],
[edit logical-systems *logical-system-name* routing-options resolution],
[edit routing-instances *routing-instance-name* routing-options resolution],
[edit routing-options resolution]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Specify a routing table name for route resolution.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

rib-group (Routing Options)

Syntax rib-group *group-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options interface-routes],
[edit logical-systems *logical-system-name* routing-options interface-routes],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static],
[edit logical-systems *logical-system-name* routing-options static],
[edit routing-instances *routing-instance-name* routing-options interface-routes],
[edit routing-options interface-routes],
[edit routing-options rib *routing-table-name* static],
[edit routing-options static]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure which routing table groups interface routes are imported into.

Options ***group-name***—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Importing Direct and Static Routes Into a Routing Instance on page 42
- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table on page 46
- interface-routes on page 203
- rib-groups on page 240

rib-groups

Syntax

```
rib-groups {
  group-name {
    export-rib group-name;
    import-policy [ policy-names ];
    import-rib [ group-names ];
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the **import-rib** statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the **export-rib** statement).

The first routing table you specify is the *primary routing table*, and any additional routing tables are the *secondary routing tables*.

The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify **inet.0** as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify **inet6.0** as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (**inet6.*x***).

In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the **import-rib** statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same **import-rib** statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the **inet6.3** routing table with IPv6 addresses that are compatible with IPv4. Specify **inet.0** as the primary routing table, and specify **inet6.3** as a secondary routing table.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

NOTE: **If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.**

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Exporting Specific Routes from One Routing Table Into Another Routing Table on page 46
- **rib-group on page 239**

route (Access)

Syntax

```
route ip-prefix</prefix-length> {
  metric route-cost;
  next-hop next-hop;
  preference route-distance;
  qualified-next-hop next-hop;
  tag tag-number;
}
```

Hierarchy Level [edit routing-options access]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the parameters for access routes.

Options *ip-prefix</prefix-length>*—Specific route prefix that you want to assign to the access route.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

route (Access-Internal)

Syntax

```
route  ip-prefix</prefix-length> {
  next-hop next-hop;
  qualified-next-hop next-hop;
}
```

Hierarchy Level [edit routing-options access-internal]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the parameters for internal access routes.

Options *ip-prefix</prefix-length>*—Specific route prefix that you want to assign to the internal access route.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

route-distinguisher-id

Syntax route-distinguisher-id *ip-address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Automatically assign a route distinguisher to the routing instance.

If you configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**.

NOTE: To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the **route-distinguisher** statement is different from the first half of the route distinguisher obtained by configuring the **route-distinguisher-id** statement.

Options *ip-address*—Address for routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring BGP Route Target Filtering for VPNs*
- *Configuring Routing Instances on PE Routers in VPNs*

route-record

Syntax route-record;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Export the AS path and routing information to the traffic sampling process.

Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Flow Aggregation*
- *Junos OS Services Interfaces Library for Routing Devices*

router-id

Syntax

router-id *address*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify the routing device's IP address.

The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

NOTE: **We strongly recommend that you configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.**

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.

Options

address—IP address of the routing device.

Default: Address of the first interface encountered by Junos OS

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring External BGP Peering*
- *Examples: Configuring Internal BGP Peering*

routing-options

Syntax routing-options { ... }

Hierarchy Level [edit],
[edit logical-systems *logical-system-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure protocol-independent routing properties.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Protocol-Independent Routing Properties Feature Guide for Routing Devices*

scope

Syntax

```
scope scope-name {
  interface [ interface-names ];
  prefix destination-prefix;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure multicast scoping.

Options ***scope-name***—Name of the multicast scope.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Creating a Named Scope for Multicast Scoping*

source-address (Routing Options)

Syntax source-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*,
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-options dynamic-tunnels *tunnel-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the source address for the generic routing encapsulation (GRE) tunnels. The source address specifies the address used as the source for the local tunnel endpoint. This address can be any local address on the router, typically the router ID or the loopback address.

Options ***address***—Name of the source address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring GRE Tunnels for Layer 3 VPNs*

source-routing

Syntax

```
source-routing {
  (ip | ipv6)
}
```

Hierarchy Level [edit routing-options]

Release Information Statement for IPv6 introduced in Junos OS Release 8.2.
Statement for IPv4 introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Enable source routing.

Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.

NOTE: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Filter-Based Forwarding on the Source Address*

ssm-groups

Syntax ssm-groups [*ip-addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure source-specific multicast (SSM) groups.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.

Options *ip-addresses*—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

static (Routing Options)

Syntax

```
static {
  defaults {
    static-options;
  }
  rib-group group-name;
  route destination-prefix {
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      local-address ip-address;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      minimum-receive-ttl number;
      multiplier number;
      neighbor address;
      no-adaptation;
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      version (1 | automatic);
    }
    next-hop address;
    next-hop options;
    qualified-next-hop address {
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          key-chain key-chain-name;
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
```

```
            metric metric;
            preference preference;
        }
        static-options;
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos 9.6.
Support for BFD authentication introduced in Junos 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure static routes to be installed in the routing table. You can specify any number of routes within a single **static** statement, and you can specify any number of **static** options in the configuration.

Options **defaults**—(Optional) Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the **static** statement.

NOTE: Specifying the global static route options does not create default routes. These options only set default attributes inherited by all newly created static routes.

route—Configure individual static routes. In this part of the **static** statement, you optionally can configure static route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the **static** statement.

- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

When you configure an individual static route in the **route** part of the **static** statement, specify the destination of the route (in **route** ***destination-prefix***) in one of the following ways:

- ***network/mask-length***, where ***network*** is the network portion of the IP address and ***mask-length*** is the destination prefix length.
- **default** if this is the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.

NOTE: IPv4 packets with a destination of 0.0.0.0 (the obsoleted limited broadcast address) and IPv6 packets with a destination of 0::0 are discarded by default. To forward traffic destined to these addresses, you can add a static route to 0.0.0.0/32 for IPv4 or 0::0/128 for IPv6.

- ***nsap-prefix***—***nsap-prefix*** is the network service access point (NSAP) address for ISO.
- **next-hop** ***address***—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).

  IPv4 or IPv6 address of the next hop to the destination, specified as:

  - IPv4 or IPv6 address of the next hop
  - Interface name (for point-to-point interfaces only)
  - ***address*** or ***interface-name*** to specify an IP address of a multipoint interface or an interface name of a point-to-point interface.

  NOTE: If an interface becomes unavailable, all configured static routes on that interface are withdrawn from the routing table.

NOTE: Load balancing is not supported on management and internal Ethernet (fxo) interfaces because this type of interface cannot handle the routing process. On fxp interfaces, you cannot configure multiple next hops and enable load balancing.

next-hop options—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.
- ***iso-net***—Reach the next-hop routing device by specifying an ISO NSAP.

- **next-table** ***routing-table-name***—Name of the next routing table to the destination.

  If you use the **next-table** action, the configuration must include a term qualifier that specifies a different table than the one specified in the **next-table** action. In other words, the term qualifier in the **from** statement must exclude the table in the **next-table** action. In the following example, the first term contains **rib vrf-customer2.inet.0** as a matching condition. The action specifies a next-hop in a different routing table, **vrf-customer1.inet.0**. The second term does the opposite by using **rib vrf-customer1.inet.0** in the match condition and **vrf-customer2.inet.0** In the **next-table** action.

```
term 1 {
  from {
    protocol bgp;
    rib vrf-customer2.inet.0;
    community customer;
  }
  then {
    next-hop next-table vrf-customer1.inet.0;
  }
}
term 2 {
  from {
    protocol bgp;
    rib vrf-customer1.inet.0;
    community customer;
  }
  then {
    next-hop next-table vrf-customer2.inet.0;
  }
}
```

NOTE: Within a routing instance, you cannot configure a static route with the **next-table inet.0** statement if any static route in the main routing instance is already configured with the **next-table** statement to point to the **inet.0** routing table of the routing instance. For example, if you configure on the main routing instance a static route **192.168.88.88/32** with the **next-table test.inet.0** statement and the routing instance test is also configured with a static route **192.168.88.88/32** with the **next-table inet.0** statement, the commit operation fails. Instead, you must configure a routing table group both on the main instance and on the routing instance, which enables you to install the static route into both routing tables.

- **receive**—Install a route for this next-hop destination into the routing table.

  The **receive** option forces the packet to be sent to the Routing Engine.

  The **receive** option can be useful in the following cases:

  - For receiving MPLS packets destined to a VRF instance's loopback address
  - For receiving packets on a link's subnet address, with zeros in the host portion of the address

- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in ***static-options***. Each of the options is explained separately.

- **(active | passive)**;
- **as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number in-address*>**;
- **community [*community-ids*]**;
- **(install | no-install)**;
- **(metric | metric2 | metric3 | metric4) *value* <type *type*>**;
- **(preference | preference2 | color | color2) *preference* <type *type*>**;
- **(readvertise | no-readvertise)**;
- **(resolve | no-resolve)**;
- **(retain | no-retain)**;
- **tag *string***;

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring Static Routes on page 9

tag (Routing Options)

Syntax

tag *string*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Associate an OSPF tag with a static, aggregate, or generated route.

Default

No OSPF tag strings are associated with static routes.

Options

string—OSPF tag string.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9
- Example: Summarizing Routes Through Route Aggregation on page 82
- Example: Conditionally Generating Static Routes on page 88
- aggregate on page 160
- generate on page 189
- static on page 250

tag (Access)

Syntax tag *tag-number*;

Hierarchy Level [edit routing-options access route *ip-prefix</prefix-length>*]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure a tag for an access route.

Options ***tag-number***—Tag number for the access route.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Examples: Configuring Static Routes on page 9

threshold (Multicast Forwarding Cache)

Syntax

```
threshold {
  log-warning value;
  suppress value <reuse value>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache family (inet | inet6)],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache (inet | inet6)],
[edit routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache family (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the **threshold** statement globally for the multicast forwarding cache or including the **family** statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

To confirm the configured threshold values, use the **show multicast forwarding-cache statistics** command.

Options

reuse ***value***—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the **suppress** value.

Range: 1 through 200,000

suppress ***value***—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the **reuse** value.

Range: 1 through 200,000

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring the Multicast Forwarding Cache*

traceoptions (Routing Options)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options],
[edit routing-options flow],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 7.4.
nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.
nsr-synchronization and **nsr-packet** flags for BFD sessions added in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.
nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.
nsr-synchronization flag for PIM added in Junos OS Release 9.3.
nsr-synchronization flag for MPLS added in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
nsr-synchronization flag for MSDP added in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Define tracing operations that track all routing protocol functionality in the routing device.

To specify more than one tracing operation, include multiple **flag** statements.

Default

If you do not include this statement, no global tracing operations are performed.

Options

Values:

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation
- Example: Tracing Global Routing Protocol Operations on page 152
- *Tracing Nonstop Active Routing Synchronization Events*

unicast-reverse-path

Syntax
unicast-reverse-path (active-paths | feasible-paths);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-instances *routing-instance-name* instance-type *name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information
Statement introduced before Junos OS Release 7.4.
Support for routing instances added in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description
Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.

Options
active-paths—Consider only active paths during the unicast reverse-path check.

feasible-paths—Consider all feasible paths during the unicast reverse-path check.

Default: If you omit the **unicast-reverse-path** statement, only the active paths to a particular destination are considered.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Unicast Reverse-Path-Forwarding Check on page 134
- *Enabling Unicast Reverse-Path Forwarding Check for VPNs*

PART 3

Administration

CHAPTER 9

Operational Commands

show bfd session

Syntax

show bfd session
<brief | detail | extensive | summary>
<address *address*>
<discriminator *discriminator*>
<logical-system (all | *logical-system-name*)>
<prefix *address*>

Syntax (EX Series Switch and QFX Series)

show bfd session
<brief | detail | extensive | summary>
<address *address*>
<discriminator *discriminator*>
<prefix *address*>

Release Information

Command introduced before Junos OS Release 7.4.
Options **discriminator** and **address** introduced in Junos OS Release 8.2.
Option **prefix** introduced in Junos OS Release 9.0.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display information about active Bidirectional Forwarding Detection (BFD) sessions.

Options

none—(Same as **brief**) Display information about active BFD sessions.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

address *address*—(Optional) Display information about the BFD session for the specified neighbor address.

discriminator *discriminator*—(Optional) Display information about the BFD session using the specified local discriminator.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix *address*—(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC).

Required Privilege Level

view

Related Documentation

- *clear bfd session*
- Examples: Configuring BFD for Static Routes on page 99
- *Example: Configuring BFD for OSPF*
- *Example: Configuring BFD for BGP*
- *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*
- *Example: Configuring BFD for IS-IS*

List of Sample Output

show bfd session on page 271
show bfd session brief on page 271
show bfd session detail on page 271
show bfd session detail (with Authentication) on page 271
show bfd session address extensive on page 271
show bfd session extensive on page 272
show bfd session extensive (with Authentication) on page 272
show bfd session summary on page 273

Output Fields

Table 4 on page 267 describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

Table 4: show bfd session Output Fields

Field Name	Field Description	Level of Output
Address	Address on which the BFD session is active.	**brief detail extensive** none
State	State of the BFD session: **Up**, **Down**, **Init** (initializing), or **Failing**.	**brief detail extensive** none
Interface	Interface on which the BFD session is active.	**brief detail extensive** none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	**brief detail extensive** none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	**brief detail extensive** none
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	**detail extensive**
Session up time	How long a BFD session has been established.	**detail extensive**
Client	Protocol for which the BFD session is active: **ISIS**, **OSPF**, or **Static**.	**detail extensive**
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	**brief detail extensive** none
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	**brief detail extensive** none
Authenticate	Indicates that BFD authentication is configured.	**detail extensive**
keychain	Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**

Table 4: show bfd session Output Fields *(continued)*

Field Name	Field Description	Level of Output
algo	BFD authentication algorithm being used for a specific client: **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**
mode	Level of BFD authentication enforcement being used by a specific client: **strict** or **loose**. Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**
Local diagnostic	Local diagnostic information about failing BFD sessions.	**detail extensive**
Remote diagnostic	Remote diagnostic information about failing BFD sessions.	**detail extensive**
Remote state	Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	**detail extensive**
Version	BFD version: **0** or **1**.	**extensive**
Replicated	The **replicated** flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine.	**detail extensive**
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.	**extensive**
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.	**extensive**
Adaptive async TX interval	Transmission interval being used because of adaptation.	**extensive**
RX interval	Minimum required receive interval.	**extensive**
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	**extensive**
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	**extensive**
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	**extensive**
Remote min TX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	**extensive**

Table 4: show bfd session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Threshold transmission interval	Threshold for notification if the transmission interval increases.	**extensive**
Threshold for detection time	Threshold for notification if the detection time increases.	**extensive**
Local discriminator	Authentication code used by the local system to identify that BFD session.	**extensive**
Remote discriminator	Authentication code used by the remote system to identify that BFD session.	**extensive**
Echo mode	Information about the state of echo transmissions on the BFD session.	**extensive**
Prefix	LDP FEC address associated with the BFD session.	All levels
Egress, Destination	Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.	All levels
Remote is control-plane independent	The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.	**extensive**
Authentication	Summary status of BFD authentication: • **status**—**enabled/active** indicates authentication is configured and active. **enabled/inactive** indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. • **keychain**—Name of the security authentication keychain associated with the specified BFD session. • **algo**—BFD authentication algorithm being used: **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**. • **mode**—Level of BFD authentication enforcement: **strict** or **loose**. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. This information is only shown if BFD authentication is configured.	**extensive**
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	**detail extensive**
sessions	Total number of active BFD sessions.	All levels
clients	Total number of clients that are hosting active BFD sessions.	All levels

Table 4: show bfd session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	All levels
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	All levels
Multi-hop, min-recv-TTL	Minimum time to live (TTL) accepted if the session is configured for multihop.	**extensive**
route table	Route table used if the session is configured for multihop.	**extensive**
local address	Local address of the source used if the session is configured for multihop.	**extensive**

Sample Output

show bfd session

```
user@host> show bfd session
                                                      Transmit
Address                  State     Interface     Detect Time  Interval  Multiplier
10.9.1.33                Up        so-7/1/0.0          0.600     0.200      3
10.9.1.29                Up        ge-4/0/0.0          0.600     0.200      3

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command. For sample output, see show bfd session on page 271.

show bfd session detail

```
user@host> show bfd session detail
                                                      Transmit
Address                  State     Interface     Detect Time  Interval  Multiplier
10.9.1.33                Up        so-7/1/0.0          0.600     0.200      3
  Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:34
  Local diagnostic None, remote diagnostic None
  Remote state Up, version 1
  Replicated
10.9.1.29                Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:29, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session detail (with Authentication)

```
user@host>  show bfd session detail
                                                      Transmit
Address                  State     Interface     Detect Time  Interval  Multiplier
10.9.1.33                Up        so-7/1/0.0          0.600     0.200      3
  Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate
  Session up time 3d 00:34
  Local diagnostic None, remote diagnostic None
  Remote state Up, version 1
  Replicated
10.9.1.29                Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:29, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session address extensive

```
user@host> show bfd session 10.255.245.212 extensive
                                                      Transmit
Address                  State     Interface     Detect Time  Interval  Multiplier
10.255.245.212           Up                            1.200     0.400      3
  Client Static, TX interval 0.400, RX interval 0.400, multiplier 3
```

```
  Session up time 00:17:03, previous down time 00:00:14
  Local diagnostic CtlExpire, remote diagnostic NbrSignal
  Remote state Up, version 1
  Replicated
  Min async interval 0.400, min slow interval 1.000
  Adaptive async tx interval 0.400, rx interval 0.400
  Local min tx interval 0.400, min rx interval 0.400, multiplier 3
  Remote min tx interval 0.400, min rx interval 0.400, multiplier 3
  Threshold transmission interval 0.000, Threshold for detection time 0.000
  Local discriminator 6, remote discriminator 16
  Echo mode disabled/inactive
  Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205


1 sessions, 1 clients
Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps
```

show bfd session extensive

```
user@host> show bfd session extensive
10.31.1.2                Up        ge-2/1/8.0     0.030     0.010        3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:13
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 12, remote discriminator 4
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x201

                                                         Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.31.2.2                Up        ge-2/1/4.0     0.030     0.010        3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:14
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 13, remote discriminator 5
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x202

2 sessions, 2 clients
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps
```

show bfd session extensive (with Authentication)

```
user@host> show bfd session extensive
                                                         Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.208.26           Up        so-1/0/0.0     2.400     0.800        10
 Client Static, TX interval 0.600, RX interval 0.600, Authenticate
       keychain bfd, algo keyed-md5, mode loose
 Session up time 00:18:07
```



```
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.600, min slow interval 1.000
 Adaptive async TX interval 0.600, RX interval 0.600
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10
 Remote min TX interval 0.800, min RX interval 0.800, multiplier 3
 Local discriminator 2, remote discriminator 3
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose

1 sessions, 1 clients
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps
```

show bfd session summary

```
user@host> show bfd session summary
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show route

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

Required Privilege Level

view

Related Documentation

- *Example: Configuring RIP*
- *Example: Configuring RIPng*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*

List of Sample Output

show route on page 278
show route on page 278
show route destination-prefix on page 278
show route extensive on page 278

Output Fields Table 5 on page 275 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 5: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).

destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:
	• *MPLS-label* (for example, 80001).
	• *interface-name* (for example, ge-1/0/2).
	• *neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id* :*source* (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96):
	  ◦ *neighbor-address*—Address of the neighbor.
	  ◦ *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ *encapsulation type*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ *vc-id*—Virtual circuit identifier.
	  ◦ *source*—Source of the advertisement: **Local** or **Remote**.

[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Table 5: show route Output Fields *(continued)*

Field Name	Field Description
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.

Table 5: show route Output Fields *(continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```

show route

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 70.0.0.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

70.0.0.0/24        @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                   #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route destination-prefix

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
          PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1
```



```
Next hop type: Indirect
Address: 0x92455b8
Next-hop reference count: 2
Source: 10.0.0.30
Protocol next hop: 10.0.0.40
Indirect next hop: 2 no-forward
State: <Active Int Ext>
          Local AS: 65500 Peer AS: 65500
Age: 3  Metric2: 1
Task: BGP_65500.10.0.0.30+179
Announcement bits (2): 0-PIM.v1 1-mvpn global task
AS path: I (Originator) Cluster list:  10.0.0.30
AS path:  Originator ID: 10.0.0.40
Communities: target:65520:100
Import Accepted
Localpref: 100
Router ID: 10.0.0.30
Primary Routing Table bgp.mvpn.0
Indirect next hops: 1
     Protocol next hop: 10.0.0.40 Metric: 1
     Indirect next hop: 2 no-forward
     Indirect path forwarding next hops: 1
           Next hop type: Router
           Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
     10.0.0.40/32 Originating RIB: inet.3
           Metric: 1                       Node path count: 1
           Forwarding nexthops: 1
              Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route active-path

Syntax

show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route active-path
<brief | detail | extensive | terse>

Release Information

Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options

none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route active-path on page 281
show route active-path brief on page 281
show route active-path detail on page 281
show route active-path extensive on page 282
show route active-path terse on page 284

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route active-path

user@host> **show route active-path**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24       *[Direct/0] 00:18:36
                    > via so-2/1/3.0
100.1.2.2/32       *[Local/0] 00:18:41
                      Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                      Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see show route active-path on page 281.

show route active-path detail

user@host> **show route active-path detail**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 100.1.2.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

100.1.2.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
```



```
                Next hop: via so-2/1/3.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

100.1.2.2/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

user@host> **show route active-path extensive**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
```



```
                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 100.1.2.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

100.1.2.2/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
```

```
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path terse

user@host> **show route active-path terse**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                       >lo0.0
* 10.255.71.50/32    I  15        10             >100.1.2.1
* 100.1.2.0/24       D   0                       >so-2/1/3.0
* 100.1.2.2/32       L   0                        Local
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.70.19/32   L   0                        Local
```


show route all

Syntax

show route all
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route all

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display information about all routes in all routing tables, including private, or internal, tables.

Options

none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route all on page 286

Output Fields

In Junos OS Release 9.5 and later, only the output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop
```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
800017             *[VPLS/7] 1d 13:54:49
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                    > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```


show route best

Syntax

show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route best on page 288
show route best detail on page 288
show route best extensive on page 289
show route best terse on page 289

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
```



```
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
                Age: 2d 1:44:20
                Task: IF
                AS path: I
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 2d 1:44:20
                Task: IF
                AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see show route best detail on page 288.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                  so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                       >fxp2.0
                     D   0                       >fxp1.0
```

show route brief

Syntax
show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route brief
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display brief information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route brief on page 291

Output Fields
For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
224.0.0.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```


show route detail

Syntax

show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route detail
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display detailed information about the active entries in the routing tables.

Options

none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route detail on page 301
show route detail (with BGP Multipath) on page 306
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 306

Output Fields

Table 6 on page 292 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 6: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Table 6: show route detail Output Fields *(continued)*

Field Name	Field Description

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 7 on page 296.

Table 6: show route detail Output Fields *(continued)*

Field Name	Field Description
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 8 on page 298.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Table 6: show route detail Output Fields *(continued)*

Field Name	Field Description
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 9 on page 300 for all possible values for this field.

Table 6: show route detail Output Fields *(continued)*

Field Name	Field Description
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 7 on page 296 describes all possible values for the Next-hop Types output field.

Table 7: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.

Table 7: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 8 on page 298 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 8: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).

Table 8: State Output Field Values *(continued)*

Value	Description
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 9 on page 300 describes the possible values for the Communities output field.

Table 9: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
```

```
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```



```
47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
        *RSVP   Preference: 7
                Next hop type: Flood
                Next-hop reference count: 130
                Flood nexthop branches exceed maximum
                Address: 0x8ea65d0
...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
```

```
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
```



```
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 1:25:49    Metric2: 1
                AIGP 210
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...
```

```
l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route detail (with BGP Multipath)

user@host> **show route detail**

```
10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:      1 Peer AS:     2
                Age: 5:04:43
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 2 I
                Accepted Multipath
                Localpref: 100
                Router ID: 1.1.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS:      1 Peer AS:     2
                Age: 5:04:43
                Task: BGP_2.10.1.1.6+58198
                AS path: 2 I
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 1.1.1.3
```

show route label detail (Multipoint LDP Inband Signaling for

user@host> **show route label 299872 detail**

```
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
```


Point-to-Multipoint LSPs)

```
                    Next hop type: Flood
                    Next-hop reference count: 3
                    Address: 0x9097d90
                    Next hop: via vt-0/1/0.1
                    Next-hop index: 661
                    Label operation: Pop
                    Address: 0x9172130
                    Next hop: via so-0/0/3.0
                    Next-hop index: 654
                    Label operation: Swap 299872
                    State: **Active Int>
                    Local AS:  1001
                    Age: 8:20       Metric: 1
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2
```

show route exact

Syntax

show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display only the routes that exactly match the specified address or range of addresses.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route exact on page 309
show route exact detail on page 309
show route exact extensive on page 309
show route exact terse on page 309

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route exact

user@host> **show route exact 207.17.136.0/24**

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                    > to 192.168.71.254 via fxp0.0
```

show route exact detail

user@host> **show route exact 207.17.136.0/24 detail**

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

user@host> **show route exact 207.17.136.0/24 extensive**

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

user@host> **show route exact 207.17.136.0/24 terse**

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                       >192.168.71.254
```

show route export

Syntax

show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options

none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route export inet** command).

Required Privilege Level

view

List of Sample Output

show route export on page 311
show route export detail on page 311
show route export instance detail on page 311

Output Fields

Table 10 on page 310 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

Table 10: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none

Table 10: show route export Output Fields *(continued)*

Field Name	Field Description	Level of Output
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates *instance*.inet.0. • **multicast**—Indicates *instance*.inet.2. • **unicast multicast**—Indicates *instance*.inet.0 and *instance*.inet.2.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

Sample Output

show route export

```
user@host> show route export
Table                            Export           Routes
inet.0                           N                     0
black.inet.0                     Y                     3
red.inet.0                       Y                     4
```

show route export detail

```
user@host>  show route export detail
inet.0                                          Routes:        0
black.inet.0                                    Routes:        3
  Import: [ inet.0 ]
red.inet.0                                      Routes:        4
  Import: [ inet.0 ]
```

show route export instance detail

```
user@host>  show route export instance detail
Instance: master                        Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                         Type: non-forwarding
Instance: red                           Type: non-forwarding
```

show route hidden

Syntax

show route hidden
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display only hidden route information. A hidden route is unusable, even if it is the best path.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Understanding Hidden Routes*

List of Sample Output

show route hidden on page 313
show route hidden detail on page 313
show route hidden extensive on page 314
show route hidden terse on page 314

Output Fields

For information about output fields, see the output field table for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route hidden

user@host> **show route hidden**

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32        [Direct/0] 04:26:38
                    > via lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.1.0/24        [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.80.4/30       [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: I
                      Unusable
...
```

show route hidden detail

user@host> **show route hidden detail**

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
```



```
                    Primary Routing Table bgp.l3vpn.0
...
```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see show route hidden detail on page 313.

show route hidden terse

```
user@host> show route hidden terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                       >lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.5.5.5/32        B 170        100             Unusable        100 I
  10.12.1.0/24       B 170        100             Unusable        100 I
  10.12.80.4/30      B 170        100             Unusable        I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.4.4.4:4:10.5.5.5/32
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.1.0/24
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.80.4/30
                     B 170        100             Unusable        I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path

Syntax

show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route inactive-path
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options

none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route inactive-path on page 316
show route inactive-path detail on page 316
show route inactive-path extensive on page 317
show route inactive-path terse on page 317

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
```



```
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 4:40:52
                Task: IF
                AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-path detail on page 316.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                       >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route instance

Syntax

show route instance
<brief | detail | summary>
<*instance-name*>
<logical-system (all | *logical-system-name*)>
<operational>

Syntax (EX Series Switches and QFX Series)

show route instance
<brief | detail | summary>
<*instance-name*>
<operational>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display routing instance information.

Options

none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show route instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level

view

List of Sample Output

show route instance on page 321
show route instance detail (Graceful Restart Complete) on page 321
show route instance detail (Graceful Restart Incomplete) on page 323
show route instance detail (VPLS Routing Instance) on page 324
show route instance operational on page 325
show route instance summary on page 325

Output Fields

Table 11 on page 319 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 11: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels

Table 11: show route instance Output Fields *(continued)*

Field Name	Field Description	Level of Output
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, **virtual-router**, or **vrf**.	All levels
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is **300**.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	**brief detail** none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                     Active/holddown/hidden
master               forwarding
         inet.0                                          16/0/1
         iso.0                                           1/0/0
         mpls.0                                          0/0/0
         inet6.0                                         2/0/0
         l2circuit.0                                     0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                     12/0/0
         __juniper_private1__.inet6.0                    1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                 : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                 : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```



```
    BGP-L.mpls.0           : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn            State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf              State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf              State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf              State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf              State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
```



```
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                 : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                 : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
```



```
LDP:
  Router ID: 10.69.105.1
  Type: vrf                State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf                State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf                State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf                State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls               State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
```



```
  Vrf-export-target: [ target:300:1 ]
  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance            Type               Primary rib            Active/holddown/hidden
master              forwarding
                                       inet.0                 15/0/1
                                       iso.0                  1/0/0
                                       mpls.0                 35/0/0
                                       l3vpn.0                0/0/0
                                       inet6.0                2/0/0
                                       l2vpn.0                0/0/0
                                       l2circuit.0            0/0/0
BGP-INET            vrf
                                       BGP-INET.inet.0        5/0/0
                                       BGP-INET.iso.0         0/0/0
                                       BGP-INET.inet6.0       0/0/0
BGP-L               vrf
                                       BGP-L.inet.0           5/0/0
                                       BGP-L.iso.0            0/0/0
                                       BGP-L.mpls.0           4/0/0
                                       BGP-L.inet6.0          0/0/0
L2VPN               l2vpn
                                       L2VPN.inet.0           0/0/0
                                       L2VPN.iso.0            0/0/0
                                       L2VPN.inet6.0          0/0/0
                                       L2VPN.l2vpn.0          2/0/0
LDP                 vrf
                                       LDP.inet.0             4/0/0
                                       LDP.iso.0              0/0/0
                                       LDP.mpls.0             0/0/0
                                       LDP.inet6.0            0/0/0
                                       LDP.l2circuit.0        0/0/0
OSPF                vrf
                                       OSPF.inet.0            7/0/0
                                       OSPF.iso.0             0/0/0
                                       OSPF.inet6.0           0/0/0
RIP                 vrf
                                       RIP.inet.0             6/0/0
                                       RIP.iso.0              0/0/0
                                       RIP.inet6.0            0/0/0
STATIC              vrf
                                       STATIC.inet.0          4/0/0
                                       STATIC.iso.0           0/0/0
                                       STATIC.inet6.0         0/0/0
```


show route label-switched-path

Syntax
show route label-switched-path *path-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route label-switched-path *path-name*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description
Display the routes used in an MPLS label-switched path (LSP).

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

path-name—LSP tunnel name.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route label-switched-path on page 327

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                      to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route martians

Syntax

show route martians
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches)

show route martians
<table *routing-table-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the martian (invalid and ignored) entries associated with each routing table.

Options

none—Display standard information about route martians for all routing tables.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

table ***routing-table-name***—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route martians table inet** command).

Required Privilege Level

view

Related Documentation

- Example: Configuring Martian Addresses on page 145

List of Sample Output

show route martians on page 329

Output Fields

Table 12 on page 328 lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

Table 12: show route martians Output Fields

Field Name	Field Description
table-name	Name of the route table in which the route martians reside.
destination-prefix	Route destination.
match value	Route match parameter.
status	Status of the route: **allowed** or **disallowed**.

Sample Output

show route martians

```
user@host> show route martians

inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.1:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed

inet.2:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.3:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed
...

inet6.0:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.1:
             ::1/128 exact -- disallowed

inet6.2:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.3:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
...
```


show route next-hop

Syntax

show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that are being sent to the specified next-hop address.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level

view

List of Sample Output

show route next-hop on page 331
show route next-hop detail on page 331
show route next-hop extensive on page 333
show route next-hop terse on page 335

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route next-hop

user@host> **show route next-hop 192.168.71.254**

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32  *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

user@host> **show route next-hop 192.168.71.254 detail**

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
```

```
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I
```

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```



```
red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

user@host> **show route next-hop 192.168.71.254 terse**

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route protocol

Syntax

show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
opsf2 and **ospf3** options introduced in Junos OS Release 9.2.
opsf2 and **ospf3** options introduced in Junos OS Release 9.2 for EX Series switches.
flow option introduced in Junos OS Release 10.0.
flow option introduced in Junos OS Release 10.0 for EX Series switches.

Description

Display the route entries in the routing table that were learned from a particular protocol.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **arp**—Route learned through the Address Resolution Protocol
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route
- **frr**—Precomputed protection route or backup route used when a link goes down
- **isis**—Intermediate System-to-Intermediate System
- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit
- **l2vpn**—Layer 2 virtual private network

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level view

List of Sample Output

show route protocol access on page 339
show route protocol access-internal extensive on page 339
show route protocol arp on page 339
show route protocol bgp on page 340
show route protocol bgp detail on page 340
show route protocol bgp extensive on page 340
show route protocol bgp terse on page 341
show route protocol direct on page 341
show route protocol frr on page 341
show route protocol l2circuit detail on page 342
show route protocol l2vpn extensive on page 343
show route protocol ldp on page 343
show route protocol ldp extensive on page 344
show route protocol ospf (Layer 3 VPN) on page 345
show route protocol ospf detail on page 345
show route protocol rip on page 346
show route protocol rip detail on page 346
show route protocol ripng table inet6 on page 346
show route protocol static detail on page 346

Output Fields For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the *show route extensive* command, or the **show route terse** command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24  (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:42:42        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:44:45        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
```



```
        Accepted
        Localpref: 100
        Router ID: 207.17.136.192
        Indirect next hops: 1
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                Indirect path forwarding next hops: 1
                        Next hop type: Router
                        Next hop: 192.168.167.254 via fxp0.0
                192.168.0.0/16 Originating RIB: inet.0
                  Node path count: 1
                  Forwarding nexthops: 1
                        Nexthop: 192.168.167.254 via fxp0.0
```

show route protocol bgp terse

user@host> **show route protocol bgp 192.168.64.0/21 terse**

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >100.1.3.2       10023 21 I
```

show route protocol direct

user@host> **show route protocol direct**

```
inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24         *[Direct/0] 17w0d 10:31:49
                    > via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
                    > via lo0.0
30.30.30.0/24      *[Direct/0] 17w0d 23:06:26
                    > via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
                    > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0
```

show route protocol frr

user@host> **show route protocol frr**

```
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol l2circuit detail

user@host> **show route protocol l2circuit detail**

```
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
```



```
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop      Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
                Protocol next hop: 10.255.14.220
                Push 800000 Offset: -4
                 Indirect next hop: 85142a0 288
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:69:1 Layer2-info: encaps:PPP,
                control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0
```

```
private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

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

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
```



```
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
224.0.0.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
224.0.0.5/32       *[OSPF/10] 20:26:20, metric 1
```

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
```

```
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32       *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
```

```
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.13.10.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

show route resolution

Syntax

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Syntax (EX Series Switches)

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options

none—Display standard information about all entries in the next-hop resolution database.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

index *index*—(Optional) Show the index of the resolution tree.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix network/destination-prefix—(Optional) Display database entries for the specified address.

table *routing-table-name*—(Optional) Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled.

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

List of Sample Output

show route resolution detail on page 350
show route resolution summary on page 350
show route resolution unresolved on page 350

Output Fields Table 13 on page 349 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 13: show route resolution Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
 Contributing routing tables: :voice.inet.0 :voice.inet.3
 Tree Index: 2, Nodes 2, Reference Count 1
 Contributing routing tables: inet.3
 Tree Index: 3, Nodes 43, Reference Count 1
 Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
        Protocol Nexthop: 10.255.71.238 Push 800000
        Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
        Protocol Nexthop: 10.255.70.103 Push 800008
        Indirect nexthop: 0 -
Tree Index 2
```

show route summary

Syntax show route summary
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches) show route summary

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display summary statistics about the entries in the routing table.

CPU utilization might increase while the device learns routes. We recommend that you use the **show route summary** command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a "timeout communicating with routing daemon" error when using the **show route summary** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).

Options **none**—Display summary statistics about the entries in the routing table.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

table ***routing-table-name***—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route summary table inet** command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.

Required Privilege Level view

List of Sample Output show route summary on page 353
show route summary table on page 353
show route summary table (with Route Limits Configured for the Routing Table) on page 354

Output Fields Table 14 on page 351 lists the output fields for the **show route summary** command. Output fields are listed in the approximate order in which they appear.

Table 14: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.
routing-table-name	Name of the routing table (for example, **inet.0**).

Table 14: show route summary Output Fields *(continued)*

Field Name	Field Description
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.
Limit/Threshold	Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. • **destinations**—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • **routes**—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**.

Sample Output

show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
              Direct:      6 routes,      5 active
               Local:      4 routes,      4 active
                OSPF:      5 routes,      4 active
              Static:      7 routes,      7 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
                RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
                MPLS:      3 routes,      3 active
                VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      2 routes,      2 active
                 PIM:      2 routes,      2 active
                 MLD:      1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      4 routes,      4 active
```

show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
              Direct:      6 routes,      5 active
               Local:      9 routes,      9 active
                OSPF:      3 routes,      1 active
              Static:     13 routes,     13 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active
```

```
inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
               Local:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
              Direct:      2 routes,      2 active
               Local:      1 routes,      1 active
                OSPF:      4 routes,      3 active
                 BGP:      4 routes,      4 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active
```


show route table

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

Related Documentation

- show route summary on page 351

List of Sample Output

show route table inetflow detail on page 364

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the **show route extensive** command, or the show route terse command.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.245.12:1
          Source: 10.255.245.12
          Next hop: 192.168.208.66 via fe-0/0/0.0, selected
          Label operation: Push 182449
          Protocol next hop: 10.255.245.12
          Push 182449
          Indirect next hop: 863a630 297
          State: <Active Int Ext>
          Local AS:     35 Peer AS:    35
          Age: 12:19      Metric2: 1
          Task: BGP_35.10.255.245.12+179
          Announcement bits (1): 0-BGP.0.0.0.0+179
          AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

          Communities: 2914:420 target:11111:1 origin:56:78
          VPN Label: 182449
          Localpref: 100
           Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.245.12:1
          Source: 10.255.245.12
          Next hop: 192.168.208.66 via fe-0/0/0.0, selected
          Label operation: Push 182465
          Protocol next hop: 10.255.245.12
          Push 182465
```

```
         Indirect next hop: 863a8f0 305
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 86bd210 330
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

         6496 I
         Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
         VPN Label: 182465
         Localpref: 100
         Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 86bd210 330
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

         6496 I
         Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
         VPN Label: 182465
         Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route

```
user@host> show route table bgp.rtarget.o
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


Target Filtering Is Configured)

```
100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                     Local
12.12.12.21/32     *[Local/0] 00:51:57
                     Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                     Local
13.13.13.21/32     *[Local/0] 00:51:58
                     Local
13.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                     Local
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```


show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 65002 Peer AS: 65000
                Age: 4
                Task: BGP_65000.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 65000 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 65002
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```


show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:45:09, metric 1
                      Receive
1                  *[MPLS/0] 00:45:09, metric 1
                      Receive
2                  *[MPLS/0] 00:45:09, metric 1
                      Receive
100000             *[L2VPN/7] 00:43:04
                    > via so-0/1/0.1, Pop
100001             *[L2VPN/7] 00:43:03
                    > via so-0/1/0.2, Pop       Offset: 4
100002             *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100002(S=0)        *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100003             *[LDP/9] 00:43:22, metric 1
                    > via so-0/1/2.0, Swap 100002
                      via so-0/1/3.0, Swap 100002
100004             *[LDP/9] 00:43:16, metric 1
                      via so-0/1/2.0, Swap 100049
                    > via so-0/1/3.0, Swap 100049
so-0/1/0.1         *[L2VPN/7] 00:43:04
                    > via so-0/1/2.0, Push 100001, Push 100049(top)
                      via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2         *[L2VPN/7] 00:43:03
                      via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                    > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4
```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
```



```
13                   *[MPLS/0] 00:37:31, metric 1
                        Receive
300352               *[RSVP/7/1] 00:08:00, metric 1
                      > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)          *[RSVP/7/1] 00:08:00, metric 1
                      > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384               *[RSVP/7/2] 00:05:20, metric 1
                      > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)          *[RSVP/7/2] 00:05:20, metric 1
                      > to 8.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                     *[VPN/7] 05:48:27
                       Discard
192.168.24.1:1:2:1/96
                     *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                        AS path: I
                      > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                     *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                        AS path: I
                      > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                     *[MVPN/70] 01:23:05, metric2 1
                        Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                     *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                        AS path: I
                      > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                     *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
                        AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```


show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                            AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
```

```
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65535:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
6:10.255.2.203:65535:65000:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 65002 Peer AS: 65000
                Age: 4
                Task: BGP_65000.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 65000 I
                Communities: traffic-rate:0:0
```



```
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 65002
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

```
user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
1.1.1.4:100:1.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
1.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard
```

```
user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
22.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
                Source: 2.2.0.0
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
```

```
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Protocol next hop: 2.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     2 Peer AS:     2
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: BGP_2.2.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 2.2.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 2.2.0.0 Metric: 35
                        Push 16
                        Composite next hop: 0x25805988 - INH Session ID: 0x193c
                        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.1.1.1 via ge-1/1/9.0
                                Session Id: 0x17d8
                        2.2.0.0/32 Originating RIB: inet.3
                          Metric: 35                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.1.1.1 via ge-1/1/9.0
         BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x9347028
                Next-hop reference count: 3
                Source: 2.3.0.0
                Next hop type: Router, Next hop index: 702
                Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                Label operation: Push 634278
                Label TTL action: prop-ttl
                Session Id: 0x17d9
                Protocol next hop: 2.3.0.0
                Push 16
                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                Inactive reason: Not Best in its group - IGP metric
                Local AS:     2 Peer AS:     2
                Age: 3:34       Metric2: 70
                Validation State: unverified
                Task: BGP_2.2.3.0.0+32805
                Announcement bits (2): 0-KRT 1-BGP_RT_Background
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 2.3.0.0
```



```
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 2.3.0.0 Metric: 70
                        Push 16
                        Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                        Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.1.4.2 via ge-1/0/0.0
                                Session Id: 0x17d9
                        2.3.0.0/32 Originating RIB: inet.3
                          Metric: 70                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.1.4.2 via ge-1/0/0.0
        #Multipath Preference: 255
                Next hop type: Indirect
                Address: 0x24afca30
                Next-hop reference count: 1
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Next hop type: Router, Next hop index: 702
                Next hop: 10.1.4.2 via ge-1/0/0.0
                Label operation: Push 634278
                Label TTL action: prop-ttl
                Session Id: 0x17d9
                Protocol next hop: 2.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
               Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                Protocol next hop: 2.3.0.0
                Push 16
                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
               Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
 0x4000
                State: <ForwardingOnly Int Ext>
                Inactive reason: Forwarding use only
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: RT
                AS path: I
                Communities: target:2:1
```


show route terse

Syntax

show route terse
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route terse

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display a high-level summary of the routes in the routing table.

NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the **show route extensive** command.

Options

none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route terse on page 370

Output Fields

Table 15 on page 368 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 15: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Table 15: show route terse Output Fields *(continued)*

Field Name	Field Description
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • **?**—Not evaluated. Indicates that the route was not learned through BGP. • **I**—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Table 15: show route terse Output Fields *(continued)*

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 1.0.1.1/32         O  10          1            >10.0.0.2
  ?                    B 170        100                             I
  unverified                                       >10.0.0.2
* ? 1.1.1.1/32         D   0                       >lo0.2
* V 2.2.0.2/32         B 170        110                             200 I
  valid                                            >10.0.0.2
* ? 10.0.0.0/30        D   0                       >lt-1/2/0.1
  ?                    B 170        100                             I
  unverified                                       >10.0.0.2
* ? 10.0.0.1/32        L   0                        Local
* ? 10.0.0.4/30        B 170        100                             I
  unverified                                       >10.0.0.2
* ? 10.0.0.8/30        B 170        100                             I
  unverified                                       >10.0.0.2
* I 172.16.1.1/32      B 170         90                             200 I
  invalid                                          >10.0.0.2
* N 192.168.2.3/32     B 170        100                             200 I
  unknown                                          >10.0.0.2
* ? 224.0.0.5/32       O  10          1             MultiRecv
```

PART 4

Troubleshooting

CHAPTER 10

Network Troubleshooting

Working with Problems on Your Network

Problem This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

Solution

Table 16: Checklist for Working with Problems on Your Network

<table>
<tr><th>Tasks</th><th>Command or Action</th></tr>
<tr><td colspan="2">“Isolate a Broken Network Connection” on page 373</td></tr>
<tr><td>1. Identify the Symptoms on page 375</td><td>ping (ip-address | hostname)
show route (ip-address | hostname)
traceroute (ip-address | hostname)</td></tr>
<tr><td>2. Isolate the Causes on page 376</td><td>show < configuration | interfaces | protocols | route ></td></tr>
<tr><td>3. Take Appropriate Action on page 376</td><td>[edit]
delete routing options static route destination-prefix
commit and-quit
show route destination-prefix</td></tr>
<tr><td>4. Evaluate the Solution on page 377</td><td>show route (ip-address | hostname)
ping (ip-address | hostname) count 3
traceroute (ip-address | hostname)</td></tr>
</table>

Isolate a Broken Network Connection

Purpose By applying the standard four-step process illustrated in Figure 17 on page 374, you can isolate a failed node in the network.

Figure 17: Process for Diagnosing Problems in Your Network

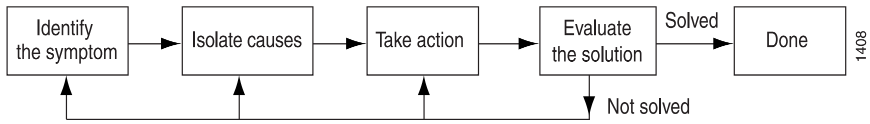

Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic in your approach to the maintenance and monitoring of your network. To prepare for problems on your network, understand how the network functions under normal conditions, have records of baseline network activity, and carefully observe the behavior of your network during a problem situation.

Figure 18 on page 374 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 18: Network with a Problem

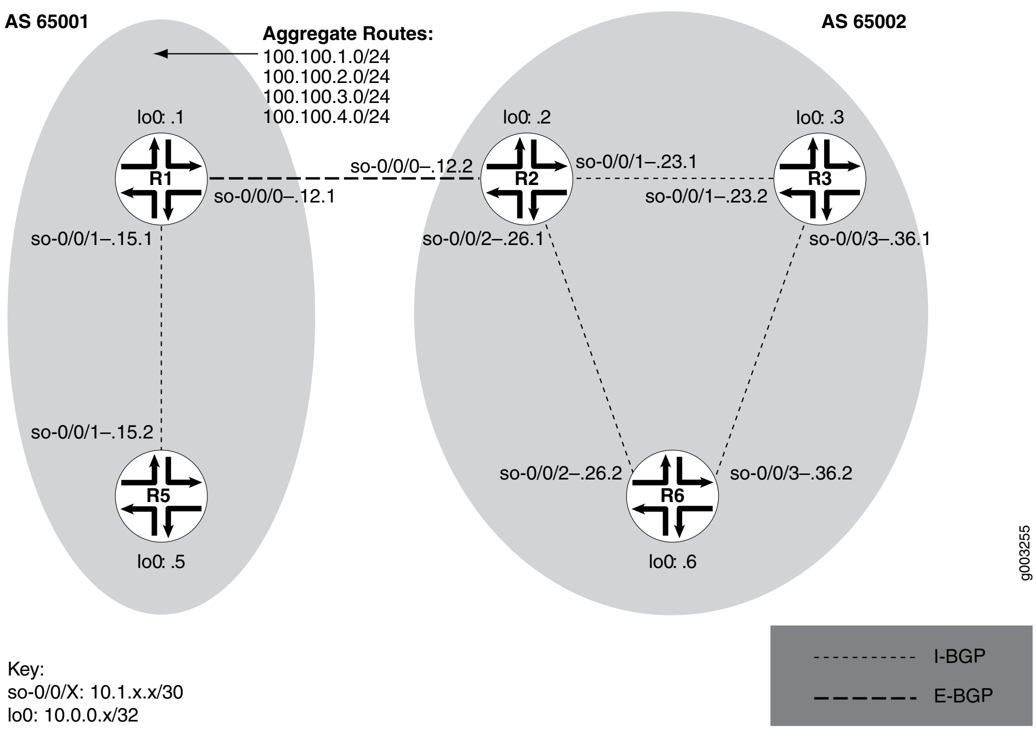

The network in Figure 18 on page 374 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The problem in this network is that **R6** does not have access to **R5** because of a loop between **R2** and **R6**.

To isolate a failed connection in your network, follow these steps:

Identify the Symptoms

Problem

The symptoms of a problem in your network are usually quite obvious, such as the failure to reach a remote host.

Solution

To identify the symptoms of a problem on your network, start at one end of your network and follow the routes to the other end, entering all or one of the following Junos OS command-line interfaces (CLI) operational mode commands:

user@host> **ping (*ip-address* | *host-name*)**
user@host> **show route (*ip-address* | *host-name*)**
user@host> **traceroute (*ip-address* | *host-name*)**

Sample Output

```
user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2db   0 0000  01  01 a8c6 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2de   0 0000  01  01 a8c3 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2e2   0 0000  01  01 a8bf 10.1.26.2  10.0.0.5

^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[IS-IS/165] 00:02:39, metric 10
                     > to 10.1.26.1  via so-0/0/2.0

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.649 ms  0.521 ms  0.490 ms
 2  10.1.26.2 (10.1.26.2)  0.521 ms  0.537 ms  0.507 ms
 3  10.1.26.1 (10.1.26.1)  0.523 ms  0.536 ms  0.514 ms
 4  10.1.26.2 (10.1.26.2)  0.528 ms  0.551 ms  0.523 ms
 5  10.1.26.1 (10.1.26.1)  0.531 ms  0.550 ms  0.524 ms
```

Meaning

The sample output shows an unsuccessful **ping** command in which the packets are being rejected because the time to live is exceeded. The output for the **show route** command shows the interface (**10.1.26.1**) that you can examine further for possible problems. The **traceroute** command shows the loop between **10.1.26.1** (**R2**) and **10.1.26.2** (**R6**), as indicated by the continuous repetition of the two interface addresses.

Isolate the Causes

Problem

A particular symptom can be the result of one or more causes. Narrow down the focus of your search to find each individual cause of the unwanted behavior.

Solution

To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

user@host> **show < configuration | bgp | interfaces | isis | ospf | route >**

Your particular problem may require the use of more than just the commands listed above. See the appropriate command reference for a more exhaustive list of commonly used operational mode commands.

Sample Output

```
user@R6> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.56.2/30
                                   iso
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.26.2/30
                                   iso
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.2/30
                                   iso
[...Output truncated...]
```

The following sample output is from R2:

```
user@R2>  show route 10.0.0.5

inet.0: 22 destinations, 25 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[Static/5] 00:16:21
                     > to 10.1.26.2 via so-0/0/2.0
                     [BGP/170] 3d 20:23:35, MED 5, localpref 100
                       AS path: 65001 I
                     > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows that all interfaces on **R6** are up. The output from **R2** shows that a static route [**Static/5**] configured on **R2** points to **R6** (**10.1.26.2**) and is the preferred route to **R5** because of its low preference value. However, the route is looping from **R2** to **R6**, as indicated by the missing reference to **R5** (**10.1.15.2**).

Take Appropriate Action

Problem

The appropriate action depends on the type of problem you have isolated. In this example, a static route configured on **R2** is deleted from the [**routing-options**] hierarchy level. Other appropriate actions might include the following:

Solution

- Check the local router's configuration and edit it if appropriate.
- Troubleshoot the intermediate router.
- Check the remote host configuration and edit it if appropriate.
- Troubleshoot routing protocols.
- Identify additional possible causes.

To resolve the problem in this example, enter the following Junos OS CLI commands:

[edit]
user@R2# **delete routing-options static route *destination-prefix***
user@R2# **commit and-quit**
user@R2# **show route *destination-prefix***

Sample Output

```
[edit]
user@R2# delete routing-options static route 10.0.0.5/32

[edit]
user@R2# commit and-quit
commit complete
Exiting configuration mode

user@R2> show route 10.0.0.5

inet.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[BGP/170] 3d 20:26:17, MED 5, localpref 100
                        AS path: 65001 I
                      > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows the static route deleted from the [**routing-options**] hierarchy and the new configuration committed. The output for the **show route** command now shows the BGP route as the preferred route, as indicated by the asterisk (*).

Evaluate the Solution

Problem

If the problem is solved, you are finished. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In relation to the network in "Isolate a Broken Network Connection" on page 373, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you have reason to believe that the problem is related to a known issue, such as a recent change in configuration.

Solution

To evaluate the solution, enter the following Junos OS CLI commands:

user@host> **show route (*ip-address* |*host-name*)**
user@host> **ping (*ip-address* | *host-name*)**
user@host> **traceroute (*ip-address* | *host-name*)**

Sample Output

```
user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[BGP/170]  00:01:35, MED 5, localpref 100, from 10.0.0.2
                      AS path: 65001 I
                    > to 10.1.26.1 via so-0/0/2.0

user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=253 time=0.866 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=253 time=0.837 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=253 time=0.796 ms
^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.796/0.833/0.866/0.029 ms

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.629 ms  0.538 ms  0.497 ms
 2  10.1.12.1 (10.1.12.1)  0.534 ms  0.538 ms  0.510 ms
 3  10.0.0.5 (10.0.0.5)  0.776 ms  0.705 ms  0.672 ms
```

Meaning The sample output shows that there is now a connection between **R6** and **R5.** The **show route** command shows that the BGP route to **R5** is preferred, as indicated by the asterisk (*****). The **ping** command is successful and the **traceroute** command shows that the path from **R6** to **R5** is through **R2** (**10.1.26.1**), and then through **R1** (**10.1.12.1**).

CHAPTER 11

Routing Protocol Process Memory FAQs

Routing Protocol Process Memory FAQs Overview

Junos OS is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device's specialized environment. For example, some executables have been deleted, while other utilities were de-emphasized. Additionally, certain software processes were added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for operating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, while the communication among all the processes is still controlled by the kernel. This separation provides isolation between the processes, and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. The kernel also generates specialized processes as needed for additional functionality, such as SNMP, the Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software, which controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements routing policy, which allows you to control the routing information that is transferred between the routing protocols and the routing table. Using routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

Related Documentation

- Routing Protocol Process Memory FAQs on page 380

Routing Protocol Process Memory FAQs

The following sections present the most frequently asked questions and answers related to the routing protocol process memory utilization, operation, interpretation of related command outputs, and troubleshooting the software process.

Frequently Asked Questions: Routing Protocol Process Memory

This section presents frequently asked questions and answers related to the memory usage of the routing protocol process.

Why does the routing protocol process use excessive memory?

The routing protocol process uses hundreds of megabytes of RAM in the Routing Engine to store information needed for the operation of routing and related protocols, such as BGP, OSPF, IS-IS, RSVP, LDP and MPLS. Such huge consumption of memory is common for the process, as the information it stores includes routes, next hops, interfaces, routing policies, labels, and label-switched paths (LSPs). Because access to the RAM memory is much faster than access to the hard disk, most of the routing protocol process information is stored in the RAM memory instead of using the hard disk space. This ensures that the performance of the routing protocol process is maximized.

How can I check the amount of memory the routing protocol process is using?

You can check routing protocol process memory usage by entering the **show system processes** and the **show task memory** Junos OS command-line interface (CLI) operational mode commands.

The **show system processes** command displays information about software processes that are running on the device and that have controlling terminals. The **show task memory** command displays memory utilization for routing protocol tasks on the Routing Engine.

You can check the routing protocol process memory usage by using the **show system processes** command with the **extensive** option. The **show task memory** command displays a report generated by the routing protocol process on its own memory usage. However, this report does not display all the memory used by the process. The value reported by the routing protocol process does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the process's internal memory manager. Further, the Resident Set Size value includes shared library pages used by the routing protocol process.

For more information about checking the routing protocol process memory usage.

For more information, see the **show system processes** command and the **show task memory** command.

I just deleted a large number of routes from the routing protocol process. Why is it still using so much memory?

The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size. The **RES** value includes shared library pages used by the process. Any amount of memory freed by the process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has freed a large amount of memory.

Frequently Asked Questions: Interpreting Routing Protocol Process-Related Command Outputs

This section presents frequently asked questions and answers about the routing protocol process-related Junos OS command-line interface (CLI) command outputs that are used to display the memory usage of the routing protocol process.

How do I interpret memory numbers displayed in the show system processes extensive command output?

The **show system processes extensive** command displays exhaustive system process information about software processes that are running on the device and have controlling terminals. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

To check overall CPU and memory usage, enter the **show system processes extensive** command. Refer to Table 17 on page 382 for information about the **show system processes extensive** commands output fields.

```
user@host> show system processes extensive
last pid:   544;  load averages:  0.00,  0.00,  0.00    18:30:33
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 184K Cache, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  544 root     30   0   604K   768K RUN      0:00  0.00%  0.00% top
    3 root     28   0     0K    12K psleep   0:00  0.00%  0.00% vmdaemon
    4 root     28   0     0K    12K update   0:03  0.00%  0.00% update
  528 aviva    18   0   660K   948K pause    0:00  0.00%  0.00% tcsh
  204 root     18   0   300K   544K pause    0:00  0.00%  0.00% csh
  131 root     18   0   332K   532K pause    0:00  0.00%  0.00% cron
  186 root     18   0   196K    68K pause    0:00  0.00%  0.00% watchdog
   27 root     10   0   512M 16288K mfsidl   0:00  0.00%  0.00% mount_mfs
    1 root     10   0   620K   344K wait     0:00  0.00%  0.00% init
  304 root      3   0   884K   900K ttyin    0:00  0.00%  0.00% bash
  200 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  203 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  202 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  201 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  194 root      2   0  2248K  1640K select   0:11  0.00%  0.00% rpd
  205 root      2   0   964K   800K select   0:12  0.00%  0.00% tnp.chassisd
```



```
 189 root       2 -12   352K   740K select   0:03  0.00%  0.00% xntpd
 114 root       2   0   296K   612K select   0:00  0.00%  0.00% amd
 188 root       2   0   780K   600K select   0:00  0.00%  0.00% dcd
 527 root       2   0   176K   580K select   0:00  0.00%  0.00% rlogind
 195 root       2   0   212K   552K select   0:00  0.00%  0.00% inetd
 187 root       2   0   192K   532K select   0:00  0.00%  0.00% tnetd
  83 root       2   0   188K   520K select   0:00  0.00%  0.00% syslogd
 538 root       2   0  1324K   516K select   0:00  0.00%  0.00% mgd
  99 daemon     2   0   176K   492K select   0:00  0.00%  0.00% portmap
 163 root       2   0   572K   420K select   0:00  0.00%  0.00% nsrexecd
 192 root       2   0   560K   400K select   0:10  0.00%  0.00% snmpd
 191 root       2   0  1284K   376K select   0:00  0.00%  0.00% mgd
 537 aviva      2   0   636K   364K select   0:00  0.00%  0.00% cli
 193 root       2   0   312K   204K select   0:07  0.00%  0.00% mib2d
   5 root       2   0     0K    12K pfesel   0:00  0.00%  0.00% if_pfe
   2 root     -18   0     0K    12K psleep   0:00  0.00%  0.00% pagedaemon
   0 root     -18   0     0K     0K sched    0:00  0.00%  0.00% swapper
```

Table 17 on page 382 describes the output fields that represent the memory values for the **show system processes extensive** command. Output fields are listed in the approximate order in which they appear.

Table 17: show system processes extensive Output Fields

Field Name	Field Description
Mem	Information about physical and virtual memory allocation.
Active	Memory allocated and actively used by the program.
Inact	Memory allocated but not recently used or memory freed by the programs. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.
Wired	Memory that is not eligible to be swapped, usually used for in-kernel memory structures and/or memory physically locked by a process.
Cache	Memory that is not associated with any program and does not need to be swapped before being reused.
Buf	Size of memory buffer used to hold data recently called from the disk.
Free	Memory that is not associated with any programs. Memory freed by a process can become **Inactive**, **Cache**, or **Free**, depending on the method used by the process to free the memory.
Swap	Information about swap memory. • Total—Total memory available to be swapped to disk. • Used—Memory swapped to disk. • Free—Memory available for further swap.

The rest of the command output displays information about the memory usage of each process. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the program in physical memory, which is also known as RSS or Resident Set Size. For more information, see the **show system processes** command.

What is the difference between Active and Inact memory that is displayed by the show system processes extensive command?

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between **Active** and **Inact** memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.

The pageout process first scans the **Inact** list, and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the **Inact** list to the **Active** list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the **Inact** list, pages from the **Active** list get freed up.

Because the pageout process runs only when the system is under memory pressure, the pages on the **Inact** list remain untouched – even if they have not been accessed recently – when the amount of **Free** memory is adequate.

How do I interpret memory numbers displayed in the show task memory command output?

The **show task memory** command provides a comprehensive picture of the memory utilization for routing protocol tasks on the Routing Engine. The routing protocol process is the main task that uses Routing Engine memory.

To check routing process memory usage, enter the **show task memory** command. Refer to Table 18 on page 383 for information about the **show task memory** command output fields.

```
user@host> show task memory
Memory                 Size (kB)  %Available  When
  Currently In Use:        29417          3%  now
  Maximum Ever Used:       33882          4%  00/02/11 22:07:03
  Available:              756281        100%  now
```

Table 18 on page 383 describes the output fields for the **show task memory** command. Output fields are listed in the approximate order in which they appear.

Table 18: show task memory Output Fields

Field Name	Field Description
Memory Currently In Use	Memory currently in use. Dynamically allocated memory plus the **DATA** segment memory in kilobytes.
Memory Maximum Ever Used	Maximum memory ever used.
Memory Available	Memory currently available.

The **show task memory** command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the **TEXT** and

STACK segments, or the memory used by the routing protocol process's internal memory manager.

Why is the Currently In Use value less than the RES value?

The **show task memory** command displays a **Currently In Use** value measured in kilobytes. This value represents the memory currently in use. It is the dynamically allocated memory plus the **DATA** segment memory. The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size.

The **Currently In Use** value does not account for all of the memory that the routing protocol process uses. This value does not include the memory used for the **TEXT** and the **STACK** segments, and a small percentage of memory used by the routing protocol process's internal memory manager. Further, the **RES** value includes shared library pages used by the routing protocol process.

Any amount of memory freed by the routing protocol process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the **Currently In Use** value and the **RES** value.

Frequently Asked Questions: Routing Protocol Process Memory Swapping

This section presents frequently asked questions and answers related to the memory swapping of the routing protocol process from the Routing Engine memory to the hard disk memory.

How do I monitor swap activity?

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. You can monitor the swap activity by viewing the syslog message reported by the kernel during periods of high pageout activity.

The syslog message appears as follows:

```
Mar  3 20:08:02 olympic /kernel: High pageout rate!! 277 pages/sec.
```

You can use the **vmstat -s** command to print the statistics for the swapout activity. The displayed statistics appear as follows:

```
0 swap pager pageouts
0 swap pager pages paged out
```

The **swap pager pageouts** is the number of pageout operations to the swap device, and the **swap pager pages paged out** is the number of pages paged out to the swap device.

Why does the system start swapping when I try to dump core using the request system core-dumps command?

The **request system core-dumps** command displays a list of system core files created when the device has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, path, and filename. You can use the **core-filename** option and the **core-file-info**, **brief**, and **detail** options to display more information about the specified core-dump files.

You can use the **request system core-dumps** command to perform a non-fatal core-dump without aborting the routing protocol process. To do this, the routing protocol process is forked, generating a second copy, and then aborted. This process can double the memory consumed by the two copies of the routing protocol processes, pushing the system into swap.

Why does the show system processes extensive command show that memory is swapped to disk although there is plenty of free memory?

Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the **show system processes extensive** command shows that memory is swapped to disk even though there is plenty of free memory, and such a situation is not unusual.

Frequently Asked Questions: Troubleshooting the Routing Protocol Process

This section presents frequently asked questions and answers related to a shortage of memory and memory leakage by the routing protocol process.

What does the RPD_OS_MEMHIGH message mean?

The **RPD_OS_MEMHIGH** message is written into the system message file if the routing protocol process is running out of memory. This message alerts you that the routing protocol process is using the indicated amount and percentage of Routing Engine memory, which is considered excessive. This message is generated either because the routing protocol process is leaking memory or the use of system resources is excessive, perhaps because routing filters are misconfigured or the configured network topology is very complex.

When the memory utilization for the routing protocol process is using all available Routing Engine DRAM memory (Routing Engines with maximum 2 GB DRAM) or reaches the limit of 2 GB of memory (Routing Engines with 4 GB DRAM), a message of the following form is written every minute in the syslog message file:

```
RPD_OS_MEMHIGH: Using 188830 KB of memory, 100 percent of available
```

This message includes the amount, in kilobytes and/or the percentage, of the available memory in use.

This message should not appear under normal conditions, as any further memory allocations usually require a portion of existing memory to be written to swap. As a recommended solution, increase the amount of RAM in the Routing Engine. For more information, go to http://kb.juniper.net/InfoCenter/index?page=content&id=KB14186 .

What can I do when there is a memory shortage even after a swap?

It is not recommended for the system to operate in this state, notwithstanding the existence of swap. The protocols that run in the routing protocol process usually have a real-time requirement that cannot reliably withstand the latency of being swapped to hard disk. If the memory shortage has not resulted from a memory leak, then either a reduction in the memory usage or an upgrade to a higher memory-capacity Routing Engine is required.

How do I determine whether there is a memory leak in the routing protocol process?

Memory leaks are typically the result of a seemingly unbounded growth in the memory usage of a process as reported by the **show system processes extensive** command.

There are two classes of memory leaks that the routing protocol process can experience.

- The first class occurs when the allocated memory that is no longer in use is not freed. This class of leak can usually be fixed by taking several samples of the **show task memory detail** command over a period of time and comparing the deltas.
- The second class occurs when there is a late access to freed memory. If the access is not outside the mapped address space, the kernel backfills the accessed page with real memory. This backfill is done without the knowledge of the routing protocol process's internal memory allocator, which makes this class of leak much more difficult to resolve. If a memory leak of this class is suspected, writing the state of the system to a disk file (creating a core file) is suggested.

A large discrepancy between the **RES** value and the **Currently In Use** value might indicate a memory leak. However, large discrepancies can also occur for legitimate reasons. For example, the memory used for the **TEXT** and **STACK** segments or the memory used by the routing protocol process's internal memory manager might not be displayed. Further, the **RES** value includes shared library pages used by the process.

What is the task_timer?

The source of a routing protocol process memory leak can usually be identified by dumping the timers for each task. You can use the **show task *task-name*** command to display routing protocol tasks on the Routing Engine. Tasks can be baseline tasks performed regardless of the device's configuration, and other tasks that depend on the device configuration.

For more information, see the **show task** command.

Related Documentation

- Routing Protocol Process Memory FAQs Overview on page 379

PART 5

Index

Index

Symbols

A

B

C

I

K

L

M

N

S

T

