Engineering Simplicity

# Junos® OS

# Protocol-Independent Routing Properties User Guide

Published
2020-03-26

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



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

# About the Documentation

IN THIS SECTION



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.

Use this guide to configure, monitor, and troubleshoot protocol-independent routing properties on your Juniper Network devices.

Junos OS Routing Protocols Library for Routing Devices

## 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 https://www.juniper.net/documentation/.

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 https://www.juniper.net/books.

# 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 CLI Explorer.

## 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. |
|  | Tip | Indicates helpful information. |
|  | Best practice | Alerts you to a recommended use or implementation. |

Table 2 on page xv 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:<br><br>user@host> **configure** |
| `Fixed-width text like this` | Represents output that appears on the terminal screen. | `user@host>` **show chassis alarms**<br><br>`No alarms currently active` |
| *Italic text like this* | • Introduces or emphasizes important new terms.<br>• Identifies guide names.<br>• Identifies RFC and Internet draft titles. | • A policy *term* is a named structure that defines match conditions and actions.<br>• *Junos OS CLI User Guide*<br>• RFC 1997, *BGP Communities Attribute* |

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

| Convention | Description | Examples |
|---|---|---|
| *Italic text like this* | Represents variables (options for which you substitute a value) in commands or configuration statements. | Configure the machine's domain name:<br>[edit]<br>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.<br>• The console port is labeled **CONSOLE**. |
| < > (angle brackets) | Encloses 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**<br>**(*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) | Encloses a variable for which you can substitute one or more values. | **community name members [ *community-ids* ]** |
| Indention and braces ( { } ) | Identifies a level in the configuration hierarchy. | [edit]<br>routing-options {<br>static {<br>route default {<br>nexthop *address*;<br>retain;<br>}<br>}<br>} |
| ; (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**.<br>• 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 so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

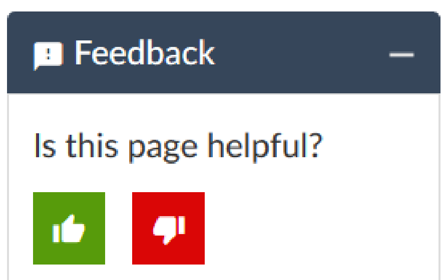


  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software 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 Juniper Care or Partner Support Services 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 https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://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: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

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

### Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- 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 https://support.juniper.net/support/requesting-support/.

# 1

PART

# Overview

CHAPTER 1

# Introduction to Protocol-Independent Routing Properties

IN THIS CHAPTER



## 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

## Routing Table Features in Junos OS

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, 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 3.

**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 |

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

| Description | Static | Aggregate | Generated |
|---|---|---|---|
| Preference values | X | X | X |
| Additional preference values | X | X | X |
| Independent preference (**qualified-next-hop** statement) | X | – | – |
| 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 | – | – |

2

PART

# Configuring Protocol-Independent Routing Properties

CHAPTER 2

# Configuring Routing Tables and Static Routes

IN THIS CHAPTER



## 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.

RELATED DOCUMENTATION



# Example: Configuring a Basic Set of Static Routes for Connecting to Stub Networks

IN THIS SECTION



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 8 shows the sample network.

**Figure 1: Customer Routes Connected to a Service Provider**

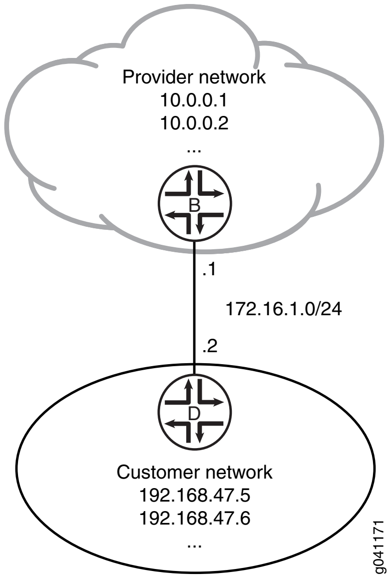


## Configuration

### 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 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
```

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

**IN THIS SECTION**



Confirm that the configuration is working properly.

***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
```

RELATED DOCUMENTATION

Understanding Basic Static Routing | 6
Verifying the Static Route Configuration | 57
Example: Configuring IPv6 Static Routes | 13

## Example: Configuring IPv6 Static Routes

IN THIS SECTION

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 14 shows the sample network.

**Figure 2: Customer Routes Connected to a Service Provider**

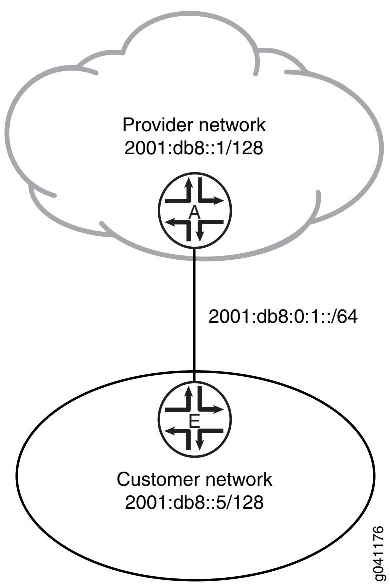


### Configuration

#### 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]
   user@A# set fe-1/2/0 unit 1 description to-E
   user@A# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1:2a0:a502:0:1da/64
   user@A# set lo0 unit 1 family inet6 address 2001:db8::1/128 primary
   user@A# set lo0 unit 1 family inet6 address 2001:db8::2/128
   user@A# set 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]
user@A# 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]
user@E# set interfaces fe-1/2/0 unit 25 description to-A
user@E# set interfaces fe-1/2/0 unit 25 family inet6 address 2001:db8:0:1:2a0:a502:0:19da/64
user@E# 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]
user@E# set 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

**IN THIS SECTION**



Confirm that the configuration is working properly.

***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
```

RELATED DOCUMENTATION


Example: Configuring a Basic Set of Static Routes for Connecting to Stub Networks | 7

Example: Configuring Static Routes Between Logical Systems Within the Same Router | 20


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

**IN THIS SECTION**



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

You must connect the logical systems by using logical tunnel (**lt**) interfaces. See *Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches*.

### Overview

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**

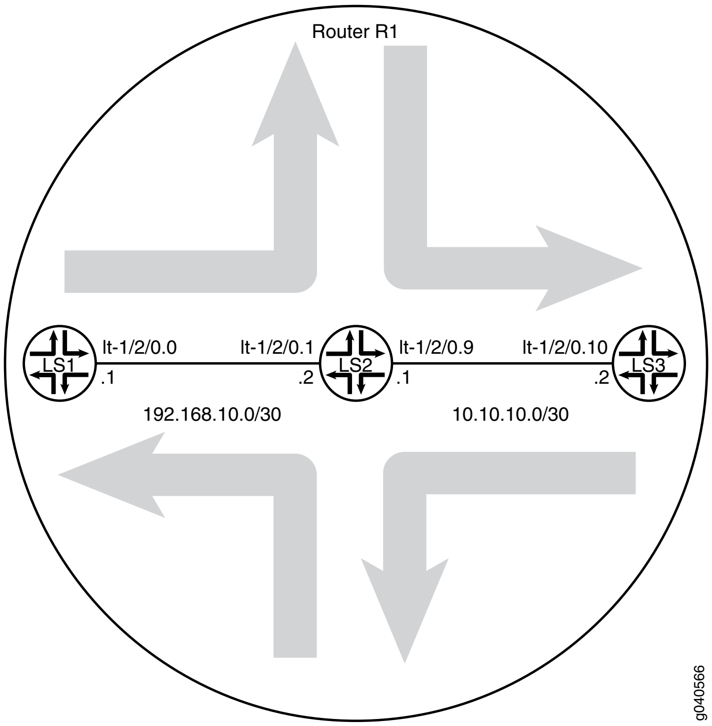


## Configuration

### 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

IN THIS SECTION



Confirm that the configuration is working properly.

***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

*Example: Creating an Interface on a Logical System*

*Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches*

## 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, the Junos OS chooses in a random fashion one of the next-hop addresses to install into the forwarding table.

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 6;
    }
    preference 5;
}
```

In this example, the qualified next hop 10.10.10.7 is assigned the preference 6, and the next-hop 10.10.10.10 is assigned the preference 5.

**NOTE:** The **preference** and **metric** options in the [**edit route *route* qualified-next-hop**] hierarchy 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 only, similar to how the route preference overrides the default preference and metric (for that specific route).

**NOTE:** Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

## RELATED DOCUMENTATION

Example: Configuring Static Route Preferences and Qualified Next Hops to Control Static Route Selection | 29

# Example: Configuring Static Route Preferences and Qualified Next Hops to Control Static Route Selection

IN THIS SECTION



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 30.

**Figure 4: Controlling Static Route Selection**

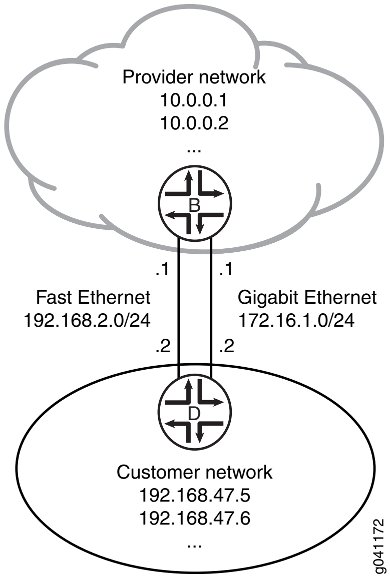


## Configuration

### 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 in Provider Network

```
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 in Customer Network

```
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

**IN THIS SECTION**



Confirm that the configuration is working properly.

### *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

Understanding Static Route Preferences and Qualified Next Hops | 27
Verifying the Static Route Configuration | 57

## Understanding Static Route Control in Routing and Forwarding Tables

IN THIS SECTION

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.

RELATED DOCUMENTATION

Example: Preventing a Static Route from Being Readvertised | 38

# Example: Preventing a Static Route from Being Readvertised

IN THIS SECTION



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 38 shows the sample network.

**Figure 5: Customer Routes Connected to a Service Provider**

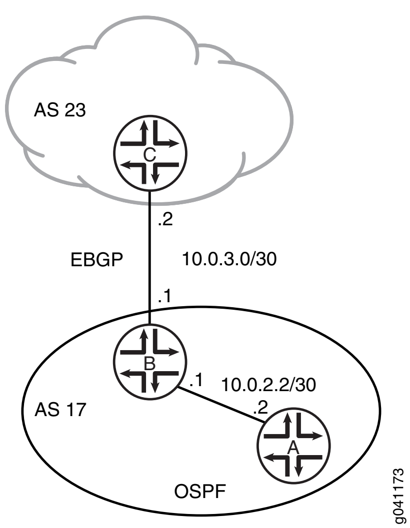

## Configuration

### 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

Understanding Static Route Control in Routing and Forwarding Tables | 36
Verifying the Static Route Configuration | 57

## 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.
- **inet6.1**—For IPv6 multicast forwarding cache. This table stores the IPv6 (S,G) group entries that are dynamically created as a result of join state information.
- ***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 51.

  > 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.
- **l2circuit.0**—For l2circuit routes learned from LDP. Routes in this table are used to send or receive l2circuit signaling messages.
- **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.

## RELATED DOCUMENTATION

Example: Creating Routing Tables | 48

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table | 51

# Example: Creating Routing Tables

IN THIS SECTION



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

**IN THIS SECTION**



Confirm that the configuration is working properly.

***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.

RELATED DOCUMENTATION

Understanding Junos OS Routing Tables | 46

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

IN THIS SECTION



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. 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

### IN THIS SECTION



### 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

## Verifying the Static Route Configuration

### Purpose

Verify that the static routes are in the routing table and that those routes are active.

### Action

From the CLI, enter the **show route terse** command.

## Sample Output

user@host> **show route terse**

```
inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
A Destination          P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.47.5/32          S   5                        Reject
* 172.16.0.0/12        S   5                       >192.168.71.254
* 192.168.0.0/18       S   5                       >192.168.71.254
* 192.168.40.0/22      S   5                       >192.168.71.254
```

```
* 192.168.64.0/18    S   5                      >192.168.71.254
* 192.168.64.0/21    D   0                      >fxp0.0
* 192.168.71.246/32  L   0                       Local
* 192.168.220.4/30   D   0                      >ge-0/0/1.0
* 192.168.220.5/32   L   0                       Local
* 192.168.220.8/30   D   0                      >ge-0/0/2.0
* 192.168.220.9/32   L   0                       Local
* 192.168.220.12/30  D   0                      >ge-0/0/3.0
* 192.168.220.13/32  L   0                       Local
* 192.168.220.17/32  L   0                       Reject
* 192.168.220.21/32  L   0                       Reject
* 192.168.220.24/30  D   0                      >at-1/0/0.0
* 192.168.220.25/32  L   0                       Local
* 192.168.220.28/30  D   0                      >at-1/0/1.0
* 192.168.220.29/32  L   0                       Local
* 224.0.0.9/32       R 100          1            MultiRecv
```

### Meaning

The output shows a list of the routes that are currently in the **inet.0** routing table. Verify the following information:

- Each configured static route is present. Routes are listed in ascending order by IP address. Static routes are identified with an **S** in the protocol (**P**) column of the output.
- Each static route is active. Routes that are active show the next-hop IP address in the **Next hop** column. If a route's next-hop address is unreachable, the next-hop address is identified as **Reject**. These routes are not active routes, but they appear in the routing table because the **passive** attribute is set.
- The preference for each static route is correct. The preference for a particular route is listed in the **Prf** column of the output.

## RELATED DOCUMENTATION

show route terse | 708 in the CLI Explorer

# Conserving IP Addresses Using Static Routes

IN THIS SECTION



Hosting providers host multiple servers for multiple customers and want to conserve the usage of their IP address space. Traditionally, when a hosting provider client adds new servers, the servers are allocated a small block of IP addresses, such as a /29 block, and the client's servers are all located in that block of IP addresses.

## The Issue, Illustrated

For example, Customer A might need three servers and is assigned the block 10.3.3.0/29 (10.3.3.0 through 10.3.3.7). In this scenario, several IP addresses are consumed. These include the network and broadcast IP addresses (10.3.3.0 and 10.3.3.7), the addresses for the router gateway that the servers are connected to, and the addresses of the individual servers. To allocate three servers, eight IP addresses have to be allocated. Breaking up a single /24 network into 32 /29 networks results in 96 IP addresses out of the 256, in that /24 is being consumed by the network, broadcast, and gateway addresses. When this effect is multiplied across thousands of hosting providers, IP address space is far from being used efficiently. Figure 6 on page 60 illustrates the issue.

**Figure 6: Inefficient Use of IP Address Space**

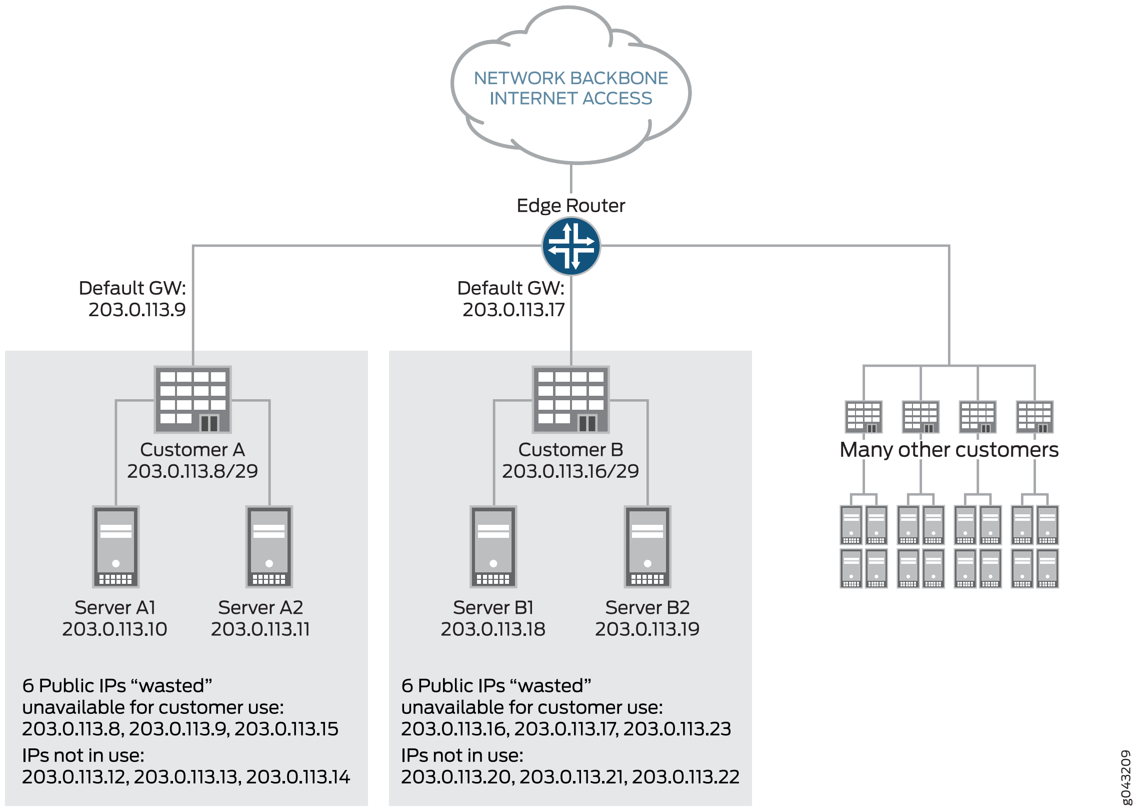


In this configuration, each customer is allocated a /29 block of address space. For each block, the network, broadcast, and gateway addresses are not available for server IP addressing, which results in three IP addresses being used inefficiently. In addition, the blocks consume unused IP addresses for future expansion.

## Solution

This issue can be resolved by configuring the interface on the router with an address from the reserved IPv4 prefix for shared address space (RFC 6598) and by using static routes pointed at interfaces. IANA has recorded the allocation of an IPv4 /10 for use as shared address space. The shared address space address range is 100.64.0.0/10.

The interface in the router gets allocated an IP address from the RFC 6598 space, so it is not consuming publicly routable address space, and connectivity is handled with static routes on an interface. The interface in the server is configured with a publicly routable address, but the router interfaces are not. Network and broadcast addresses are consumed out of the RFC 6598 space rather than the publicly routable address space.

This feature is supported on QFX10000 switches starting with Junos OS 17.1R1.

Figure 7 on page 61 shows the efficient use of IP address space.

**Figure 7: Configuration Using the Shared Address Space**

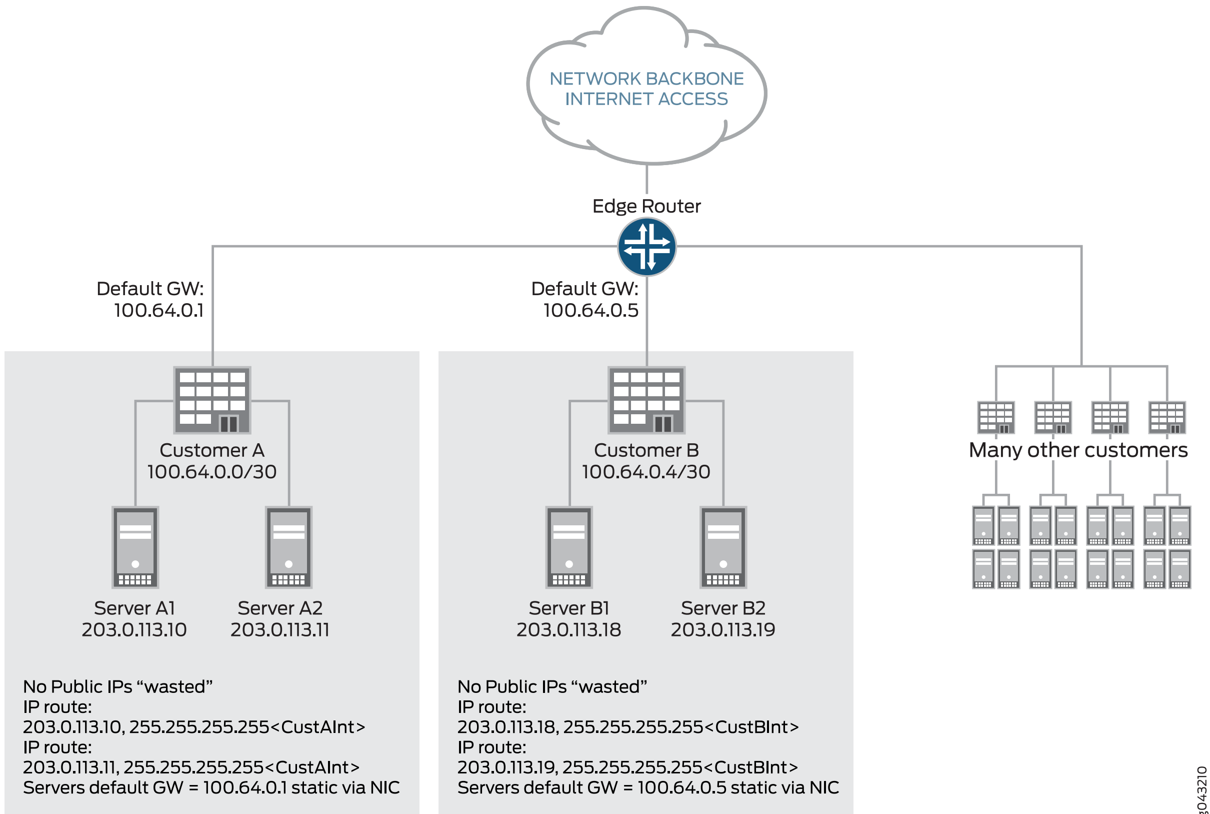


In this configuration, each customer gets allocated individual IP addresses per server. There is a static route that can be configured as a host route. The interface in the router gets allocated an IP address from the RFC 6598 space, so it does not consume publicly routable address space, and connectivity is handled with static routes out to an interface.

## Configuration

The configuration would look like this for Customer A on the gateway router:

```
interfaces {
    ge-1/0/1 {
        unit 0 {
            family inet {
                address 100.64.0.1/30;
            }
        }
    }
}
```

```
routing-options {
  static {
    route 203.0.113.10/32 {
      qualified-next-hop ge-1/0/1.0;
    }
    route 203.0.113.11 {
      qualified-next-hop ge-1/0/1.0;
    }
  }
}
```

With this configuration, no publicly routable IP addresses are wasted. It is worth noting that when a packet is forwarded in this configuration from the router to the server of Customer A's server 203.0.113.10, the route is forwarded out to the interface ge-1/0/1.0 which has an IP address of **100.64.0.1**.

The servers for customer A would be configured as follows:

```
ifconfig eth0 203.0.113.10 netmask 255.255.255.255
route add -host 100.64.0.1/32 dev eth0 route add default gw 100.64.0.1
```

```
ifconfig eth0 203.0.113.11 netmask 255.255.255.255
route add -host 100.64.0.1/32 dev eth0 route add default gw 100.64.0.1
```

This example shows a single host route per server, which is a 1:1 mapping. This could equate to a large number of static host routes, if maintained. For scaling purposes, we need to support nonhost routes in this environment. For example, if there were a Customer C in this configuration that had eight servers, it would be much more efficient to allocate a /29 route on the router that points out the interface on which the eight servers are connected. If Customer C were allocated server IPs from 203.0.114.8 through 203.0.114.15 and these were connected via interface ge-1/0/2.0, this would look like:

```
user@host# set routing-options static route 203.0.114.8/29 qualified-next-hop ge-1/0/2.0
```

**Release History Table**

| Release | Description |
|---|---|
| 17.1R1 | This feature is supported on QFX10000 switches starting with Junos OS 17.1R1. |

CHAPTER 3

# Configuring Static Routes for CLNS

IN THIS CHAPTER



## 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.

RELATED DOCUMENTATION

Example: Configuring Static Routes for CLNS When No IGP is Present | 63

## Example: Configuring Static Routes for CLNS When No IGP is Present

IN THIS SECTION

This example shows how to configure static routes for CLNS.

## Requirements

Before you begin, configure the network interfaces. See *Interfaces User 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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
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

### *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.

## RELATED DOCUMENTATION

*CLNS Configuration Overview*

Understanding Static Routes for CLNS | 63

CHAPTER 4

# Configuring Route Aggregation

**IN THIS CHAPTER**



## Understanding Route Aggregation

The route aggregation methodology helps minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route. The aggregation methodology does not help reduce the size of the routing-table on the router that does the aggregation. When you configure an export policy that only advertises the aggregate but not the contributing routes anymore, you then have the aggregation effect on the routers that receive updates.

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 **192.168.0.0/16**, routes to **192.168.192.0/19** and **192.168.67.0/24** are contributing routes, but routes to **192.168.0.0./8** and **192.168.0.0/16** are not.

A route can only contribute 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 **192.168.0.0/16** can contribute to an aggregate route to **192.168.0.0/13**.

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 metric type number;
        }
    }
}
```

### 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 Policies, Firewall Filters, and Traffic Policers User Guide*. 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.

> **NOTE:** Starting with Junos OS 13.2R1, a BGP route is hidden when the AS path of an aggregate route—built from contributing routes— is more than half of the maximum BGP packet size (4096 bytes). Such AS paths have the OverflowASPathSize flag set for them. If you would like to leak such a BGP route, whose AS path length can overflow, we recommend to add the AS path statically in the default route configuration. For example:
>
> ```
> [edit routing-instances instance-name routing options]
> user@host# set aggregate route 0.0.0.0/0 as-path path 1267
> ```

### 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 a Tag Value for Aggregate Routes

By default, no tag values are associated with aggregate routes. You can specify a tag value by including the **tag** option:

```
aggregate (defaults | route) {
  tag metric type number;
}
```

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;
}
```

**Release History Table**

| Release | Description |
|---|---|
| 13.2R1 | Starting with Junos OS 13.2R1, a BGP route is hidden when the AS path of an aggregate route—built from contributing routes— is more than half of the maximum BGP packet size (4096 bytes). |

RELATED DOCUMENTATION



## Example: Summarizing Static Routes Through Route Aggregation

IN THIS SECTION



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 77 shows the sample network.

**Figure 8: Aggregate Route Advertised to an ISP**

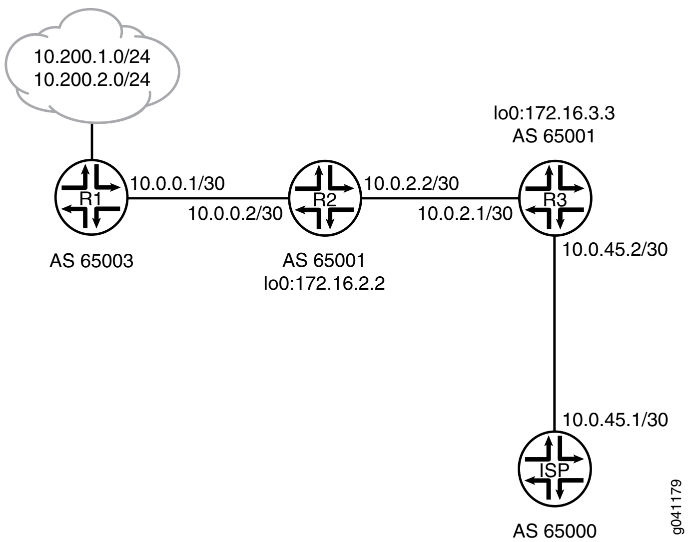


## Configuration

### 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

**IN THIS SECTION**



Confirm that the configuration is working properly.

***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

Understanding Route Aggregation | 67

CHAPTER 5

# Configuring Bidirectional Forwarding Detection for Static Routes

**IN THIS CHAPTER**



## Understanding BFD for Static Routes for Faster Network Failure Detection

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 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.

**NOTE:** On MX Series devices, multihop BFD is not supported on a static route if the static route is configured with more than one next hop. It is recommended that you avoid using multiple next hops when a multihop BFD is required for a static route.

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

**NOTE:** Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, the **bfd-liveness-detection** command includes the description field. The description is an attribute under the **bfd-liveness-detection** object and it is supported only on SRX Series devices. This field is applicable only for the static routes.

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 also supported for the eBGP protocol.

**NOTE:** Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2. Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.

There are three types of BFD sessions based on the source from which BFD packets are sent to the neighbors. Different types of BFD sessions and their descriptions are:

| Type of BFD session | Description |
|---|---|
| Non-distributed BFD | BFD sessions running completely on the Routing Engine. |
| Distributed BFD | BFD sessions running completely on the Packet Forwarding Engine. |
| Inline BFD<br><br>**NOTE:** Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured **enhanced-ip**.<br><br>**NOTE:** Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces. | BFD sessions running on the FPC hardware. |

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:** QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.

**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.

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.

**Release History Table**

| Release | Description |
|---|---|
| 16.1R1 | Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces. |
| 15.1X49-D70 | Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, the **bfd-liveness-detection** command includes the description field. The description is an attribute under the **bfd-liveness-detection** object and it is supported only on SRX Series devices. This field is applicable only for the static routes. |
| 15.1F6 | Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2. |
| 15.1F3 | Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2. |
| 13.3 | Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured **enhanced-ip**. |

RELATED DOCUMENTATION

*Enabling Dedicated and Real-Time BFD*

Example: Configuring BFD for Static Routes for Faster Network Failure Detection | 90
Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection | 110

# Example: Configuring BFD for Static Routes for Faster Network Failure Detection

IN THIS SECTION

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 9 on page 91 shows the sample network.

**Figure 9: Customer Routes Connected to a Service Provider**

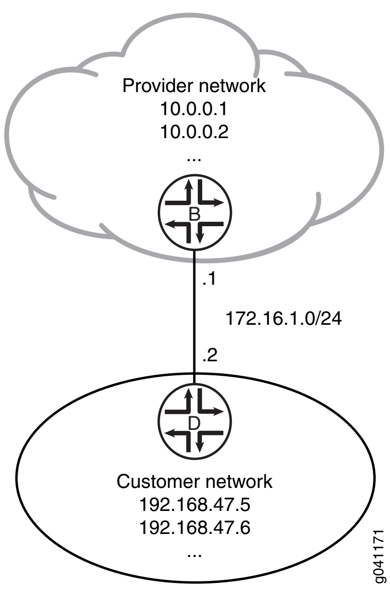


### Configuration

#### 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 description Site-xxx
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
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description Site-xxx
```

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 {
            description Site- xxx;
            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 {
            description Site - xxx;
            minimum-interval 1000;
        }
    }
}
```

## Verification

IN THIS SECTION



Confirm that the configuration is working properly.

***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, description Site-xxx, 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
```

**NOTE:** The **description Site- *<xxx>*** is supported only on the SRX Series devices.

If each client has more than one description field, then it displays "and more" along with the first description field.

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.

## RELATED DOCUMENTATION

Understanding BFD for Static Routes for Faster Network Failure Detection | 85

# Understanding BFD Authentication for Static Route Security

IN THIS SECTION



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.

> **NOTE:** EX3300 supports BFD over static routes only.

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.

**NOTE:** QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.

### Security Authentication Keychains

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

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

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

### Strict Versus Loose Authentication

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

RELATED DOCUMENTATION

Example: Configuring BFD Authentication for Securing Static Routes | 101

## Example: Configuring BFD Authentication for Securing Static Routes

IN THIS SECTION

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 10 on page 103 shows the sample network.

Figure 10: Customer Routes Connected to a Service Provider

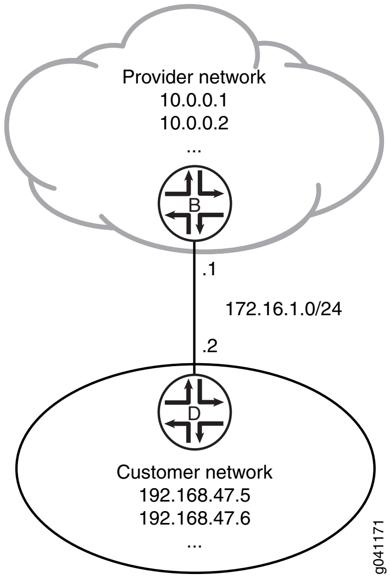


## Configuration

**IN THIS SECTION**



### 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 description Site-xxx
```

```
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 "$ABC123$ABC123$ABC123"
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 "$ABC123$ABC123$ABC123"
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
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description Site-xxx
```

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 "$ABC123$ABC123$ABC123"
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 {
            description Site- xxx;
```

```
            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 "$ABC123$ABC123$ABC123"; ## SECRET-DATA
            start-time "2011-1-1.12:00:00 -0800";
        }
    }
}
```

## Verification

IN THIS SECTION



Confirm that the configuration is working properly.

***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, description Site-xxx, 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.

> **NOTE:** The **description Site- <*xxx*>** is supported only on the SRX Series devices.
>
> If each client has more than one description field, then it displays "and more" along with the first description field.

RELATED DOCUMENTATION

Understanding BFD Authentication for Static Route Security | 99

# Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection

**IN THIS SECTION**



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 111.

**Figure 11: BFD Enabled on Qualified Next Hops**

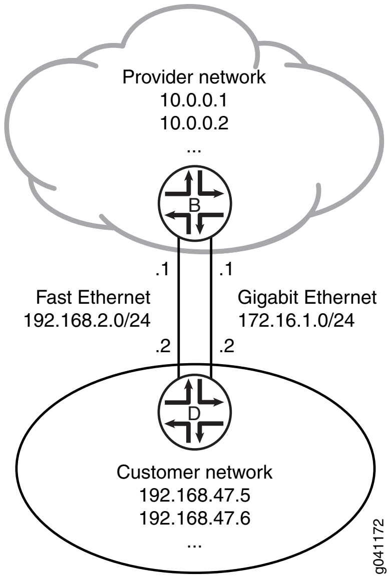


## Configuration

### 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

IN THIS SECTION



Confirm that the configuration is working properly.

### *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 Static Route Preferences and Qualified Next Hops to Control Static Route Selection | 29

Understanding Static Route Preferences and Qualified Next Hops | 27

Understanding BFD for Static Routes for Faster Network Failure Detection | 85

Verifying the Static Route Configuration | 57

CHAPTER 6

# Configuring Packet Forwarding Behavior for Protocol-Independent Routing

IN THIS CHAPTER



## Understanding Default Routing Table Groups for Interface Routes on PTX 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 for VPNs*

*Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Routers*

## 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 12 on page 121 illustrates the route to forwarding next-hop bindings with indirect next hop disabled.

**Figure 12: Route to Forwarding Next-Hop Bindings**

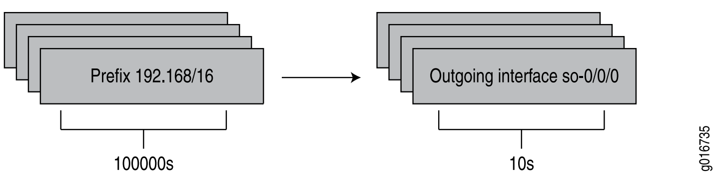


You can enable Junos OS to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. As a result, fewer route to forwarding next-hop bindings need to be updated, which improves the route convergence time. Figure 13 on page 121 illustrates the route to forwarding next-hop bindings with indirect next hop enabled.

**Figure 13: Route to Forwarding Indirect Next-Hop Bindings**

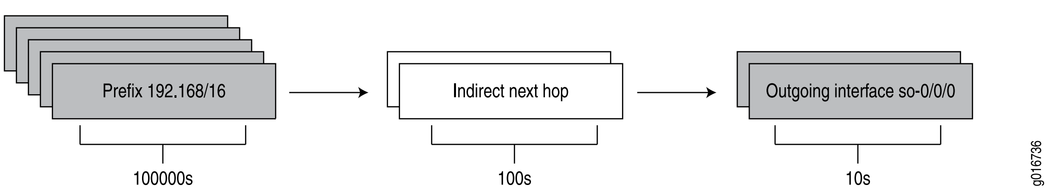


RELATED DOCUMENTATION

# Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

IN THIS SECTION



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 14 on page 123 shows the sample network.

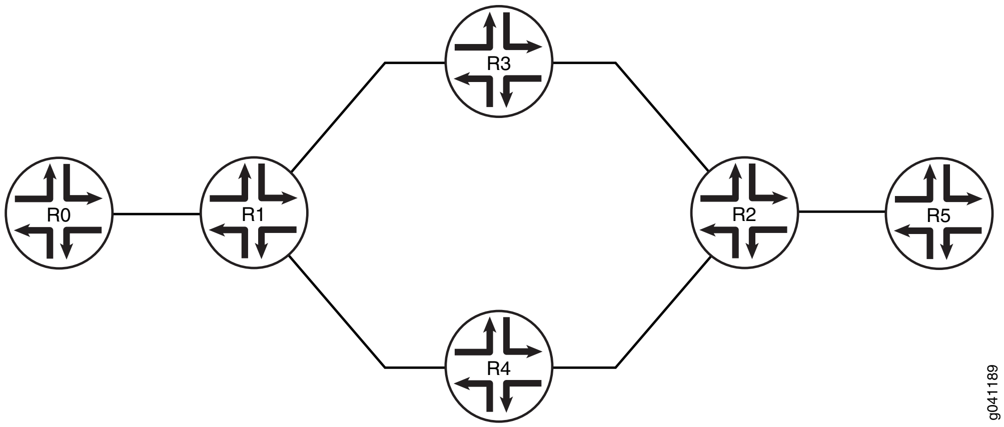


The “CLI Quick Configuration” on page 123 section shows the full configuration on all of the devices in Figure 14 on page 123. Otherwise, the example focuses on Device R0, Device R1, and Device R2.

## 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 fe-1/2/0 unit 1 family inet address 10.0.0.1/30
user@R0# set lo0 unit 1 family inet address 1.1.0.1/32
user@R0# set lo0 unit 1 family inet address 1.1.0.2/32
user@R0# set lo0 unit 1 family inet address 1.1.0.3/32
user@R0# set lo0 unit 1 family inet address 1.1.0.4/32
user@R0# set lo0 unit 1 family inet address 1.1.0.5/32
user@R0# set lo0 unit 1 family inet address 1.1.0.6/32
user@R0# set lo0 unit 1 family inet address 1.1.0.7/32
user@R0# set lo0 unit 1 family inet address 1.1.0.8/32
user@R0# set 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 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 fe-1/2/0 unit 2 family inet address 10.0.0.2/30
user@R1# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30
user@R1# set fe-1/2/2 unit 9 family inet address 10.0.0.9/30
user@R1# set lo0 unit 2 family inet address 1.1.1.1/32
```

2. Configure BGP.

```
[edit protocols]
user@R1# set bgp export send-local
user@R1# set bgp export send-static
user@R1# set bgp group int type internal
user@R1# set bgp group int local-address 1.1.1.1
user@R1# set bgp group int neighbor 2.2.2.2
```

3. Configure OSPF.

```
[edit protocols]
user@R1# set ospf area 0.0.0.0 interface fe-1/2/1.5
user@R1# set ospf area 0.0.0.0 interface fe-1/2/2.9
user@R1# set ospf area 0.0.0.0 interface lo0.2
```

4. Configure the routing policies.

```
[edit]
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.

```
[edit]
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.

```
[edit]
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 fe-1/2/0 unit 14 family inet address 10.0.0.14/30
user@R2# set fe-1/2/1 unit 18 family inet address 10.0.0.18/30
user@R2# set fe-1/2/2 unit 21 family inet address 10.0.0.21/30;
user@R2# set lo0 unit 3 family inet address 2.2.2.2/32
```

2. Configure BGP.

```
[edit]
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]
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.

```
[edit]
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.

```
[edit]
user@R2# set routing-options autonomous-system 65500
```

6. Enable indirect next hops in the forwarding plane.

```
[edit]
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 123.

## 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**

```
show krt indirect-next-hop
Indirect Nexthop:
Index: 1048575 Protocol next-hop address: 10.255.3.1
  RIB Table: __mpls-oam__.mpls.0
  Label: Swap 299968
  Policy Version: 0                         References: 1
  Locks: 2                                  0x95bc514
  Flags: 0x3
  INH Session ID: 0xa
  INH Version ID: 1
  Ref RIB Table: unknown
        Next hop: 50.50.244.9 via ge-2/0/2.0
        Label operation: Swap 299968, Push 299792(top)
        Label TTL action: no-prop-ttl, no-prop-ttl(top)
        Session Id: 0x9
      IGP FRR Interesting proto count : 0
```

#### Meaning

The **0x3** 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 from the configuration, this flag changes to **0x2**. Junos MX series routers with Trio Modular Port Concentrator (MPC) chipset supports indirect-next-hop by default and can not be disabled. Thus, even if **indirect-next-hop** is not configured under **forwarding-options**, the feature will work by default. Thus, **0x3** flag is not applicable for Trio Modular Port Concentrator (MPCs).

> **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

Understanding Indirect Next Hops | 121

## Understanding Unicast RPF (Switches)

IN THIS SECTION

To protect against IP spoofing, and some types of denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks, unicast reverse-path-forwarding (RPF) verifies that packets are arriving from a legitimate path. It does this by checking the source address of each packet that arrives on an untrusted ingress interface and, comparing it to the forwarding-table entry for its source address. If the packet is from a valid path, that is, one that the sender would use to reach the destination, the device forwards the packet to the destination address. If it is not from a valid path, the device discards the packet. Unless it is protected against, IP spoofing can be an effective way for intruders to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination.

Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family. Unicast RPF is not supported on interfaces configured as tunnel sources. This affects only the transit packets exiting the tunnel.

There are two modes of unicast RPF, *strict mode*, and *loose mode*. The default is strict mode, which means the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. Strict mode is especially useful on untrusted interfaces (where untrusted users or processes can place packets on the network segment), and for symmetrically routed interfaces (see "When to Enable Unicast RPF" on page 138.) For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at http://www.ietf.org/rfc/rfc3704.txt.

To enable strict mode unicast RPF on a selected customer-edge interface:

[edit interfaces]
user@switch# **set *interface-name* unit 0 family inet *rpf-check***

The other mode is loose mode, which means the system checks to see if the packet has a source address with a corresponding prefix in the routing table, but it does not check whether the receiving interface is the best return path to the packet's unicast source address.

To enable unicast RPF loose mode, enter:

[edit interfaces]
user@switch# **set *interface-name* unit 0 family inet rpf-check mode loose**

**NOTE:** On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see "Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches" on page 140.

### Unicast RPF for Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces. The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

### Unicast RPF Implementation

IN THIS SECTION



#### *Unicast RPF Packet Filtering*

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

#### *Bootstrap Protocol (BOOTP) and DHCP Requests*

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

#### *Default Route Handling*

If the best return path to the source is the default route (**0.0.0.0**) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

**NOTE:** On the EX4300, the default route is not used when the switch is configured in unicast RPF strict mode.

### When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces, and as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled. Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches, as shown in Figure 15 on page 138. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered. A symmetrically routed interface uses the same route in both directions between the source and the destination.

Unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, to with these devices, be sure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.

**Figure 15: Symmetrically Routed Interfaces**

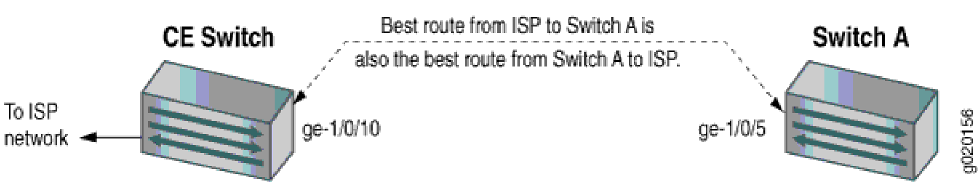


The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

### When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in Figure 16 on page 139. This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

**Figure 16: Asymmetrically Routed Interfaces**

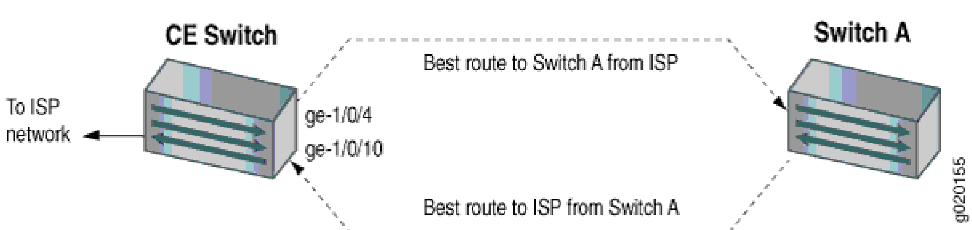

> **NOTE:** Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

### Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.

> **NOTE:** You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches, OCX switches, and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

RELATED DOCUMENTATION

*Example: Configuring Unicast RPF (On a Switch)*

*Troubleshooting Unicast RPF*

# Understanding and Preventing Unknown Unicast Forwarding

**IN THIS SECTION**



Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that traverse a VLAN to all interfaces that are members of that VLAN. Forwarding this type of traffic can create unnecessary traffic that leads to poor network performance or even a complete loss of network service. This flooding of packets is known as a traffic storm.

To prevent a traffic storm, you can disable the flooding of unknown unicast packets to all VLAN interfaces by configuring specific VLANs or all VLANs to forward all unknown unicast traffic traversing them to a specific interface. You can configure multiple VLANs to forward unknown unicast packets to the same interface or configure different interfaces for different VLANs. This channels the unknown unicast traffic traversing VLANs to specific interfaces instead of flooding all interfaces.

## Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface

### Purpose

Verify that a VLAN is forwarding all unknown unicast packets (those with unknown destination MAC addresses) to a single interface instead of flooding unknown unicast packets across all interfaces that are members of that VLAN.

> **NOTE:** This procedure uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, See: "Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface" on page 145. For ELS details see: *Using the Enhanced Layer 2 Software CLI*.

### Action

(EX4300 Switches) Display the forwarding interface for unknown unicast packets for a VLAN (here, the VLAN name is v1):

user@switch> **show configuration switch-options**

```
unknown-unicast-forwarding {
    vlan v1 {
        interface ge-0/0/7.0;
    }
}
```

(EX9200 Switches) Display the forwarding interface for unknown unicast packets:

user@switch> **show forwarding-options**

```
next-hop-group uuf-nhg {
    group-type layer-2;
    interface ge-0/0/7.0;
}
```

#### Meaning

The sample output from the **show** commands show that the unknown unicast forwarding interface for VLAN **v1**is interface **ge-0/0/7**.

### Configuring Unknown Unicast Forwarding (ELS)

IN THIS SECTION



**NOTE:** This task uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see "Configuring Unknown Unicast Forwarding (CLI Procedure)" on page 147. For ELS details, see *Using the Enhanced Layer 2 Software CLI*

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets that traverse a VLAN to all interfaces associated with that VLAN. This flooding of packets is known as a traffic storm and can negatively impact network performance.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN to a specific interface. You can configure each VLAN to divert unknown unicast traffic to a different interface or use the same interface for multiple VLANs.

***Configuring Unknown Unicast Forwarding on EX4300 Switches***

To configure unknown unicast forwarding options on EX4300 switches:

- Configure unknown unicast forwarding for a specific VLAN and specify the interface to which all unknown unicast traffic will be forwarded:

  ```
  [edit switch-options]
  user@switch# set unknown-unicast-forwarding vlan vlan-name interface interface-name
  ```

- Configure unknown unicast forwarding for all VLANs and specify the interface to which all unknown unicast traffic will be forwarded:

  ```
  [edit switch-options]
  user@switch# set unknown-unicast-forwarding vlan all interface interface-name
  ```

***Configuring Unknown Unicast Forwarding on EX9200 Switches***

To configure unknown unicast forwarding on EX9200 switches, you must configure a flood filter and apply it to VLANs for which you want to configure unknown unicast forwarding. Flood filters are firewall filters that are applied only to broadcast, unknown unicast, and multicast (BUM) traffic. If a flood filter is configured, only traffic packets that are of the packet type **unknown-unicast** are forwarded to the interface on which unicast forwarding is configured. A next-hop group redirects the packets according to the action specified in the flood filter.

To configure the next-hop group that receives Layer 2 packets and then configure the interface to which these packets are forwarded:

1. Configure the **next-hop-group** action for the Layer 2 interface expected to receive unknown unicast packets:

   ```
   [edit forwarding-options]
   user@switch# set next-hop-group next-hop-group-name group-type layer-2
   [edit forwarding-options]
   user@switch# set next-hop-group next-hop-group-name interface interface-name
   ```

   For example:

   ```
   [edit forwarding-options]
   user@switch# set next-hop-group uuf-nhg group-type layer-2
   [edit forwarding-options]
   user@switch# set next-hop-group uuf-nhg interface ge-3/1/7.0
   ```

2. Configure a firewall filter with family address type **ethernet-switching**:

   ```
   [edit firewall]
   user@switch# set family ethernet-switching filter filter-name
   ```

   For example:

   ```
   [edit firewall]
   user@switch# set family ethernet-switching filter uuf_filter
   ```

3. Configure a term in the firewall filter for the interface that receives unknown unicast packets (the interface specified in Step 1) to discard unknown unicast packets:

   ```
   [edit firewall family ethernet-switching filter filter-name]
   user@switch# set term term-name from interface interface-name
   user@switch# set term term-name from traffic-type unknown-unicast
   user@switch# set term term-name then discard
   ```

   For example:

   ```
   [edit firewall family ethernet-switching filter uuf_filter]
   user@switch# set term source-drop from interface ge-3/1/7.0
   user@switch# set term source-drop from traffic-type unknown-unicast
   user@switch# set term source-drop then discard
   ```

4. Configure a term in the firewall filter for unknown unicast packets to be flooded to the interface enabled for unknown unicast forwarding by using **next-hop-group** (in step 1):

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name from traffic-type unknown-unicast
user@switch# set term term-name then next-hop-group group-name
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term uuf-flood from traffic-type unknown-unicast
user@switch# set term uuf-flood then next-hop-group uuf-nhg
```

5. Configure a default term for the firewall filter to forward packets other than unknown unicast packets:

```
[edit firewall family ethernet-switching filter filter-name]
user@switch# set term term-name then accept
```

For example:

```
[edit firewall family ethernet-switching filter uuf_filter]
user@switch# set term fwd-default then accept
```

6. Apply the filter as a flood filter on the VLAN that includes the interface which will receive unknown unicast packets:

```
[edit vlans vlan-name]
user@switch# set forwarding-options flood input filter-name
```

For example:

```
[edit vlans v1]
user@switch# set forwarding-options flood input uuf_filter
```

### Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface

#### Purpose

Verify that a VLAN is forwarding all unknown unicast packets (those with unknown destination MAC addresses) to a single trunk interface instead of flooding unknown unicast packets across all interfaces that are members of the same VLAN.

#### Action

Display the forwarding interface for unknown unicast packets for a VLAN (here, the VLAN name is **v1**):

```
user@switch> show configuration ethernet-switching-options
```

```
unknown-unicast-forwarding {
    vlan v1 {
        interface ge-0/0/7.0;
    }
}
```

Display the Ethernet switching table:

`user@switch>` **show ethernet-switching table vlan v1**

```
Ethernet-switching table: 3 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  v1                *                 Flood          - All-members
  v1                00:01:09:00:00:00 Learn         24 ge-0/0/7.0
  v1                00:11:09:00:01:00 Learn         37 ge-0/0/3.0
```

### Meaning

The sample output from the **show configuration ethernet-switching-options** command shows that the unknown unicast forwarding interface for VLAN **v1** is interface **ge-0/0/7**. The **show ethernet-switching table** command shows that an unknown unicast packet is received on interface **ge-0/0/3** with the destination MAC address (DMAC) **00:01:09:00:00:00** and the source MAC address (SMAC) of **00:11:09:00:01:00**. This shows that the SMAC of the packet is learned in the normal way (through the interface **ge-0/0/3.0**), while the DMAC is learned on interface **ge-0/0/7**.

SEE ALSO

## Configuring Unknown Unicast Forwarding (CLI Procedure)

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets to all interfaces associated with a VLAN. Forwarding such traffic to interfaces on the switch can create a security issue.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN out to a specific trunk interface. From there, the destination MAC address can be learned and added to the Ethernet switching table. You can configure each VLAN to divert unknown unicast traffic to different trunk interfaces or use one trunk interface for multiple VLANs.

> **NOTE:** For Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style, see "Configuring Unknown Unicast Forwarding (ELS)" on page 142.

To configure unknown unicast forwarding options:

> **NOTE:** Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

1. Configure unknown unicast forwarding for a specific VLAN (here, the VLAN name is **employee**):

```
[edit ethernet-switching-options]
user@switch# set unknown-unicast-forwarding vlan employee
```

2. Specify the trunk interface to which all unknown unicast traffic will be forwarded:

```
[edit ethernet-switching-options]
user@switch# set unknown-unicast-forwarding vlan employee interface ge-0/0/3.0
```

RELATED DOCUMENTATION

Understanding and Preventing Unknown Unicast Forwarding | 141

*Understanding Storm Control*

*Configuring Autorecovery for Port Security Events*

## Ensuring That Distributed ppm Is Not Disabled

By default, the router's period packet management process (**ppm**) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, **ppm** handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With **ppm** processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run such processes as BFD on the Packet Forwarding Engine.

**Distributed ppm Required for ETH-DM**

Ethernet frame delay measurement requires that **ppm** remains distributed to the Packet Forwarding Engine. If **ppm** is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is *NOT* present:

```
[edit]
routing-options {
    ppm {
        no-delegate-processing;
    }
}
```

If distributed **ppm** processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

**Procedure to Ensure that Distributed ppm is Not Disabled**

To ensure that distributed **ppm** is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed **ppm** is disabled.

   - In the following example, distributed **ppm** is enabled on the router. In this case, you do not need to modify the router configuration:

```
[edit]
user@host# show routing-options
ppm;
```

- In the following example, distributed **ppm** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

```
[edit]
user@host show routing-options
ppm {
    no-delegate-processing;
}
```

2. Modify the router configuration to re-enable distributed **ppm** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppm** is disabled (as determined in the previous step).

   a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

      Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

      - All connectivity fault management (CFM) sessions re-establish.
      - All ETH-DM requests on the router terminate.
      - All ETH-DM statistics and frame counts reset to 0.

   b. Modify the router configuration to re-enable distributed **ppm**. For example:

```
[edit]
user@host# delete routing-options ppm no-delegate-processing
```

   c. Commit the updated router configuration. For example:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
```

   d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management <gracefully | immediately | soft>** operational mode command. For example:

```
user@host> restart ethernet-connectivity-fault-management
Connectivity fault management process started, pid 9893
```

Connectivity fault management (CFM) sessions operate in centralized mode over AE interfaces by default. Y.1731 performance monitoring (PM) is supported on centralized CFM sessions over AE interfaces. Also, distribution of CFM session over AE interfaces to line cards is supported from Junos OS Release 13.3. To enable the distribution of CFM sessions and to operate in centralized mode, include the **ppm delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The mechanism that enables distribution of CFM sessions over AE interfaces provides the underlying infrastructure to support PM over AE interfaces. In addition, periodic packet management (PPM) handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With PPM processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run performance monitoring processes on the Packet Forwarding Engine.

RELATED DOCUMENTATION

*Periodic Packet Management*

*Understanding Periodic Packet Management on MX Series Routers*

CHAPTER 7

# Configuring a Juniper Entropy Beacon

IN THIS CHAPTER



## Understanding the Juniper Entropy Beacon

IN THIS SECTION



### Juniper Entropy Beacon Overview

The Juniper Entropy Beacon (JEB) feeds high quality entropy over the network from a SRX345 Services Gateway to entropy-starved clients. This entropy can be used by any entropy consuming application or device. Entropy is a key component for cryptographic security systems. Generation of symmetric and asymmetric cryptographic keys requires entropy, but low or poor entropy leads to predictable keys, compromising the encryption and overall security of the system. The SRX Services Gateway can produce large amounts of entropy continuously and quickly, making it a reliable source for network based entropy sharing.

Starting in Junos OS release 19.1R1, devices can use JEB to talk to an SRX345 Services Gateway device and request up to 64 KB of entropy at a time. JEB uses its own protocol to respond to these requests. The protocol is wrapped by TLSv1.2, making it cryptographically secure. JEB also forces clients to authenticate using X.509 certificates, so it can detect API abuse and pin incoming requests to respective clients.

JEB makes use of a simple request-response Layer 7 protocol wrapped by TLSv1.2. To send an entropy seed request to a JEB server, connect to the server through TLSv1.2 and send a command in the following format: **{S: *size*}**. The *size* variable is the number of entropy bytes you are requesting. The JEB server will respond by sending an entropy seed of the specified size, as long the size falls within the server's configured parameters.

### Configuring a Juniper Entropy Beacon Server

IN THIS SECTION



This example shows how to configure a Juniper Entropy Beacon (JEB) server on a SRX345 Services Gateway. A JEB server can send high quality entropy over the network to entropy consuming applications and devices.

***Requirements***

This example uses the following hardware and software components:

- A SRX345 Services Gateway
- Junos OS release 19.1R1 or later

You must load in your own certificate-key pairs as well as a trusted CA bundle on to the SRX345 Services Gateway in order to configure the JEB server.

***Overview***

A JEB server can send out entropy seeds through a request-response Layer 7 protocol wrapped by TLSv1.2. You can configure the maximum size of the entropy seeds, as well as the type of random bit generator the JEB server will use for entropy generation. Once the JEB server is properly configured, users with the proper certificate can connect via TLSv1.2 and send requests for entropy seeds. The JEB server will send out an entropy seed of the requested size if it is within the configured parameters.

***Configuration***

#### CLI Quick Configuration

```
user@jebserver# show system services
jeb {
  max-seed-size 4096;
  port 57005;
```

```
    rbg default-rng;
    tls {
        cert-bundle /path/to/cert-bundle;
        certificate /path/to/certificate;
        key /path/to/private-key;
    }
}
```

***Configuring a JEB Server***

**Step-by-Step Procedure**

1. Specify the port that the JEB server will run on. The value can be in a range from 1025 through 65535, with a default value of 57005.

   `[edit system services]`
   `user@jebserver#` **set jeb port *port-number***

2. Configure the maximum entropy seed size that the server will send out in bytes. The range is 1 through 65536 and the default value is 4096.

   `[edit system services]`
   `user@jebserver#` **set jeb max-seed-size *bytes***

3. Choose a random bit generator (RBG) to use for generating entropy seeds. There are 3 options, **default-rng**, **hmac-drbg**, and **jrbc**. The **default-rng** option is a cryptographically secure pseudorandom number generator (CSPRNG), the **hmac-drbg** option is a deterministic RBG detailed in NIST SP 800-90A, and the **jrbc** option is the Juniper random bit conditioner.

   `[edit system services]`
   `user@jebserver#` **set jeb rbg default-rng**

   > NOTE: If you want to use **hmac-drbg**, it must be configured under the **[edit system rng]** hierarchy before being configured for the JEB server.

4. Specify a path to the certificate bundle that is used by the server to authenticate the client.

   `[edit system services]`
   `user@jebserver#` **set jeb tls cert-bundle */path/to/cert-bundle***

5. Configure a path to the server certificate.

   `[edit system services]`

```
user@jebserver# set jeb tls certificate /path/to/certificate
```

6. Set a path to the server key.

```
[edit system services]
user@jebserver# set jeb tls key /path/to/private-key
```

***Results***

After configuring the JEB server, you can request entropy seeds by connecting to the JEB server through TLSv1.2 and sending a request in the following format: **{S: *size*}**. The *size* variable is the number of entropy bytes you are requesting.

If the request is successful the server will respond with the following message:

**response: OK**
**response_code:200**

If the request is unsuccessful the server will respond with one of the following messages:

**response: ERROR**
**response_code: 501**
**data: SizeUnsupportedError**

or

**response: ERROR**
**response_code: 501**
**data: MalformedPacketError**

The **SizeUnsupportedError** error indicates that the requested size is over the configured size of the **max-seed-size** statement, and the **MalformedPacketError** error indicates that the packet was not properly formed.

***Verification***

**IN THIS SECTION**



Once the JEB server is configured, you can check to verify that the process is running by using the **show system processes** command or checking for JEB-related syslog messages.

*Verifying the JEB process is running*

**Purpose**

Verify that your JEB server is configured properly by using the **show system processes** command.

**Action**

`user@jebserver>` **show system processes | grep jeb**

```
37356  ??  S      0:07.00 /usr/bin/python /usr/sbin/jeb -d -f /var/etc/jeb_conf
```

**Meaning**

The output contains references to **/usr/sbin/jeb** and **/var/etc/jeb_conf**, indicating that the JEB process is running.

*Verifying JEB syslog messages*

**Purpose**

Check the syslog messages to see if **init** was able to start JEB. This will also provide the process ID (PID) for JEB.

**Action**

`user@jebserver>` **show log messages | grep jeb**

```
Jan  2 11:12:28  cartier init: jeb (PID 37356) started
```

**Meaning**

The syslog message should indicate that the JEB process has started.

CHAPTER 8

# Reconfiguring Martian Addresses to be Recognized for Routing

IN THIS CHAPTER



## 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
```

RELATED DOCUMENTATION

Example: Removing the Class E Prefix on Martian Addresses | 159

## Example: Removing the Class E Prefix on Martian Addresses

IN THIS SECTION



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 E addresses in the default unicast routing table.

```
[edit routing-options]
user@host# set martians 240.0.0.0/4 orlonger allow
```

2. Allow Class E 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 E 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 E 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

**IN THIS SECTION**



Confirm that the configuration is working properly.

***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

Understanding Martian Addresses | 157

*Example: Creating an Interface on a Logical System*

*Example: Configuring an OSPF Default Route Policy on Logical Systems*

CHAPTER 9

# Protocol-Independent Routing Properties for QFX Series Switches

IN THIS CHAPTER



## Overview of Hierarchical ECMP Groups on QFX5200 Switches

Starting with Junos OS Releases 15.1X53-D30 and 17.2R1, QFX5200 switches support hierarchical equal-cost multipath (ECMP) groups.Hierarchical ECMP is enabled by default and provides for two-level route resolution automatically through the Packet Forwarding Engine. Two-level route resolution through ECMP groups enhances load balancing of traffic. The partition of hierarchical ECMP groups is fixed and cannot be modified. As a result, if groups designated for one type of prefixes, for example, overlay routes are not used, they cannot be designated for use by other types of traffic.

Starting with Junos OS Releases 15.1X53-D210 and 17.2R1, you can disable hierarchical ECMP groups at system start time on QFX5200 switches. Doing so makes all the ECMP groups available for route resolution for all types of prefixes, effectively increasing the number of ECMP groups. To disable hierarchical ECMP, include **no-hierarchical-ecmp** at the **[edit forwarding-options]** hierarchy level. Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart.

**Release History Table**

| Release | Description |
|---|---|
| 15.1X53-D30 | Starting with Junos OS Releases 15.1X53-D30 and 17.2R1, QFX5200 switches support hierarchical equal-cost multipath (ECMP) groups. |
| 15.1X53-D210 | Starting with Junos OS Releases 15.1X53-D210 and 17.2R1, you can disable hierarchical ECMP groups at system start time on QFX5200 switches. |

RELATED DOCUMENTATION

no-hierarchical-ecmp | 307

# Understanding the Unified Forwarding Table

IN THIS SECTION



Traditionally, forwarding tables have been statically defined and have supported only a fixed number of entries for each type of address. The Unified Forwarding Table (UFT) feature enables you to allocate forwarding table resources to optimize the memory available for different address types based on the needs of your network. You can choose to allocate a higher percentage of memory for one type of address or another.

## Using the Unified Forwarding Table to Optimize Address Storage

On QFX5100, EX4600, QFX5200 and OCX1100 switches, you can control the allocation of forwarding table memory available to store the following:

- MAC addresses—In a Layer 2 environment, the switch learns new MAC addresses and stores them in a MAC address table
- Layer 3 host entries–In a Layer 2 and Layer 3 environment, the switch learns which IP addresses are mapped to which MAC addresses; these key-value pairs are stored in the Layer 3 host table.
- Longest prefix match (LPM) table entries—In a Layer 3 environment, the switch has a routing table and the most specific route has an entry in the forwarding table to associate a prefix or netmask to a next hop. Note, however, that all IPv4 /32 prefixes and IPv6 /128 prefixes are stored in the Layer 3 host table.

> **NOTE:** Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range.

**NOTE:** Platform support depends on the Junos OS release in your installation.

### When to Use a United Forwarding Table

This feature gives you the flexibility to configure your switch to match the needs of your particular network environment. For example, you might configure the switch to store more MAC addresses in a Layer 2 network, such as a virtualized network with many servers and virtualized machines. On the other hand, if your switch is located in the routing core of a network or participates in an IP fabric, you probably want to maximize the number of routing table entries it can store. In this case, you would configure it to use the **lpm-profile**, which provides the most longest prefix match table entries.

UFT essentially combines the three distinct forwarding tables to create one table with flexible resource allocation. You can select one of five forwarding table profiles that best meets your network needs. Each profile is configured with different maximum values for each type of address. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized machines, you would likely choose a profile that allocates a higher percentage of memory to MAC addresses. For a switch that operates in the core of a network, participates in an IP fabric, you probably want to maximize the number of routing table entries it can store. In this case, you would choose a profile that allocates a higher percentage of memory to longest match prefixes. The QFX5200 switch also supports a custom profile that allows you to partition the four available shared memory banks with a total of 128,000 entries among MAC addresses, Layer 3 host addresses, and LPM prefixes.

**NOTE:** Support for QFX5200 switches was introduced in Junos OS Release 15.1x53-D30. The QFX5200 switch is not supported on Junos OS Release 16.1R1.

### MAC Address and Host Address Memory Allocation

There are several profiles that allocate memory differently for MAC addresses and host addresses. You configure the mix that best meets your needs by choosing the appropriate profile. Table 4 on page 168 lists the profiles you can choose and the associated maximum values for the MAC address and host table entries on QFX3500, QFX3600, and QFX5100 switches.

Table 4: Unified Forwarding Table Profiles on QFX3500, QFX3600, and QFX5100 Switches

| Profile Name | MAC Table | Host Table (unicast and multicast addresses) | | | | | |
|---|---|---|---|---|---|---|---|
| | MAC Addresses | IPv4 unicast | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) |
| l2-profile-one | 288K | 16K | 8K | 8K | 8K | 4K | 4K |
| l2-profile-two | 224K | 80K | 40K | 40K | 40K | 20K | 20K |
| l2-profile-three (default) | 160K | 144K | 72K | 72K | 72K | 36K | 36K |
| l3-profile | 96K | 208K | 104K | 104K | 104K | 52K | 52K |
| lpm-profile | 32K | 16K | 8K | 8K | 8K | 4K | 4K |
| lpm-profile with unicast-in-lpm option | 32K | (stored in LPM table) | (stored in LPM table) | 8K | 8K | 4K | 4K |

**NOTE:** On QFX5100 and QFX5200 switches, IPv4 and IPv6 host routes with ECMP next hops are stored in the host table. On QFX3500 and QFX3600 switches, these routes are stored in the LPM table.

The QFX5100, and QFX5200 switches all support five profiles, each of which allocates different amounts of memory for Layer 2 or Layer 3 entries, enabling you choose one that best suits the needs of your network. The QFX5200 switch, however, supports different maximum values for each profile from the other switches. The QFX5200 switch also supports a custom profile. For more information about the custom profile, see *Configuring the Unified Forwarding Table on Switches*.

Table 5 on page 168 lists the profiles you can choose and the associated maximum values for the MAC address and host table entries on QFX5200 switches.

Table 5: Unified Forwarding Table Profiles on QFX5200 Switches

| Profile Name | MAC Table | Host Table (unicast and multicast addresses) | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | MAC Addresses | IPv4 unicast | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) | Exact-Match |

Table 5: Unified Forwarding Table Profiles on QFX5200 Switches *(continued)*

| Profile Name | MAC Table | Host Table (unicast and multicast addresses) | | | | | | |
|---|---|---|---|---|---|---|---|---|
| l2-profile-one | 136K | 8K | 4K | 4K | 4K | 2K | 2K | 0 |
| l2-profile-two | 104K | 40K | 20K | 20K | 20K | 10K | 10K | 0 |
| l2-profile-three (default) | 72K | 72K | 36K | 36K | 36K | 18K | 18K | 0 |
| l3-profile | 40K | 104K | 52K | 52K | 52K | 26K | 26K | 0 |
| exact-match-profile | 8K | 8K | 4K | 4K | 4K | 2K | 2K | 64K |
| lpm-profile | 8K | 8K | 4K | 4K | 4K | 2K | 2K | 0 |
| lpm-profile with unicast-in-lpm option | ? | (stored in LPM table) | (stored in LPM table) | ? | ? | ? | ? | 0 |

Note that all entries in the host table share the same memory space. If the host table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate *any* entries of any other type. As you can see, different entry types occupy different amounts of memory. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address. Table 6 on page 169 lists various valid combinations that the host table can store if you use the **l2-profile-one** profile on a QFX5100 switch. Each row in the table represents a case in which the host table is full and cannot accommodate any more entries.

Table 6: Example Host Table Combinations Using l2-profile-one on a QFX5100 Switch

| IPv4 unicast | IPv6 unicast | IPv4 multicast (*, G) | IPv4 multicast (S, G) | IPv6 multicast (*, G) | IPv6 multicast (S, G) |
|---|---|---|---|---|---|
| 16K | 0 | 0 | 0 | 0 | 0 |
| 12K | 2K | 0 | 0 | 0 | 0 |
| 12K | 0 | 2K | 2K | 0 | 0 |
| 8K | 4K | 0 | 0 | 0 | 0 |
| 4K | 2K | 2K | 2K | 0 | 0 |

Table 6: Example Host Table Combinations Using l2-profile-one on a QFX5100 Switch *(continued)*

| IPv4 unicast | IPv6 unicast | IPv4 multicast (*, G) | IPv4 multicast (S, G) | IPv6 multicast (*, G) | IPv6 multicast (S, G) |
|---|---|---|---|---|---|
| 0 | 4K | 0 | 0 | 1K | 1K |

**NOTE:** On QFX5100 and QFX5200 switches, IPv4 and IPv6 host routes with ECMP next hops are stored in the host table.

**NOTE:** The default profile is **l2-profile-three**, which allocates equal space for MAC Addresses and Layer 3 host addresses and space equal to 16,000 IPv4 entries for the LPM table.

**BEST PRACTICE:** If the host or LPM table stores the maximum number of entries for any given type of entry, the entire shared table is full and is unable to accommodate *any* entries of any other type. Different entry types occupy different amounts of memory. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address.

### LPM Table Memory Allocation

You configure the memory allocation for LPM table entries differently depending on which version of Junos OS you use. To learn how to configure memory allocation for LPM table entries see *Configuring the Unified Forwarding Table on Switches*. Starting with Junos OS 14.1X53-D30 you can free memory in the host table by using the **unicast-in-lpm** option with the **lpm-profile** to store IPv4 and IPv6 unicast addresses in the LPM table instead of the host table. See "Configuring the lpm-profile With Junos OS Release 14.1x53-D30 and Later" on page 177.

**Release History Table**

| Release | Description |
|---|---|
| 14.1X53-D30 | Starting with Junos OS 14.1X53-D30 you can free memory in the host table by using the **unicast-in-lpm** option with the **lpm-profile** to store IPv4 and IPv6 unicast addresses in the LPM table instead of the host table. |
| 13.2X51-D15 | Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range. |

RELATED DOCUMENTATION

*Configuring the Unified Forwarding Table on Switches*

## Configuring the Unified Forwarding Table on Switches

**IN THIS SECTION**



Traditionally, forwarding tables have been statically defined and have supported only a fixed number of entries for each type of address stored in the tables. The Unified Forwarding Table feature lets you optimize how your switch allocates forwarding-table memory for different types of addresses. You can choose one of five unified forwarding table profiles. Each profile allocates a different maximum amount of memory for Layer 2, Layer 3 host, and longest prefix match (LPM) entries. In addition to selecting a profile, you can also select how much additional memory to allocate for LPM entries.

Two profiles allocate higher percentages of memory to Layer 2 addresses. A third profile allocates a higher percentage of memory to Layer 3 host address, while a fourth profile allocates a higher percentage of memory to LPM entries. There is a default profile configured that allocates an equal amount of memory to Layer 2 and Layer 3 host addresses with the remainder allocated to LPM entries. For a switch in a virtualized network that handles a great deal of Layer 2 traffic, you would choose a profile that allocates a higher percentage of memory to Layer 2 addresses. For a switch that operates in the core of the network, you would choose a profile that allocates a higher percentage of memory to LPM entries.

On QFX5200 and QFX5210-64C switches only, you can also configure a custom profile that allows you to partition shared memory banks among the different types of forwarding table entries. On QFX5200 switches, these shared memory banks have a total memory equal to 128,000 IPv4 unicast addresses. On QFX5210 switches, these shared memory banks have a total memory equal to 256,000 IPv4 unicast addresses. For more information about configuring the custom profile, see *Example: Configuring a Unified Forwarding Table Custom Profile*.

### Configuring a Unified Forwarding Table Profile

To configure a unified forwarding table profile:

Specify a forwarding-table profile.

```
[edit chassis forwarding-options]
user@switch# set profile-name
```

For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

```
[edit chassis forwarding-options]
user@switch# set l2-profile-one
```

**CAUTION:** When you configure and commit a profile, in most cases the Packet Forwarding Engine automatically restarts and all the data interfaces on the switch go down and come back up (the management interfaces are unaffected).

Starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. This behavior avoids Virtual Chassis or VCF instability after the change propagates to member switches and multiple Packet Forwarding Engines automatically restart at the same time. Instead, a message is displayed at the CLI prompt and logged to the switch’s system log to notify you that the profile change does not take effect until the next time you reboot the Virtual Chassis or VCF. We recommend that you plan to make profile changes only when you can perform a Virtual Chassis or VCF system reboot immediately after committing the configuration update. Otherwise, the Virtual Chassis or VCF could become inconsistent if one or more members have a problem and restart with the new configuration before a planned system reboot activates the change on all members.

**NOTE:** You can configure only one profile for the entire switch.

**NOTE:** The **l2-profile-three** is configured by default.

**NOTE:** If the host table stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. Keep in mind that an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address..

### Configuring the Memory Allocation for Longest Prefix Match Entries

**IN THIS SECTION**



In addition to choosing a profile, you can further optimize memory allocation for longest prefix match (LPM) entries by configuring how many IPv6 prefixes to store with lengths from /65 through /127. The switch uses LPM entries during address lookup to match addresses to the most-specific (longest) applicable prefix. Prefixes of this type are stored in the space for ternary content addressable memory (TCAM). Changing the default parameters makes this space available for LPM entries. Increasing the amount of memory available for these IPv6 prefixes reduces by the same amount how much memory is available to store IPv4 unicast prefixes and IPv6 prefixes with lengths equal to or less than 64.

The procedures for configuring the LPM table are different, depending on which version of Junos OS you are using. In the initial releases that UFT is supported, Junos OS Releases 13.2X51-D10 and 13.2X52-10, you can only increase the amount of memory allocated to IPv6 prefixes with lengths from /65 through /127 for any profile, except for **lpm-profile**. Starting with Junos OS Release 13.2X51-D15, you can also allocate either less or no memory for IPv6 prefixes with lengths in the range /65 through /127, depending on which profile is configured. For the **lpm-profie**, however, the only change you can make to the default parameters is to allocate no memory for these types of prefixes.

***Configuring the LPM Table With Junos OS Releases 13.2X51-D10 and 13.2X52-D10***

In Junos OS Releases 13.2x51-D10 and 13.2X52-D10, by default, the switch allocates memory for 16 IPv6 with prefixes with lengths in the range /65 through /127. You can configure the switch to allocate more memory for IPv6 prefixes with lengths in the range /65 through /127.

To allocate more memory for IPv6 prefixes in the range /65 through /127:

1. Choose a forwarding table profile.

   ```
   [edit chassis forwarding-options]
   user@swtitch# set profile-name
   ```

   For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

   ```
   [edit chassis forwarding-options]
   user@swtitch# set l2-profile-one
   ```

2. Select how much memory to allocate for IPv6 prefixes in the range /65 thorugh 127.

   ```
   [edit chassis forwarding-options profile-name]
   user@swtitch# set num-65-127-prefix number
   ```

   For example, to specify to allocate memory for 32 IPv6 prefixes in the range /65 through 127:

   ```
   [edit chassis forwarding-options l2-profile-one]
   user@switch# set num-65-127-prefix 2
   ```

**NOTE:** When you configure and commit the **num-65-127-prefix *number*** statement, all the data interfaces on the switch restart. The management interfaces are unaffected.

The **num-65-127-prefix *number*** statement is not supported on the **lpm-profile**.

***Configuring the LPM Table With Junos OS Release 13.2x51-D15 and Later***

**IN THIS SECTION**

***Configuring Layer 2 and Layer 3 Profiles With Junos OS Release 13.2x51-D15 or Later***

Starting in Junos OS Release 13.2X51-D15, you can configure the switch to allocate forwarding table memory for as many as 4,000 IPv6 prefixes with lengths in the range /65 through /127 for any profile other than the **lpm-profile** or **custom-profile**. You can also specify to allocate no memory for these IPv6 entries. The default is 1,000 entries for IPv6 prefixes with lengths in the range /65 through /127. Previously, the maximum you could configure was for 2,048 entries for IPv6 prefixes with lengths in the range /65 through /127. The minimum number of entries was previously 16, which was the default.

To specify how much forwarding table memory to allocate for IPv6 prefixes with length in the range /65 through /127:

1. Choose a forwarding table profile.

   ```
   [edit chassis forwarding-options]
   user@swtitch# set profile-name
   ```

   For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

   ```
   [edit chassis forwarding-options]
   user@swtitch# set l2-profile-one
   ```

2. Select how much memory to allocate for IPv6 prefixes in the range /65 thorugh 127.

   ```
   [edit chassis forwarding-options profile-name]
   user@swtitch# set num-65-127-prefix number
   ```

   For example, to specify to allocate memory for 2,000 IPv6 prefixes in the range /65 through 127:

   ```
   [edit chassis forwarding-options l2-profile-one]
   user@switch# set num-65-127-prefix 2
   ```

Starting with Junos OS Release 13.2X51-D15, you can use the **num-65-127-prefix** statement to allocate entries. Table 7 on page 176 shows the numbers of entries that you can allocate. Each row represents a case in which the table is full and cannot accommodate any more entries.

Table 7: LPM Table Combinations for L2 and L3 profiles With Junos OS 13.2X51-D15 and Later

| num-65-127-prefix Value | IPv4 Entries | IPv6 Entries (Prefix <= 64) | IPv6 Entries (Prefix >= 65) |
|---|---|---|---|
| **0** | 16K | 8K | 0K |
| **1** (default) | 12K | 6K | 1K |
| **2** | 8K | 4K | 2K |
| **3** | 4K | 2K | 3K |
| **4** | 0K | 0K | 4K |

**CAUTION:** When you configure and commit a profile change with the **num-65-127-prefix** ***number*** statement, the Packet Forwarding Engine automatically restarts and all the data interfaces on the switch go down and come back up (the management interfaces are unaffected).

However, starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, Packet Forwarding Engines on switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) do not automatically restart upon configuring a unified forwarding table profile change. This behavior avoids Virtual Chassis or VCF instability after the change propagates to member switches and multiple Packet Forwarding Engines automatically restart at the same time. Instead, a message is displayed at the CLI prompt and logged to the switch's system log to notify you that the profile change does not take effect until the next time you reboot the Virtual Chassis or VCF. We recommend that you plan to make profile changes only when you can perform a Virtual Chassis or VCF system reboot immediately after committing the configuration update. Otherwise, the Virtual Chassis or VCF could become inconsistent if one or more members have a problem and restart with the new configuration before a planned system reboot activates the change on all members.

***Configuring the lpm-profile With Junos OS Release 13.2x51-D15 and Later***

Starting with Junos OS Release 13.2X51-D15 you can configure the **lpm-profile** profile not to allocate any memory for IPv6 entries with prefix lengths from /65 through /127. These are the default maximum values allocated for LPM memory for the **lpm-profile** by address type:

- 128K of IPv4 prefixes
- 16K of IPv6 prefixes (all lengths)

**NOTE:** The memory allocated for each address type represents the maximum default value for all LPM memory.

To configure the **lpm-profile** not to allocate forwarding-table memory for IPv6 entries with prefixes from /65 through /127, thus allocating more memory for IPv4:

Specify to disable forwarding-table memory for IPv6 prefixes with lengths in the range /65 through /127.

```
[edit chassis forwarding-options lpm-profile]
user@switch# set prefix-65-127-disable
```

For example, on the QFX5100 and EX4600 switches only, if you use the **prefix-65-127-disable** option, each of the following combinations are valid:

- 100K IPv4 and 28K IPv6 /64 or shorter prefixes.
- 64K IPv4 and 64K IPv6 /64 or shorter prefixes.
- 128K IPv4 and 0K IPv6 /64 or shorter prefixes.
- 0K IPv4 and 128K IPv6 /64 or shorter prefixes.

**NOTE:** On the QFX5200 switches, when you configure the **prefix-65-127-disable** statement, the maximum number of IPv6 entries with prefixes equal to or shorter than 64 is 98,000.

#### *Configuring the lpm-profile With Junos OS Release 14.1x53-D30 and Later*

Starting in Junos OS Release 15.1X53-D30, you can configure the **lpm-profile** profile to store unicast IPv4 and IPv6 host addresses in the LPM table , thereby freeing memory in the host table. Unicast IPv4 and IPv6 addresses are stored in the LPM table instead of the host table, as shown in Table 8 on page 177 for QFX5100 and EX4600 switches. (Platform support depends on the Junos OS release in your installation.) You can use this option in conjunction with the option to allocate no memory in the LPM table for IPv6 entries with prefix lengths in the range /65 through /127. Together, these options maximize the amount of memory available for IPv4 unicast entries and IPv6 entries with prefix lengths equal to or less than 64.

**Table 8: lpm-profile with unicast-in-lpm Option for QFX5100 and EX4600 Switches**

| prefix-65-127-disable | MAC Table | Host Table (multicast addresses) | | | | | | LPM Table unicast addresses) | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | MAC | IPv4 unicast | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) | IPv4 unicast | IPv6 unicast (</65) | IPv6 unicast (>/64) |

Table 8: lpm-profile with unicast-in-lpm Option for QFX5100 and EX4600 Switches *(continued)*

| prefix-65-127-disable | MAC Table | Host Table (multicast addresses) | | | | | | LPM Table unicast addresses) | | |
|---|---|---|---|---|---|---|---|---|---|---|
| No | 32K | 0 | 0 | 8K | 8K | 4K | 4K | 128K | 16K | 16K |
| Yes | 32K | 0 | 0 | 8K | 8K | 4K | 4K | 128K | 128K | 0 |

Starting with Junos Release 18.1R1, you cannot set configure a prefix for the **num-65-127-prefix** statement on non-LPM profiles. You can only enable or disable the **prefix-65-127-disable** statement for the **lpm-profile**.

Table 9 on page 178 lists the situations in which the **prefix-65-127-disable** statement should be enabled or disabled.

Table 9: LPM Table Size Variations on QFX5200-48Y Switches

| Profile Name | Prefix Entries | | |
|---|---|---|---|
| **num-65-127-prefix** | **IPv4 <= /32** | **IPv6 <= /64** | **IPv6 > /64** |
| **Enabled** | > 128K (minimum guaranteed) | 98K | 0K |
| **Disabled** | 128K | 16K | 16K |

On QFX5120 and EX4600 switches, you cannot set configure a prefix for the **num-65-127-prefix** statement on non-LPM profiles. You can only enable or disable the **prefix-65-127-disable** statement for the **lpm-profile**

Table 10 on page 178 lists the situations in which the **prefix-65-127-disable** statement should be enabled or disabled.

Table 10: LPM Table Size Variations on QFX5120 and EX4650 Switches

| Profile Name | Prefix Entries | | |
|---|---|---|---|
| **prefix-65-127-disable** | **IPv4 <= /32** | **IPv6 <= /64** | **IPv6 > /64** |
| **Enabled** | 351K (360,000 approximate) | 168K (172,000 approximate) | 0K |
| **Disabled** | 168K (172,000 approximate) | 64K (65,524 approximate) | 64K (65,524 approximate) |

Note that all entries in each table share the same memory space. If a table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate any entries of any other type. For example, if you use the the **unicast-in-lpm** option and there are 128K IPv4 unicast addresses stored in the LPM table, the entire LPM table is full and no IPv6 addresses can be stored. Similarly, if you use the **unicast-in-lpm** option but do not use the **prefix-65-127-disable** option, and 16K IPv6 addresses with prefixes shorter than /65 are stored, the entire LPM table is full and no additional addresses (IPv4 or IPv6) can be stored.

To configure the **lpm-profile** to store unicast IPv4 entries and IPv6 entries with prefix lengths equal to or less than 64 in the LPM table:

1. Specify the option to store these entries in the LPM table.

   ```
   [edit chassis forwarding-options lpm-profile]
   user@switch# set unicast-in-lpm
   ```

2. (Optional) Specify to allocate no memory for in the LPM table for IPv6 prefixes with length in the range /65 through /127:

   ```
   [edit chassis forwarding-options lpm-profile]
   user@switch# set prefix-65-127-disable
   ```

***Configuring Non-LPM Profiles on QFX5120 and EX4650 Switches***

For non-LPM profiles, each profile provides the option of reserving a portion of the 16K L3-defip table to store IPv6 Prefixes > 64. Because these are 128-bit prefixes, you can have maximum of 8k IPv6/128 entries in the l3-defip table.

1. Choose a forwarding table profile.

   ```
   [edit chassis forwarding-options]
   user@swtitch# set profile-name
   ```

   For example, to specify the profile that allocates the highest percentage of memory to Layer 3 traffic:

   ```
   [edit chassis forwarding-options]
   user@swtitch# set l3-profile
   ```

2. Select how much memory to allocate for IPv6 prefixes in the range /65 thorugh 127.

   ```
   [edit chassis forwarding-options profile-name]
   user@swtitch# set num-65-127-prefix number
   ```

   For example, to specify to allocate memory for 2,000 IPv6 prefixes in the range /65 through 127:

   You can choose between 0 and 4, 1 being the default.

```
[edit chassis forwarding-options l3-profile]
user@switch# set num-65-127-prefix 1
```

**Release History Table**

| Release | Description |
|---|---|
| 18.1R1 | Starting with Junos Release 18.1R1, you cannot set configure a prefix for the **num-65-127-prefix** statement on non-LPM profiles. You can only enable or disable the **prefix-65-127-disable** statement for the **lpm-profile**. |
| 14.1X53-D40 | Starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. |
| 13.2X51-D15 | Starting with Junos OS Release 13.2X51-D15, you can also allocate either less or no memory for IPv6 prefixes with lengths in the range /65 through /127, depending on which profile is configured. |
| 13.2X51-D15 | Starting in Junos OS Release 13.2X51-D15, you can configure the switch to allocate forwarding table memory for as many as 4,000 IPv6 prefixes with lengths in the range /65 through /127 for any profile other than the **lpm-profile** or **custom-profile**. |
| 13.2X51-D15 | Starting with Junos OS Release 13.2X51-D15, you can use the **num-65-127-prefix** statement to allocate entries. |

# Example: Configuring a Unified Forwarding Table Custom Profile

**IN THIS SECTION**

Traditionally, forwarding tables have been statically defined and have supported only a fixed number of entries for each type of address. The Unified Forwarding Table (UFT) feature enables you to optimize how forwarding-table memory is allocated to best suit the needs of your network. This example shows how to configure a Unified Forwarding Table profile that enables you to partition four shared hash memory banks among three different types of forwarding-table entries: MAC addresses, Layer 3 host addresses, and longest prefix match (LPM).

The UFT feature also supports five profiles that each allocate a specific maximum amount of memory for each type of forwarding table entry. Some profiles allocate more memory to Layer 2 entries, while other profiles allocate more memory to Layer 3 or LPM entries. The maximum values for each type of entry are fixed in these profiles. With the custom profile, you can designate one or more shared memory banks to store a specific type of forwarding-table entry. You can configure as few as one or as many as four memory banks in a custom profile. The custom profile thus provides even more flexibility in enabling you to allocate forwarding-table memory for specific types of entries.

## Requirements

This example uses the following hardware and software components:

- One QFX5200 switch
- Junos OS Release 15.1x53-D30 or later.

Before you configure a custom profile, be sure you have:

- Configured interfaces

## Overview

The Unified Forwarding Table custom profile enables you to allocate forwarding-table entries among four banks of shared hash tables with a total memory equal to 128,000 unicast IPv4 addresses, or 32,000 entries for each bank. Specifically, you can allocate one or more of these shared banks to store a specific type of forwarding-table entry. The custom profile does not affect the dedicated hash tables. Those tables remain fixed with 8,000 entries allocated to Layer 2 addresses, the equivalent of 8,000 entries allocated to IPv4 addresses, and the equivalent of 16,000 entries allocated to longest prefix match (LPM) addresses.

In this example, you allocate two memory banks to Layer 3 host addresses, and two memory banks to LPM entries. This means that no shared hash table memory is allocated for Layer 2 addresses. Only the dedicated hash table memory is allocated for Layer 2 addresses in this scenario.

## Configuration

IN THIS SECTION



To configure a custom profile for the Unified Forwarding Table feature on a QFX5200 switch that allocates two shared memory banks for Layer 3 host address and two shared memory banks for LPM entries, perform these tasks:

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode. A commit check is performed to ensure that you have allocated forwarding-table space for no more than four memory banks.

> **CAUTION:** When you configure and commit a profile, the Packet Forwarding Engine restarts and all the data interfaces on the switch go down and come back up.

```
user@switch# set chassis forwarding-options custom-profile
user@switch# set chassis forwarding-options custom-profile l2-entries num-banks 0
user@switch# set chassis forwarding-options custom-profile l3-entries num-banks 2
user@switch# set chassis forwarding-options custom-profile lpm-entries num-banks 2
```

#### *Configuring the Custom Profile*

### Step-by-Step Procedure

To create the custom profile:

1. Specify the **custom-profile** option.

```
[edit chassis forwarding-options]
user@switch# set custom-profile
```

*Configuring the Allocation of Shared Memory Banks*

### Step-by-Step Procedure

To allocate memory for specific types of entries for the shared memory banks:

1. Specify to allocate no shared bank memory for Layer 2 entries.

```
[edit chassis forwarding-options custom-profile]
user@switch# set l2-entries num-banks 0
```

2. Specify to allocate two shared memory banks (or the equivalent of 64,000 IPv4 entries) for Layer 3 host entries.

```
[edit chassis forwarding-options custom-profile]
user@switch# set l3-entries num-banks 2
```

3. Specify to allocate two shared memory banks (or the equivalent of 64,000 IPv4 entries) for LPM entries.

```
[edit chassis forwarding-options custom-profile]
user@switch# set lpm-entries numer-banks 2
```

## Results

From configuration mode, confirm your configuration by entering the show chassis forwarding-options command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch# show chassis forwarding-profile
custom-profile {
  l2-entries {
    num-banks 0;
  }
  l3-entries {
    num-banks 2;
  }
  lpm-entries {
    num-banks 2
  }
}
```

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

**CAUTION:** The Packet Forwarding Engine will restart and all the data interfaces on the switch will go down and come back up.

## Verification

**IN THIS SECTION**



Confirm that the configuration is working properly.

***Checking the Parameters of the Custom Profile***

### Purpose

Verify that the custom profile is enabled.

### Action

```
user@switch> show chassis forwarding-options
```

```
UFT Configuration:
custom-profile
Configured custom scale:
Entry type                   Total scale(K)
L2(mac)                      8
L3 (unicast & multicast)          72
Exact Match                   0
Longest Prefix Match (lpm)          80
num-65-127-prefix = 1K
-------------Bank details for various types of entries------------
Entry type                     Dedicated Bank Size(K)      Shared Bank Size(K)
L2 (mac)                      8                         32 * num shared banks
L3 (unicast & multicast          8                           32 * num shared banks
Exact match                    0                           16 * num shared banks
Longest Prefix match(lpm)           16                          32 * num shared banks
```

### Meaning

The output shows that the custom profile is enabled as configured with two shared memory banks designated for Layer 3 host entries; two shared memory banks designated for LPM entries; and no shared memory allocated for Layer 2 entries.

The total scale(K) field shows the total allocation of memory, that is, the amount allocated through the shared memory banks plus the amount allocated through the dedicated hash tables. The amount allocated through the dedicated hash tables is fixed and cannot be changed. Therefore, Layer 2 entries have 8K of memory allocated only through the dedicated hash table. Layer 3 host entries have 64K of memory allocated through two shared memory banks plus 8K through the dedicated hash table, for a total of 72K of memory. LPM entries have 64K of memory allocated through two shared memory banks plus 16K through the dedicated hash table, for a total of 80K of memory.

# 3
PART

# Troubleshooting

CHAPTER 10

# Troubleshooting Network Issues

IN THIS CHAPTER



## Working with Problems on Your Network

**Problem**

**Description:** This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

**Solution**

**Table 11: Checklist for Working with Problems on Your Network**

| Tasks | Command or Action |
|---|---|
| **"Isolating a Broken Network Connection" on page 188** | |
| 1. Identifying the Symptoms of a Broken Network Connection on page 189 | **ping (*ip-address* \| *hostname*)**<br>**show route (*ip-address* \| *hostname*)**<br>**traceroute (*ip-address* \| *hostname*)** |
| 2. Isolating the Causes of a Network Problem on page 191 | **show < configuration \| interfaces \| protocols \| route >** |
| 3. Taking Appropriate Action for Resolving the Network Problem on page 192 | **[edit]**<br>**delete routing options static route *destination-prefix***<br>**commit and-quit**<br>**show route *destination-prefix*** |

Table 11: Checklist for Working with Problems on Your Network *(continued)*

| Tasks | Command or Action |
|---|---|
| 4. Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 193 | **show route (*ip-address* \| *hostname*)**<br>**ping (*ip-address* \| *hostname*) count 3**<br>**traceroute (*ip-address* \| *hostname*)** |

## Isolating a Broken Network Connection

By applying the standard four-step process illustrated in Figure 17 on page 188, you can isolate a failed node in the network. Note that the functionality described in this section is not supported in versions 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Figure 17: Process for Diagnosing Problems in Your Network

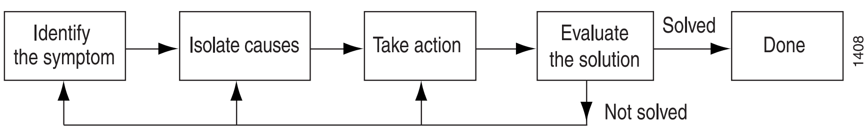


Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic 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 189 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 18: Network with a Problem

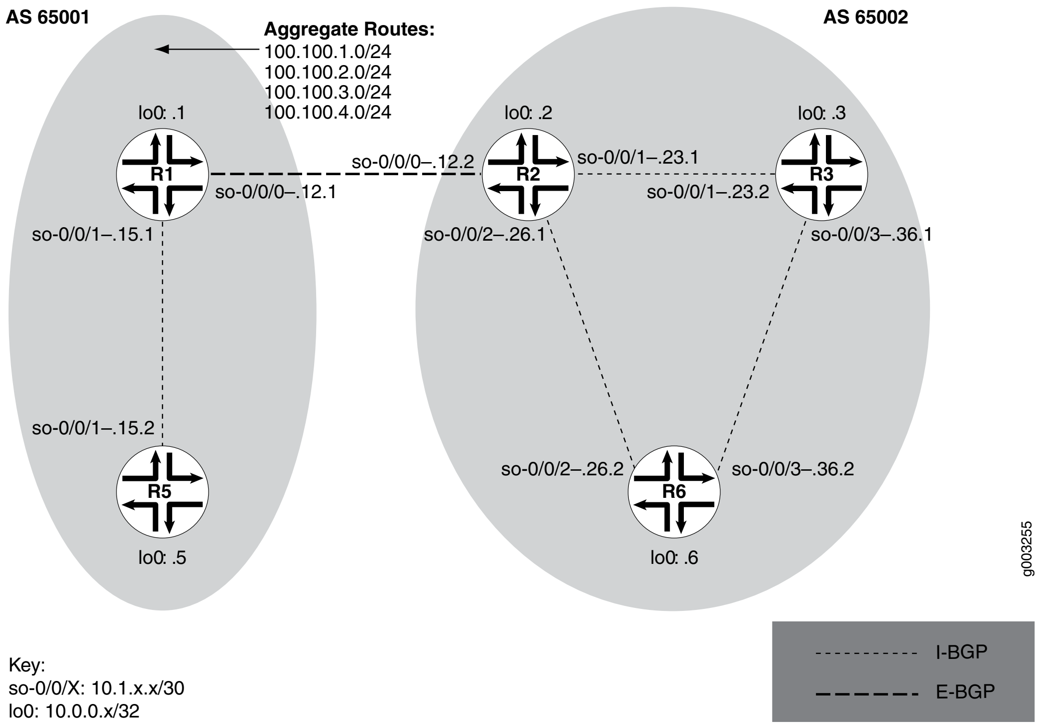


The network in Figure 18 on page 189 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The problem in this network is that **R6** does not have access to **R5** because of a loop between **R2** and **R6**.

To isolate a failed connection in your network, follow the steps in these topics:

- Isolating the Causes of a Network Problem on page 191
- Taking Appropriate Action for Resolving the Network Problem on page 192
- Taking Appropriate Action for Resolving the Network Problem on page 192
- Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 193

## Identifying the Symptoms of a Broken Network Connection

**Problem**

**Description:** The symptoms of a problem in your network are usually quite obvious, such as the failure to reach a remote host.

**Solution**

To identify the symptoms of a problem on your network, start at one end of your network and follow the routes to the other end, entering all or one of the following Junos OS command-line interfaces (CLI) operational mode commands:

```
user@host> ping (ip-address | host-name)
user@host> show route (ip-address | host-name)
user@host> traceroute (ip-address | host-name)
```

**Sample Output**

```
user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2db   0 0000  01  01 a8c6 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2de   0 0000  01  01 a8c3 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2e2   0 0000  01  01 a8bf 10.1.26.2  10.0.0.5

^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[IS-IS/165] 00:02:39, metric 10
                      > to 10.1.26.1   via so-0/0/2.0

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.649 ms  0.521 ms  0.490 ms
 2  10.1.26.2 (10.1.26.2)  0.521 ms  0.537 ms  0.507 ms
 3  10.1.26.1 (10.1.26.1)  0.523 ms  0.536 ms  0.514 ms
 4  10.1.26.2 (10.1.26.2)  0.528 ms  0.551 ms  0.523 ms
 5  10.1.26.1 (10.1.26.1)  0.531 ms  0.550 ms  0.524 ms
```

#### Meaning

The sample output shows an unsuccessful **ping** command in which the packets are being rejected because the time to live is exceeded. The output for the **show route** command shows the interface (**10.1.26.1**) that you can examine further for possible problems. The **traceroute** command shows the loop between **10.1.26.1** (**R2**) and **10.1.26.2** (**R6**), as indicated by the continuous repetition of the two interface addresses.

## Isolating the Causes of a Network Problem

#### Problem

**Description:** A particular symptom can be the result of one or more causes. Narrow down the focus of your search to find each individual cause of the unwanted behavior.

#### Solution

To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

```
user@host> show < configuration | bgp | interfaces | isis | ospf | route >
```

Your particular problem may require the use of more than just the commands listed above. See the appropriate command reference for a more exhaustive list of commonly used operational mode commands.

#### Sample Output

```
user@R6> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.56.2/30
                                   iso
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.26.2/30
                                   iso
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.2/30
                                   iso
[...Output truncated...]
```

The following sample output is from **R2**:

```
user@R2>  show route 10.0.0.5
```

```
inet.0: 22 destinations, 25 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[Static/5] 00:16:21
                    > to 10.1.26.2 via so-0/0/2.0
                    [BGP/170] 3d 20:23:35, MED 5, localpref 100
                      AS path: 65001 I
                    > to 10.1.12.1 via so-0/0/0.0
```

**Meaning**

The sample output shows that all interfaces on **R6** are up. The output from **R2** shows that a static route [**Static/5**] configured on **R2** points to **R6** (**10.1.26.2**) and is the preferred route to **R5** because of its low preference value. However, the route is looping from **R2** to **R6**, as indicated by the missing reference to **R5** (**10.1.15.2**).

## Taking Appropriate Action for Resolving the Network Problem

**Problem**

**Description:** The appropriate action depends on the type of problem you have isolated. In this example, a static route configured on **R2** is deleted from the [**routing-options**] hierarchy level. Other appropriate actions might include the following:

**Solution**

- Check the local router's configuration and edit it if appropriate.
- Troubleshoot the intermediate router.
- Check the remote host configuration and edit it if appropriate.
- Troubleshoot routing protocols.
- Identify additional possible causes.

To resolve the problem in this example, enter the following Junos OS CLI commands:

```
[edit]
user@R2# delete routing-options static route destination-prefix
user@R2# commit and-quit
user@R2# show route destination-prefix
```

### Sample Output

```
[edit]
user@R2# delete routing-options static route 10.0.0.5/32

[edit]
user@R2# commit and-quit
commit complete
Exiting configuration mode

user@R2> show route 10.0.0.5

inet.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[BGP/170] 3d 20:26:17, MED 5, localpref 100
                      AS path: 65001 I
                    > to 10.1.12.1 via so-0/0/0.0
```

### Meaning

The sample output shows the static route deleted from the [**routing-options**] hierarchy and the new configuration committed. The output for the **show route** command now shows the BGP route as the preferred route, as indicated by the asterisk (*).

## Evaluating the Solution to Check Whether the Network Problem Is Resolved

### Problem

**Description:** If the problem is solved, you are finished. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In relation to the network in "Isolating a Broken Network Connection" on page 188, 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

# Debugging and Trace Operations

IN THIS CHAPTER



## 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.

RELATED DOCUMENTATION

Example: Tracing Global Routing Protocol Operations | 196

*Junos OS Administration Library*

## Example: Tracing Global Routing Protocol Operations

IN THIS SECTION



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

Understanding Global Routing Protocol Tracing Operations | 195

CLI Explorer

4

PART

# Configuration Statements and Operational Commands

CHAPTER 12

# Configuration Statements

IN THIS CHAPTER

## 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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

# 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.

> **NOTE:** Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

### 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 and PPP Subscribers*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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

Example: Summarizing Static Routes Through Route Aggregation | 76

*Example: Configuring a Conditional Default Route Policy*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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***(— Removes inactive routes from the forwarding table.
- ***passive***— Retains inactive routes in the forwarding table.

- **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 *metric* type *number***;

**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.
- **next-table *next-table***—Specify the name of the next routing table to the destination. Forwarding for the aggregate prefix is done using this table. When you configure **next-table** to **next-hop**, route lookup is redirected to the route table that the next hop points to.

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 Static Routes Through Route Aggregation | 76

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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

Example: Summarizing Static Routes Through Route Aggregation | 76

*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. See CLI Explorer.

#### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Technology Overview: Understanding the Auto Export Feature

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## 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. See CLI Explorer.

**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:** **xk** to specify KB, **xm** to specify MB, or **xg** 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 for Faster Network Failure Detection | 90

# bfd-liveness-detection (Routing Options Static Route)

## Syntax

```
bfd-liveness-detection {
    description Site- xxx;
    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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**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 for Faster Network Failure Detection | 90

Example: Configuring BFD Authentication for Securing Static Routes | 101

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## 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 Static Routes Through Route Aggregation | 76

*Example: Configuring a Conditional Default Route Policy*

*Understanding Conditionally Generated Routes*

aggregate | 211

generate | 261

# color

### Syntax

```
color {
  metric-value;
  <type metric_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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Fine-grained preference value for a static, aggregate, or generated route.

You can also specify a primary route preference (by including the **color** statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the **color2** statement). You can also mark route preferences with additional route tiebreaker information by specifying a primary route preference and a tiebreaker route preference (by including the **preference** and **preference2** statements in the configuration).

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**

***metric_value***—The metric value for an aggregate, a generated, or a static route.

**Range:** 0 through 4,294,967,295 ($2^{32}$ – 1)

**type *metric_type***—(Optional) External metric type for routes exported by OSPF. 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 | **76**

aggregate | **211**

generate | **261**

static | **365**

preference | **322**

# community (Routing Options)

### Syntax

```
community ([ community-ids ] | no-advertise | no-export | no-export-subconfed | none | llgr-stale | no-llgr);
```

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
**llgr-stale** and **no-llgr** options added in Junos OS Release 15.1.
Support for BGP large community introduced in Junos OS Release 17.3 for MX Series, PTX Series, and QFX Series.

### Description

Associate BGP community information with a static, aggregate, or generated route.

> **NOTE:** BGP large community is available only for static routes.

### 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 Policies, Firewall Filters, and Traffic Policers User Guide*.

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 are advertised to IBGP peers with the same AS number, but not to members of other confederations.
- **llgr-stale**—Adds a community to a long-lived stale route when it is readvertised.
- **no-llgr**—Marks routes which a BGP speaker does not want to be retained by LLGR. The Notification message feature does not have any associated configuration parameters.

> **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 Policies, Firewall Filters, and Traffic Policers User Guide*.

As defined in RFC 8092, BGP large community uses 12-byte encoding and the format for BGP large ***community-ids*** is:

```
large: global-administrator:assigned-number:assigned-number
```

**large** indicates BGP large community.

**global-administrator** is the administrator. It is a 4-byte AS number.

**assigned-number** is a 4-byte value used to identify the local provider. BGP large community uses two 4-byte assigned number to identify the local provider.

**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 | 76
aggregate | 211
generate | 261
static | 365

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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*

*Understanding 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

*High Availability User Guide*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## 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 Static Routes Through Route Aggregation | **76**

aggregate | **211**

generate | **261**

# dynamic-tunnels

## Syntax

```
dynamic-tunnels tunnel-name {
    destination-networks prefix;
    gre;
    rsvp-te entry-name {
        destination-networks network-prefix;
        label-switched-path-template (Multicast) {
            default-template;
            template-name;
        }
    }
    source-address address;
    spring-te;
    traceoptions;
    tunnel-attributes name {
        dynamic-tunnel-anchor-pfe dynamic-tunnel-anchor-pfe;
        dynamic-tunnel-anti-spoof (off | on);
        dynamic-tunnel-gre-key
        dynamic-tunnel-mtu dynamic-tunnel-mtu;
        dynamic-tunnel-source-prefix dynamic-tunnel-source-prefix;
        dynamic-tunnel-type V4oV6;
    }
}
```

## 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.

**NOTE:** ACX Series routers do not support the **gre** statement.

Configure dynamic IPv4-over-IPv6 tunnels and define their attributes to forward IPv4 traffic over an IPv6-only network. IPv4 traffic is tunneled from customer premises equipment to IPv4-over-IPv6 gateways. You must also configure **extended-nexthop** option at **[edit protocols bgp family inet unicast]** hierarchy level to allow BGP to route IPv4 address families over an IPv6 session.

### Options

**gre**—Enable dynamic generic routing encapsulation type tunnel mode for IPv4

**Values:**

- **next-hop-based-tunnel**—Enable next hop base dynamic-tunnel for steering IPv4 traffic with IPv6 next hop address.

**source-address**—Specify the source address of the tunnel.

***tunnel-name***—Name of the dynamic tunnel.

The remaining statements are explained separately. See CLI Explorer.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

## RELATED DOCUMENTATION

*extended-nexthop*

*tunnel-attributes*

*Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks*

*Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the **export-rib** statement.

The **export-rib** statement specifies the source table from which routing information is advertised.

One common use of the **export-rib** statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of **export-rib** is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance*.inet.0). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
  rib-groups {
    rib-interface-routes-v4 {
      import-rib [ inet.0 VRF.inet.0 ];
    }
```

```
    rib-import-VRF-routes-to-inet0-v4 {
        export-rib VRF.inet.0;
        import-rib [ VRF.inet.0 inet.0 ];
        import-policy rib-import-VRF-routes-to-inet0-v4;
    }
    rib-import-inet0-routes-to-VRF-v4 {
        export-rib inet.0;
        import-rib [ inet.0 VRF.inet.0 ];
        import-policy rib-import-inet0-routes-to-VRF-v4;
    }
}
routing-options {
    interface-routes {
        rib-group {
            inet rib-interface-routes-v4;
        }
    }
}
protocols {
    bgp {
        group iBGP-peers {
            type internal;
            family inet {
                unicast {
                    rib-group rib-import-inet0-routes-to-VRF-v4;
                }
            }
        }
    }
}
```

```
routing-instances {
    VRF {
        routing-options {
            interface-routes {
                rib-group {
                    inet rib-interface-routes-v4;
                }
            }
        }
        protocols {
            bgp {
                group peersin-VRF {
                    family inet {
                        unicast {
                            rib-group rib-import-VRF-routes-to-inet0-v4;
                        }
                    }
                }
            }
        }
    }
}
```

**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 | **51**

*Example: Configuring a PIM RPF Routing Table*

*Example: Configuring DVMRP to Announce Unicast Routes*

*Example: Configuring a Dedicated PIM RPF Routing Table*

*Example: Configuring Any-Source Multicast for Draft-Rosen VPNs*

import-rib | **268**

passive | **315**

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 Alternate Backup Paths Using Fate Sharing*

*MPLS Applications User Guide*

# 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 Forwarding Table Filters*

# 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

For PTX Series routers, this statement appears in the default configuration, preventing installation of flow-specification firewall filters into dfwd. For other models, this setting is omitted from the default configuration, allowing installation of flow-specification firewall filters into dfwd.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

*Example: Enabling BGP to Carry Flow-Specification Routes*

*Understanding BGP Flow Routes for Traffic Filtering*

# 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 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.
**firewall-install-disable** statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## Description

Configure a flow route.

## Default

**legacy**

## Options

***actions***—An action to take if conditions match.

**firewall-install-disable**—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).

**Default:** For PTX Series routers, the **firewall-install-disable** statement appears in the default configuration, preventing installation of flow-specification firewall filters into dfwd. For other models, this setting is omitted from the default configuration, allowing installation of flow-specification firewall filters into dfwd.

***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. See CLI Explorer.

**Required Privilege Level**
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

*Example: Enabling BGP to Carry Flow-Specification Routes*

*Understanding BGP Flow Routes for Traffic Filtering*

# forwarding-table

### Syntax

```
forwarding-table {
    chained-composite-next-hop;
    ecmp-fast-reroute,
    export [ policy-name ];
    (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 14.1X53-D30 for the QFX Series.

### Description

Configure information about the routing device's forwarding table.

The remaining statements are explained separately. See CLI Explorer.

### 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*

# forwarding-options

### Syntax

```
forwarding-options {
  dhcp-security {
    arp-inspection;
    group group-name {
      interface interface-name {
        static-ip ip-address {
          mac mac-address;
        }
      }
      overrides {
        no-option82;
        (trusted | untrusted);
      }
    }
    ip-source-guard;
    no-dhcp-snooping;
    option-82 {
      circuit-id {
        prefix {
          host-name;
          logical-system-name;
          routing-instance-name;
        }
        use-interface-description (device | logical);
        use-vlan-id;
      }
      remote-id {
        host-name hostname;
        use-interface-description (device | logical);
        use-string string;
      }
      vendor-id {
        use-string string;
      }
    }
  }
  filter (VLANs) {
    input filter-name;
    output filter-name;
  }
  flood {
```

```
    input filter-name;
  }
```

**Chassis: EX4600 and QFX Series**

```
forwarding options profile-name {
  num-65-127-prefix number;
}
```

**Chassis: EX4600 and QFX Series**

```
forwarding-options lpm-profile {
  prefix-65-127-disable;
  unicast-in-lpm;
}
```

**Chassis: EX4600 and QFX Series**

```
forwarding-options custom-profile {
  l2-entries | l3-entries | lpm-entries {
    num-banks number;
  }
}
```

### Hierarchy Level

```
[edit],
[edit bridge-domains bridge-domain-name],
[edit vlans vlan-name]
```

```
[edit chassis (QFX Series)]
```

### Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.3 for QFX Series switches.
Hierarchy level **[edit vlans *vlan-name*]** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Hierarchy level **[edit bridge-domains *bridge-domain-name*]** introduced in Junos OS Release 14.1 for MX Series routers.
**custom-profile** option introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches only.

**Description**

Configure a unified forwarding table profile to allocate the amount of memory available for the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match table entries.

This feature enables you to select a profile that optimizes the amount of memory available for various types of forwarding-table entries based on the needs of your network. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized machines, you would choose the **l2-profile-one**, which allocates the highest amount of memory to MAC addresses.

You configure the memory allocation for LPM table entries differently, depending on whether you using Junos OS Release 13.2X51-D10 or Junos OS Release 13.2X51-D15 and later. For more information about configuring memory allocation for LPM table entries, see *Configuring the Unified Forwarding Table on Switches*.

The **num-65-127-prefix *number*** statement is not supported on the **custom-profile** and the **lpm-profile**. The **prefix-65-127-disable** and **unicast-in-lpm** statements are supported only on the **lpm-profile**.

When you commit a configuration with a forwarding table profile change, in most cases the Packet Forwarding Engine restarts automatically to apply the new parameters, which brings the data interfaces down and then up again.

However, starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. This behavior avoids having Virtual Chassis or VCF instability and a prolonged convergence period if a profile change is propagated to member switches and multiple Packet Forwarding Engines all restart at the same time. In this environment, instead of automatically restarting when you initially commit a profile configuration change, the message **Reboot required for configuration to take effect** is displayed at the master switch CLI prompt, notifying you that the profile change does not take effect until the next time you restart the Virtual Chassis or VCF. The profile configuration change is propagated to member switches that support this feature, and a reminder that a reboot is required to apply this pending configuration change appears in the system log of the master switch and applicable member switches. You then enable the profile change subsequently during a planned downtime period using the **request system reboot** command, which quickly establishes a stable Virtual Chassis or VCF with the new configuration.

**NOTE:** You should plan to make unified forwarding table profile changes only when you are ready to perform a Virtual Chassis or VCF system reboot *immediately* after committing the configuration update. Otherwise, in the intervening period between committing the configuration change and rebooting the Virtual Chassis or VCF, the system can become inconsistent if a member experiences a problem and restarts. In that case, the new configuration takes effect on the member that was restarted, while the change is not yet activated on all the other members.

The remaining statements are explained separately. See CLI Explorer.

**Options**

**profile-name**—name of the profile to use for memory allocation in the unified forwarding table.

Table 12 on page 259 lists the profiles you can choose that have set values and the associated values for each type of entry.

On QFX5200 Series switches only, you can also select **custom-profile**. This profile enables you to allocate from one to four banks of shared hash memory to a specific type of forwarding-table entry. Each shared hash memory bank can store a maximum of the equivalent of 32,000 IPv4 unicast addresses.

**Table 12: Unified Forwarding Table Profiles**

| Profile Name | MAC Table | Host Table (unicast and multicast addresses) | | | | | |
|---|---|---|---|---|---|---|---|
| | **MAC Addresses** | **IPv4 unicast** | **IPv6 unicast** | **IPv4 (*, G)** | **IPv4 (S, G)** | **IPv6 (*, G)** | **IPv6 (S, G)** |
| **l2-profile-one** | 288K | 16K | 8K | 8K | 8K | 4K | 4K |
| **l2-profile-two** | 224K | 80K | 40K | 40K | 40K | 20K | 20K |
| **l2-profile-three** (default) | 160K | 144K | 72K | 72K | 72K | 36K | 36K |
| **l3-profile** | 96K | 208K | 104K | 104K | 104K | 52K | 52K |
| **lpm-profile*** | 32K | 16K | 8K | 8K | 8K | 4K | 4K |

* This profile supports only IPv4 in Junos OS Release13.2X51-D10. Starting in Junos OS Relase13.2X51-D15, the **lpm-profile** supports IPv4 and IPv6 entries.

**NOTE:** If the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see "Understanding the Unified Forwarding Table" on page 166.

**l2-entries | l3-entries | lpm-entries**—(**custom-profile only**) Select a type of forwarding-table entry—Layer 2, Layer 3, or LPM—to allocate a specific number of shared memory banks. You configure the amount of memory to allocate for each type of entry separately.

**num-banks** ***number***—(**custom-profile only**) Specify the number of shared memory banks to allocate for a specific type of forwarding-table entry. Each shared memory bank stores the equivalent of 32,000 IPv4 unicast addresses.

**Range:** 0 through 4.

**NOTE:** There are four shared memory banks, which can be allocated flexibly among the three types of forwarding-table entries. To allocate no shared memory for a particular entry type, specify the number **0**. When you commit the configuration, the system issues a commit check to ensure that you have not configured more than four memory banks. You do not have to configure all four shared memory banks. By default, each entry type is allocated the equivalent of 32,000 IPv4 unicast addresses in shared memory.

**Required Privilege Level**

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding the Unified Forwarding Table | **166**

*Example: Configuring a Unified Forwarding Table Custom Profile*

*Configuring Traffic Forwarding and Monitoring*

## 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 *metric* type *number*;**

**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.

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## Description

You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. To enable graceful restart globally, include the graceful-restart statement under the **[edit routing options]** hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify the global settings at the individual protocol level.

NOTE:

- 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.
- LDP sessions flap when **graceful-restart** configurations change.

**Default**

Graceful restart is disabled by default.

**Options**

The remaining statements are explained separately. See CLI Explorer.

**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*

*Configuring Graceful Restart for QFabric Systems*

# import

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 | 51
export-rib | 245

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 | **51**

export-rib | **245**

passive | **315**

# 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.

> **NOTE:** The **[edit logical-systems]** hierarchy level is not applicable in ACX Series routers.

### 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

*Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains*

*Configuring Layer 3 VPNs to Carry IBGP Traffic*

autonomous-system | **217**

# 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 14.1X53-D30 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.
- By default, indirectly connected next hops are enabled on PTX Series routers.

### 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 | **122**

# 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

Configure the routing protocol process (rpd) to request an acknowledgement when creating a new forwarding next hop.

During an indirect next-hop change sequence, the routing device might create a new forwarding next hop that is referenced by the indirect next hop. If the **indirect-next-hop-change-acknowledgements** statement is configured, the routing protocol process requests an acknowledgement when creating the new forwarding next hop. When the routing protocol process receives the acknowledgement, this indicates that all PICs have received the new forwarding next hop and it is then safe to change the indirect next hop to reference the new forwarding next hop. This prevents packet loss when changing the indirect next hop by ensuring that all PICs have consistent state information for the new forwarding next hop.

The routing protocol process is not requesting an acknowledgement for the indirect next hop itself. Rather, the routing protocol process is requesting an acknowledgement for the new forwarding next hop that the indirect next hop is going to reference. In the case when the forwarding next hop is an existing one (meaning that it is already installed in the forwarding table), the routing protocol process does not request an acknowledgement, even if the **indirect-next-hop-change-acknowledgements** statement is configured.

We recommend that the **indirect-next-hop-change-acknowledgements** statement be configured when protection mechanisms are being used. This includes MPLS RSVP protection such as fast reroute (FRR) as well as interior gateway protocol (IGP) loop-free alternate (LFA) link or node protection. If there is no protection mechanism being used in the network, the **indirect-next-hop-change-acknowledgements** statement does not provide any benefit and might increase packet loss.

### 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 | **122**

krt-nexthop-ack-timeout | **287**

# 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 Forwarding Table Filters*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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

static | 365

## 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 Policies, Firewall Filters, and Traffic Policers User Guide*

# 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 Policies, Firewall Filters, and Traffic Policers User Guide*

# 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 | 302

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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. See CLI Explorer.

**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-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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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. See CLI Explorer.

**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*

*Example: Configuring Multiple Routing Instances of OSPF*

passive | 315

# jeb

### Syntax

```
jeb {
    max-seed-size max-seed-size;
    port port;
    rbg (default-rng | hmac-drbg);
    tls {
        cert-bundle cert-bundle;
        certificate certificate;
        key key;
    }
}
```

### Hierarchy Level

```
[edit system services]
```

### Release Information

Statement introduced in Junos OS Release 19.1R1.

### Description

Configure a Juniper Entropy Beacon (JEB) server. JEB can be used to feed high quality entropy over the network from a SRX345 Services Gateway to entropy-starved clients.

### Options

**max-seed-size**—Maximum allowed size in bytes for a requested entropy seed

**Default:** 4096

**Range:** 1 through 65536

**port**—Port to use for JEB service

**Default:** 57005

**Range:** 1025 through 65535

**rbg**—Type of random bit generator (RBG) to use for generating entropy seeds

**Values:**

- default-rng—Default cryptographically secure pseudorandom number generator (CSPRNG)
- hmac-drbg—Deterministic RBG detailed in NIST SP 800-90A

**NOTE:** If you want to use **hmac-drbg**, it must be configured under the **[edit system rng]** hierarchy before being configured for the JEB server.

**TLS**—Configure TLS attributes for JEB services

**cert-bundle**—Path to the certificate bundle that is used by the server to authenticate the client

**certificate**—Path to the server certificate

**key**—Path to the server key

**Required Privilege Level**
admin

# 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 | 122

indirect-next-hop-change-acknowledgements | 274

# longest-match (Static Routes)

### Syntax

```
longest-match;
```

### Hierarchy Level

[edit logical-systems *logical-system-name* **routing-options static** route *destination-prefix* next-hop *address* resolve]
[edit routing-instances *routing-instance-name* **routing-options static** route *destination-prefix* next-hop *address* resolve],
[edit **routing-options static** route *destination-prefix* next-hop *address* resolve],

### Release Information

Statement introduced in Junos OS Release 15.1 for EX Series switches.

### Description

Specify the static route on the device to resolve and determine the packet's next-hop interface using the Longest Match Routing Rule (most specific entry), sometimes referred to as the longest prefix match or maximum prefix length match. The Longest Match Routing Rule is an algorithm used by IP routers to select an entry from a routing table. The router uses the longest (prefix) match to determine the egress (outbound) interface and the address of the next device to which to send a packet. Typically, the static route prefers the directly connected subnet route for resolving the next hop rather than performing a longest prefix match with any other available routes.

> **NOTE:** (Required) You must include the **resolve** next-hop option to specify the **longest-match** statement. Next-hop options define additional information about static routes that are included with the route when it is installed in the routing table. You alter the default next-hop resolution behavior using the **resolve** next-hop option.

The router implements the Longest Match Routing Rule as follows:

1. The router receives a packet.

2. While processing the header, the router compares the destination IP address, bit-by-bit, with the entries in the routing table.

The entry that has the longest number of network bits that match the IP destination address is always the best match (or best path) as shown in the following example:

Longest Match Example

- The router receives a packet with a destination IP address of 192.168.1.33.
- The routing table contains the following possible matches:
  - 192.168.1.32/28
  - 192.168.1.0/24
  - 192.168.0.0/16

To determine the longest match, it's easiest to convert the IP addresses in Table 13 on page 289 to binary and compare them.

**Table 13: Converted IP Addresses**

| Address | Converted Binary Address |
|---|---|
| 192.168.1.33 (destination IP address) | **11000000.10101000.00000001.001**00001 |
| 192.168.1.32/28 | **11000000.10101000.00000001.001**00000 (<–Best match) |
| 192.168.1.0/24 | **11000000.10101000.00000001.**00000000 |
| 192.168.0.0/16 | **11000000.10101000.**00000000.00000000 |

**NOTE:** When determining the next-hop interface for customer deployments, setting the **longest-match** statement results in traffic loss.

**Required Privilege Level**
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

## RELATED DOCUMENTATION

Understanding Static Route Preferences and Qualified Next Hops | 27

# 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. See CLI Explorer.

### 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 Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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: Removing the Class E Prefix on Martian Addresses | **159**

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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

aggregate | 211
generate | 261

static | 365
Example: Summarizing Static Routes Through Route Aggregation | 76
Understanding Route Aggregation | 67

## 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 | 330
static | 365

# multicast (Routing Options)

## Syntax

```
multicast {
    cont-stats-collection-interval interval;
    forwarding-cache {
        threshold suppress value <reuse value>;
    }
    interface interface-name {
        enable;
    }
    local-address address
    omit-wildcard-address
    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. See CLI Explorer.

**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 User Guide*

(indirect-next-hop | 272 | 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.

## 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.

> **NOTE:** Starting in Junos OS Release 15.1R4, the router no longer supports a configuration where a static route points to a next hop that is tied to a subscriber. Typically, this might occur when RADIUS assigns the next hop with the Framed-IP-Address attribute. An alternative to this misconfiguration is to have the RADIUS server provide a Framed-Route attribute that matches the static route.

### 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.

# 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.

# no-hierarchical-ecmp

### Syntax

```
no-hierarchical-ecmp;
```

### Hierarchy Level

```
[edit forwarding-options]
```

### Release Information

Statement introduced in Junos OS Release 15.1X53-D210 for QFX5200 switches.

### Description

Disable hierarchical equal-cost multipath (ECMP) groups at system start time. Hierarchical ECMP is enabled by default and provides for two-level route resolution. Disabling hierarchical ECMP effectively increases the number of ECMP groups available for route resolution because hierarchical ECMP allocates separate groups for overlay and underlay routes. Increasing the number of ECMP groups available is potentially useful in a virtual extensible LAN (VXLAN) environment or when MPLS LDP forwarding equivalence classes (FECs) are configured.

> **CAUTION:** Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart.

To reenable hierarchical ECMP, use the following command: **delete forwarding-options no-hierarchical-ecmp**. Reenabling hierarhical ECMP also causes the Packet Forwarding Engine to restart.

### Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Overview of Hierarchical ECMP Groups on QFX5200 Switches | 165

# nonstop-routing

### Syntax

```
nonstop-routing;
```

### Hierarchy Level

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

**NOTE:** Although **nonstop-routing** is also a valid keyword at the **logical-systems** hierarchy level, it is not supported.

### 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.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

### 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*

# num-65-127-prefix

### Syntax

```
num-65-127-prefix number;
```

### Hierarchy Level

```
[edit chassis (QFX Series) forwarding-options profile-name]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for QFX Series switches.
Support for QFX5200 Series switches introduced in Junos OS Release 15.1X53-D30.

### Description

For the Unified Forwarding Table (UFT) feature, specify how much forwarding table memory to allocate for IPv6 entries with prefix lengths in the range of /65 through /127. The ability to allocate flexibly the memory for IPv6 entries with prefixes in this range extends the use of this memory space to accommodate the appropriate mix of longest-prefix match (LPM) entries that best suits your network. The LPM table stores IPv4 unicast prefixes, IPv6 prefixes with lengths equal to or less than 64, and IPv6 prefixes with lengths from 65 through 127. With this option, you can increase, decrease, or allocate no memory for IPv6 prefixes with lengths from 65 through 127, depending on which version of Junos OS you are using.

> **NOTE:** This statement is supported only for the following forwarding table memory profiles: **l2-profile-one**, **l2-profile-three**, **l2-profile-two**, and **l3-profile**. Do not use this statement with the **custom-profile** or the **lpm-profile** statements.

> **NOTE:** The values you can configure are different depending on the version of Junos OS you are using.

### Options

**number**—Specify a numerical value.

**Range:** (Junos OS Release 13.2x51-D10 only) 1 through 128. Each increment represents 16 IPv6 prefixes with lengths in the range of /65 through /127, for a total maximum of 2,058 prefixes (16 x 128 = 2,048).

**Default:** 1 (16 IPv6 prefixes with lengths in the range of /65 through /127).

**Range:** (Junos OS Release 13.2X51-D15 or later) 0 through 4. Each increment allocates memory for 1,000 IPv6 prefixes with lengths in the range of /65 through /127, for a maximum of 4,000 such IPv6 prefixes.

**Default:** 1 (1,000 IPv6 prefixes with lengths in the range of /65 through /127).

**Required Privilege Level**

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

## RELATED DOCUMENTATION

*Configuring the Unified Forwarding Table on Switches*

## 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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*

# p2mp-ldp-next-hop

### Syntax

```
p2mp-ldp-next-hop {
root-address root-address{
   lsp-id id;
      }
   }
```

### 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 in Junos OS Release 13.3.

### Description

Specify a point-to-multipoint LDP label-switched path (LSP) as the next hop for a static route, and configure a root and provide an lsp-id on that LDP-signalled label-switched path.

### Options

**root-address *root address***— Specify the root address of the point-to-multipoint LSP.

**lsp-id *id***— Specify the generic LSP identifier. The range is 1 through 65535.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

# 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. See CLI Explorer.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

*Configuring Static Unicast Routes for Point-to-Multipoint LSPs*

*Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP*

*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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## 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 Static Routes Through Route Aggregation | 76
aggregate | 211
generate | 261

# ppm

### Syntax

```
ppm {
    no-delegate-processing;
}
```

### Hierarchy Level

```
[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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Releases 15.1X49-D10, 15.1X49-D20, 15.1X49-D30, 15.1X49-D35, 15.1X49-D40, 15.1X49-D45, and 15.1X49-D50 for the SRX Series.

### 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 such as:

- Bidirectional Forwarding Detection (BFD)
- Connectivity Fault Management (CFM)
- Link Aggregation Control Protocol (LACP)
- Link Fault Management (LFM)
- Real-time Performance Monitoring (RPM)

- Spanning Tree Protocol (STP)
- Synchronous Ethernet (SYNCE)
- Virtual Router Redundancy Protocol (VRRP)

> **NOTE:** The protocols BFD and CFM are supported on SRX300, SRX320, SRX340, SRX345 and SRX1500 devices.
>
> The protocol LACP is supported on SRX5400, SRX 5600, and SRX5800 devices.

**Options**

**no-delegate-processing**—Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.

**redistribution-timer**— Ensures that link aggregation (and STP) work properly for the periodic packet management (PPM) daemons on the aggregation and satellite devices. A value of 120 is recommended for MXVC-ISSU.

**Required Privilege Level**

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

*Ensuring That Distributed ppm Is Not Disabled*

# precision-timers-max-period

### Syntax

```
precision-timers-max-period precision-timers-max-period;
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 16.1.

### Description

Support of precision-timers in the kernel is a feature where the kernel takes over auto-generation of BGP keepalives right after the switchover from standby to master event occurs. The kernel in the RE continues this auto-generation until the BGP protocol is able to take over the session or until a maximum period has elapsed since the switchover event occurred. The maximum period for which the kernel auto-generates keepalives on behalf of BGP after a switchover event from standby to master ranges from 60 seconds to 1800 seconds. The default value is 600 seconds.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

# 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.

# preference (Routing Options)

### Syntax

```
preference {
    metric-value;
    <type metric_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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value, as well as color values, which are even finer-grained preference values.

You can specify a primary route preference (by including the **preference** statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the **preference2** statement). You can also mark route preferences with additional route tiebreaker information by specifying a color and a tiebreaker color (by including the **color** and **color2** statements in the configuration).

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**

***metric_value***—The metric value for an aggregate, a generated, or a static route to determine the best route among multiple routes to a destination

**Range:** 0 through 4,294,967,295 ($2^{32}$ – 1)

**Default:** 5 (for static routes), 130 (for aggregate and generated routes)

**type *metric_type***—(Optional) External metric type for routes exported by OSPF. 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 | **76**

aggregate | **211**

generate | **261**

static | **365**

color | **230**

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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*

# prefix-65-127-disable

### Syntax

```
prefix-65-127-disable;
```

### Hierarchy Level

```
[edit chassis (QFX Series) forwarding-options lpm-profile]
```

### Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for QFX Series switches.
Support introduced in Junos OS Release 15.1X53-D30 for QFX5200 Series switches.
Support introduced in Junos OS Release 18.1R1 for QFX5200-48C and QFX5210 switches.

### Description

For the Unified Forwarding Table (UFT) feature, specify not to allocate any memory for IPv6 prefixes with lengths in the range /65 through /127 for longest-prefix-match (LPM) entries. Doing so increases the memory available for LPM entries for IPv4 unicast prefixes and IPv6 prefixes with lengths equal to or less than 64. The maximum default value for LPM entries is 16,000 IPv6 prefixes of all lengths.

In an environment where the switch is being used in the core of the network, for example, it might not need to store IPv6 prefixes with lengths in the range /65 through /127. IPv6 prefixes of this type are not typically used in the core.

> **NOTE:** When using this statement, IPv6 prefixes within the range /65 through /127 will still appear in the routing table, but will *not* be installed in the forwarding table; therefore, matching traffic will be dropped. Note further that if a default route is configured, traffic will be forwarded, though it will be sent through the RE and rate-limited.

> **NOTE:** On QFX5100 switches, when you configure this statement, the maximum number of LPM IPv6 entries with prefix lengths equal to or less than 64 increases to 128,000. On the QFX5200 switch, when you configure this statement, the maximum number of IPv6 entries with prefix lengths equal to or less 64 that are allocated in the LPM table increases to 98,0000.

> **NOTE:** This statement is supported only with the **lpm-profile**. No other profile is supported.

The effects of this statement can be seen on a QFX5100 as follows:

```
[edit]
```

`user@host#` **set chassis forwarding-options lpm-profile prefix-65-127-disable**

```
[edit]
```

`user@host#` **commit**

```
configuration check succeeds
commit complete
```

```
[edit]
```

`user@host#` **run show chassis forwarding-options**

```
fpc0:
--------------------------------------------------------------------------
Current UFT Configuration:
lpm-profile. (MAC: 32K L3-host: 16K LPM: 128K)
prefix-65-127 = disable
```

```
[edit]
```

`user@host#` **run show pfe route summary hw**

```
Slot 0
================ fpc0 ================

Unit: 0
Profile active: lpm-profile
Type            Max       Used      Free      % free
----------------------------------------------------
IPv4 Host       16384     20        16354     99.82
IPv4 LPM        131072    5         131065    99.99
IPv4 Mcast      8192      0         8177      99.82

IPv6 Host       8192      5         8177      99.82
IPv6 LPM(< 64)  131072    2         131065    99.99
```

```
IPv6 LPM(> 64) 0     0     0     0.00
IPv6 Mcast        4096      0          4089      99.83
```

### Options

**None**—This statement has no options.

### Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

#### RELATED DOCUMENTATION

*Configuring the Unified Forwarding Table on Switches*

Understanding the Unified Forwarding Table | 166

## 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.

# 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.

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**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. See CLI Explorer.

**Required Privilege Level**

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection | 110

# readvertise

## 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## 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: Preventing a Static Route from Being Readvertised | **38**

Understanding Static Route Control in Routing and Forwarding Tables | **36**

static | **365**

# resolution

### Syntax

```
resolution {
    rib routing-table-name {
        import [ policy-names ];
        inet-import [ policy-names ];
        inet-resolution-ribs [ routing-table-names ];
        inet6-import [ policy-names ];
        inet6-resolution-ribs [ routing-table-names ];
        iso-import [ policy-names ];
        iso-resolution-ribs [ routing-table-names ];
        resolution-family resolution-family;
        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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
inet6-import and inet6-resolution-ribs options added in Junos OS Release 17.2R1.

### 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 the second-listed routing table. Hence, only one routing table is used for each protocol next-hop 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 Overview for EX Series Switches*.

The remaining statements are explained separately. See CLI Explorer.

**Options**

**inet-import [ *policy-names* ]**—(Optional) Import policy for IPv4 family resolution tree.

**inet-resolution-ribs [ *routing-table-names***—(Optional) Specify routing tables to use for IPv4 family protocol-next-hop resolution.

**inet6-import [ *policy-names* ]**—(Optional) Import policy for IPv6 family resolution tree.

**inet6-resolution-ribs [ *routing-table-names* ]**—(Optional) Specify routing tables to use for IPv6 family protocol-next-hop resolution.

Enabling the **inet6-resolution-ribs** option causes the static LSP route resolution to happen over the more preferred resolving route (lowest protocol preference) among the longest-matching-prefix routes in both the inet6.0 and inet6.3 routing tables.

**iso-import [ *policy-names* ]**—(Optional) Import policy for ISO family resolution tree.

**iso-resolution-ribs [ *routing-table-names* ]**—(Optional) Specify routing tables to use for ISO family protocol-next-hop resolution.

**resolution-family *resolution-family***—(Optional) Specify a family of resolution tree.

**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*

*Understanding Multitopology Routing in Conjunction with PIM*

*Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 to Provide Redundancy for Multicast Traffic over Separate Network Paths*

*Understanding Multitopology Routing in Conjunction with PIM*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.

> **NOTE:** You cannot configure both **resolve** and **retain** options for a statically configured route because resolved next hops cannot be retained.

### Default

Static routes can point only to a directly connected next hop.

> **TIP:** We recommend configuring the **no-resolve** option for individual routes to override default configuration.

**Required Privilege Level**

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

static | 365

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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 Graceful Restart for MPLS-Related Protocols*

*Configuring VPN Graceful Restart*

*Configuring Graceful Restart for VPNs*

*Configuring Logical System Graceful Restart*

# 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

*High Availability User Guide*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## Description

Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally.

> **NOTE:** You cannot configure both **retain** and resolve options for a statically configured route because resolved next hops cannot be retained.

## Default

Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally.

The retention policy applied to defaults and route is as follows:

- defaults— The default route retention policy. By default, the default route retention policy is no-retain for the configured scope using the **defaults** configuration statement. This default retention policy is overridden on a per route basis using the routes option.

- route—The retention policy for this route. The default retention policy for a route is the policy of 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

static | 365

# 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>;
        }
    }
    policy-multipath;
    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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
**policy-multipath** statement introduced in Junos Release 19.1R1 for all platfroms.

#### 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. See CLI Explorer.

#### 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 | 48

passive | 315

*policy-multipath*

*Policy-Based Multipath Routes Overview*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## Description

Specify a routing table name for route resolution.

The remaining statements are explained separately. See CLI Explorer.

## 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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*

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table | 51

interface-routes | 283

rib-groups | 350

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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. See CLI Explorer.

**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 | 51
rib-group | 349

# 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. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

# 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.

# route-distinguisher-id

### Syntax

```
route-distinguisher-id ip-address;
```

### 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.

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX 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.

> **NOTE:** Starting with Junos OS Release 15.1, when you commit a minor configuration change, the routing protocol process sends only AS paths that are active routes to the FPCs. Not all known AS paths are sent to the FPC, thereby considerably reducing the memory and CPU usage, resulting in a faster route record database update.

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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.

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is **0.0.0.0**, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.

**NOTE:** If you run OSPF for IPv6 or BGP for IPv6 in a routing instance, you must configure an IPv4 router identifier (**router-id**) in the routing instance itself. In other words, the IPv4 **router-id** in the main routing instance is not inherited by other routing instances. Even if you run *only* IPv6 OSPF or BGP in a routing instance, the IPv4 **router-id** must be configured because OSPF and BGP, even when used exclusively with IPv6, use the IPv4 **router-id** for handshaking. If you do not configure the IPv4 **router-id** in the IPv6 OSPF or BGP routing instance, then the IPv6 protocols will use invalid IPv4 address **0.0.0.0** and the adjacencies and connections will fail.

#### 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 { ... }
```

For information on the complete list of **routing-options**, see the *Protocol-Independent Routing Properties User Guide* .

### Hierarchy Level

[edit],
[edit logical-systems *logical-system-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit tenants *tenant-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.
The [edit tenants *tenant-name* routing-instances *routing-instance-name*] hierarchy level introduced in Junos OS Release 18.3R1.

### 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 User Guide*

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

## Description

Configure multicast scoping.

## Options

***scope-name***—Name of the multicast scope.

The remaining statements are explained separately. See CLI Explorer.

## Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

*Configuring Multicast Snooping*

## scope-policy

### Syntax

```
scope-policy [ policy-names ];
```

### Hierarchy Level

```
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-options multicast]
```

> **NOTE:** You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the **scope** statement to a specific routing instance at the **[edit routing-instances *routing-instance-name* routing-options multicast]** or **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]** hierarchy level.

### 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 policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

### Options

***policy-names***—Name of one or more multicast scope policies.

### Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

scope | 359

# 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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 *metric* type *number***;

The remaining statements are explained separately. See CLI Explorer.

**Required Privilege Level**
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Basic Static Routing | 6
Example: Configuring a Basic Set of Static Routes for Connecting to Stub Networks | 7
Example: Configuring IPv6 Static Routes | 13

# 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.

# tag (Routing Options)

### Syntax

```
tag metric type number;
```

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Associate a tag with a static, aggregate, or generated route.

### Default

No tag strings are associated with routes.

### Options

***metric***—Tag metric.

**Range:** 0 through 4,294,967,295

**type *number***—Tag type.

**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 | 76
aggregate | 211
generate | 261
static | 365

# threshold (Multicast Forwarding Cache)

### Syntax

```
threshold {
    log-warning value;
    suppress value;
    reuse value;
    mvpn-rpt-suppress value;
    mvpn-rpt-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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Configure the 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.

When general forwarding-cache suppression is active, the multicast forwarding-cache prevents forwarding traffic on the shared RP tree (RPT). At the same time, MVPN (*,G) forwarding states are not created for new RPT c-mcast entires, and , (*,G) installed by BGP-MVPN protocol are deleted. When general forwarding-cache suppression ends, BGP-MVPN (*,G) entries are re-added in the RIB and restored to the FIB (up to the MVPN (*,G) limit).

When MVPN RPT suppression is active, for all PE routers in excess of the threshold (including RP PEs), MVPN will not add new (*,G) forwarding entries to the forwarding-cache. Changes are visible once the entries in the current forwarding-cache have timed out or are deleted.

To use **mvpn-rpt-suppress** and/or **mvpn-rpt-reuse**, you must first configure the general **suppress** threshold. If **suppress** is configured but **mvpn-rpt-suppress** is not, both **mvpn-rpt-suppress** and **mvpn-rpt-reuse** will inherit *and use* the value set for the general **suppress**.

### Options

**reuse** or **mvpn-rpt-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** or **mvpn-rpt-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

### 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*

*show multicast forwarding-cache statistics*

# traceoptions

### 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.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### 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.

**NOTE:** On Junos OS Evolved, **traceoptions** is disabled for op, event, and commit scripts. Instead, Junos OS Evolved enables default tracking and trace messages that are logged under **/var/log/traces**.

### 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**—Nonstop active routing synchronization
- **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:** **xk** to specify KB, **xm** to specify MB, or **xg** 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 | 196

# 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.
Statement introduced in Junos OS Release 11.3 for QFX Series switches.

> **NOTE:** This feature is not supported on the EX4300 switch, even though it is available on the device.

### 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 RPF (On a Router)*

*Enabling Unicast Reverse-Path Forwarding Check for VPNs*

CHAPTER 13

# Operational Commands

IN THIS CHAPTER

# clear bfd adaptation

### Syntax

```
clear bfd adaptation
<all>
<address session-address>
<discriminator discr-number>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### Description

Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this 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.

### Options

**all**—Clear adaptation for all BFD sessions.

**address *session-address***—(Optional) Clear adaptation for all BFD sessions matching the specified address.

**discriminator *discr-number***—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.

### Additional Information

For more information, see the description of the **bfd-liveness-detection** configuration statement in the *Junos Routing Protocols Configuration Guide*.

### Required Privilege Level

clear

RELATED DOCUMENTATION

show bfd session | 389

### List of Sample Output

clear bfd adaptation on page 386

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear bfd adaptation**

`user@host>` **clear bfd adaptation**

# clear bfd session

### List of Syntax

Syntax on page 387
Syntax (EX Series Switch and QFX Series) on page 387

### Syntax

```
clear bfd session
<all>
<address session-address>
<discriminator discr-number>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switch and QFX Series)

```
clear bfd session
<all>
<address session-address>
<discriminator discr-number>
```

### Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

### Description

Drop one or more Bidirectional Forwarding Detection (BFD) sessions.

### Options

**all**—Drop all BFD sessions.

**address *session-address***—(Optional) Drop all BFD sessions matching the specified address.

**discriminator *discr-number***—(Optional) Drop the local BFD session matching the specified discriminator.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

### Required Privilege Level

clear

RELATED DOCUMENTATION

show bfd session | 389

**List of Sample Output**

clear bfd session all on page 388

**Output Fields**

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear bfd session all**

`user@host>` **clear bfd session all**

# show bfd session

### List of Syntax

Syntax on page 389
Syntax (EX Series Switch and QFX Series) on page 389

### Syntax

```
show bfd session
<brief | detail | extensive | summary>
<address address>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam (brief | detail | extensive | instance instance-name
   | summary)>
<discriminator discriminator>
<logical-system (all | logical-system-name)>
<prefix address>
<subscriber (address destination-address | discriminator discriminator | extensive)>
```

### Syntax (EX Series Switch and QFX Series)

```
show bfd session
<brief | detail | extensive | summary>
<address address>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam (brief | detail | extensive | instance instance-name
   | summary)>
<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.
Option **client** introduced in Junos OS Release 12.3R3.
Option **subscriber** introduced in Junos OS Release 15.1 for the MX 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.

**client rsvp-oam**
**(brief | detail | extensive | summary)**
**| vpls-oam**
**(brief | detail | extensive | instance** ***instance-name*** **| summary)**—(Optional) Display information about RSVP-OAM or VPLS-OAM BFD sessions in the specified level of output. For VPLS-OAM, display the specified level of output or display information about all of the BFD sessions for the specified VPLS routing instance.

**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.

**<subscriber (address** ***destination-address*** **| discriminator** ***discriminator*** **| extensive)>**—(Optional) Display information about all BFD sessions for subscribers, or for a single BFD subscriber session with a particular destination address, or with a particular denominator.

**Required Privilege Level**
view

RELATED DOCUMENTATION

clear bfd session | 387
Understanding BFD for Static Routes for Faster Network Failure Detection | 85
Example: Configuring BFD for Static Routes for Faster Network Failure Detection | 90
*Understanding BFD for OSPF*
*Example: Configuring BFD for OSPF*
*Understanding BFD for BGP*
*Example: Configuring BFD on Internal BGP Peer Sessions*
*Understanding Bidirectional Forwarding Detection Authentication for PIM*
*Configuring BFD for PIM*
*Understanding BFD for IS-IS*

**List of Sample Output**
show bfd session on page 396
show bfd session brief on page 396
show bfd session detail on page 396
show bfd session detail (with Authentication) on page 396

### Output Fields

Table 14 on page 391 describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

**Table 14: 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 or process for which the BFD session is active: **ISIS**, **OSPF**, **DHCP**, **Static**, or **VGD**. | **detail extensive** |
| TX interval | Time interval, in seconds, used by the host system to transmit BFD control packets. | **brief detail extensive** none |

Table 14: show bfd session Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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.<br><br>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** |
| 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**.<br><br>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.<br><br>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.<br><br>Following are the expected values for Local Diagnostic output field:<br><br>• **None**—No diagnostic<br>• **CtlExpire**—Control detection time expired<br>• **EchoExpire**—Echo detection time expired<br>• **NbrSignal**—Neighbor signalled session down<br>• **FwdPlaneReset**—Forwarding plane reset<br>• **PathDown**—Path down<br>• **ConcatPathDown**—Concatenated path down<br>• **AdminDown**—Administratively down | **detail extensive** |

Table 14: show bfd session Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Remote diagnostic | Remote diagnostic information about failing BFD sessions.<br>Following are the expected values for Remote Diagnostic output field:<br>• **None**—No diagnostic<br>• **CtlExpire**—Control detection time expired<br>• **EchoExpire**—Echo detection time expired<br>• **NbrSignal**—Neighbor signalled session down<br>• **FwdPlaneReset**—Forwarding plane reset<br>• **PathDown**—Path down<br>• **ConcatPathDown**—Concatenated path down<br>• **AdminDown**—Administratively down | **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** |

Table 14: show bfd session Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Remote min TX interval | Minimum amount of time, in seconds, between control packet detections on the remote system. | **extensive** |
| 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.<br><br>The following BFD sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces. | **extensive** |

Table 14: show bfd session Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Authentication | Summary status of BFD authentication:<br>• **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.<br>• **keychain**—Name of the security authentication keychain associated with the specified BFD session.<br>• **algo**—BFD authentication algorithm being used: **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**.<br>• **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.<br>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 |
| 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.<br>The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address. | **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.

### 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 00:34:02
  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 00:29:04, 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 00:34:18
  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 00:29:12, 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 client rsvp-oam**

`user@host>` **show bfd session client rsvp-oam**

```
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.0.223            Up                       540.000   180.000      3

1 Up sessions, 0 Down sessions
1 sessions, 1 clients
Cumulative transmit rate 0.0 pps, cumulative receive rate 0.0 pps
```

**show bfd session client vpls-oam summary**

`user@host>` **show bfd session client vpls-oam summary**

```
1 Up sessions, 1 Down sessions
 2 sessions, 2 clients
 Cumulative transmit rate 2.0 pps, cumulative receive rate 1.0 pps
```

**show bfd session client vpls-oam instance instance-name**

`user@host>` **show bfd session client vpls-oam instance vpls**

```
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
127.0.0.1                Up        ae9.0          3.000     1.000        3

1 Up Sessions, 0 Down Sessions
1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

**show bfd session extensive**

`user@host>` **show bfd session extensive**

```
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
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
  Micro-BFD Session
                                                  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 bfd session subscriber**

user@host> **show bfd session subscriber**

```
      Detect  Transmit
Address   State  Interface  Time  Interval  Multiplier
1.0.0.2   Up   ae0.0  90.000  30.000   3
1.0.0.6   Up  ae0.1  90.000  30.000  3
1.0.0.10 Up  ae0.2  90.000 30.000   3
1.0.0.14 Up  ae0.3  90.000  30.000   3
1.0.0.18 Up   ae0.4  90.000  30.000   3

20 sessions, 20 clients
```

**show bfd session subscriber address**

user@host> **show bfd session subscriber address 1.0.0.2**

```
     Detect  Transmit
Address  State  Interface Time  Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000  30.000   3
```

```
1 sessions, 1 clients
Cumulative transmit rate 5.0 pps, cumulative receive rate 5.0 pps
```

**show bfd session subscriber extensive**

`user@host>` **show bfd session subscriber extensive**

```
     Detect   Transmit
Address  State  Interface Time   Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 20, remote discriminator 16
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x1

     Detect   Transmit
Address  State  Interface  Time   Interval  Multiplier
1.0.0.6  Up   ae0.1  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 21, remote discriminator 17
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x2
```

**show bfd session subscriber discriminator extensive**

user@host> **show bfd session subscriber discriminator 20 extensive**

```
     Detect   Transmit
Address  State  Interface  Time   Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 20, remote discriminator 16
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x1

1 sessions, 1 clients
Cumulative transmit rate 5.0 pps, cumulative receive rate 5.0 pps
```

# show as-path

### List of Syntax

Syntax on page 403
Syntax (EX Series Switches) on page 403

### Syntax

```
show as-path
<brief | detail>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show as-path
<brief | detail>
```

### 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.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

### Description

Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the **as-path-prepend** action) or through aggregation.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

### Options

**none**—Display basic information about AS paths that the local routing device is using (same as brief).

**brief | detail**—(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.

### Required Privilege Level

view

RELATED DOCUMENTATION

show as-path summary | 411

**List of Sample Output**



**Output Fields**

Table 15 on page 404 lists the output fields for the **show as-path** command. Output fields are listed in the approximate order in which they appear.

**Table 15: show as-path Output Fields**

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Total AS paths** | Total number of AS paths. | **brief** none |
| **Bucket** | Bucket number. | All levels |
| **Count** | Number of AS path entries in this bucket. | All levels |
| ***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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **?**—Incomplete; typically, the AS path was aggregated.<br>• **Atomic**—Route is an aggregate of several route prefixes.<br>• **Aggregator**—Routing device has summarized a range of prefixes. | All levels |
| **domain** | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain. | **detail** |
| **neighbor as** | AS peer address. | **detail** |
| **length** | Length of the AS path. | **detail** |
| **segments** | Length of the AS segment descriptor. | **detail** |
| **unique-count** | Number of unique autonomous systems (ASs) present in the AS path | **detail** |
| **references** | Path reference count. | **detail** |

## Sample Output

### show as-path

user@host> **show as-path**

```
Total AS paths: 30382
  Bucket 0    Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I
    14203 2914 3491 1680 33802 I
    14203 2914 3549 7908 27658 I
    14203 2914 3549 20804 30730 I
    14203 2914 7018 2687 9226 I
    14203 2914 174 9318 9318 23564 I
    14203 2914 701 3786 3786 23564 I
    14203 2914 701 4761 4795 9228 I
    14203 2914 1239 7132 5673 18444 I
    14203 2914 3491 20485 24588 24588 I
    14203 2914 5511 2200 1945 2060 I
    14203 2914 7911 14325 14325 14348 I
    14203 2914 701 4637 9230 9230 9230 I
    14203 2914 6395 14 14 14 14 I
    14203 2914 9299 6163 6163 6163 6163 9232 I
    14203 2914 3356 3356 3356 3356 3356 11955 21522 I
    14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
    14203 2914 174 30209 30222 30222 30222 ?
    14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
    14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
    14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
  Bucket 1    Count: 28
```

```
    14203 2914 35847 I
    14203 2914 174 19465 I
    14203 2914 174 35849 I
    14203 2914 2828 32777 I
    14203 2914 4323 14345 I
    14203 2914 4323 29705 I
    14203 2914 6395 32777 I

...
```

**show as-path detail**

`user@host>` **show as-path detail**

```
Total AS paths: 30410
  Bucket 0    Count: 36
    AS path: I
      domain 0, length 0, segments 0, unique-count 0, references 54
    AS path: 14203 2914 174 31752 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 701 21512 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 1239 26632 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5, references
 2
    AS path: 14203 2914 1239 29704 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 4323 10248 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6, references
 2
    AS path: 14203 2914 4766 23560 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 6395 32776 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5, references
 3
    AS path: 14203 2914 7911 11272 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6, references
 2
    AS path: 14203 2914 12180 18440 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 3, references
```

```
 3
    AS path: 14203 2914 17408 17416 I
     domain 1, neighbor as: 14203, length 4, segments 1, unique-count 8, references
 3
    AS path: 14203 2914 701 702 24586 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references
 3
    AS path: 14203 2914 1239 4657 9226 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 5, references
 7
    AS path: 14203 2914 1239 7132 16394 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 7, references
 2
    AS path: 14203 2914 1299 8308 34826 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 8, references
 2
    AS path: 14203 2914 3320 5603 28682 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 3491 1680 33802 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 14,
references 2
    AS path: 14203 2914 3549 7908 27658 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 6, references
 2
    AS path: 14203 2914 3549 20804 30730 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 24,
references 2
    AS path: 14203 2914 7018 2687 9226 I
     domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4, references
 3
    AS path: 14203 2914 174 9318 9318 23564 I
     domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 701 3786 3786 23564 I
     domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 701 4761 4795 9228 I
     domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 14
    AS path: 14203 2914 1239 7132 5673 18444 I
     domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 3491 20485 24588 24588 I
```

```
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 4
    AS path: 14203 2914 5511 2200 1945 2060 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 7911 14325 14325 14348 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4, references
 2
    AS path: 14203 2914 701 4637 9230 9230 9230 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4, references
 3
    AS path: 14203 2914 6395 14 14 14 14 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4, references
 10
...
```

# show as-path domain

### List of Syntax

Syntax on page 409
Syntax (EX Series Switches) on page 409

### Syntax

```
show as-path domain
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show as-path domain
```

### Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

### Description

Display autonomous system (AS) path domain information.

### Options

**none**—(Optional) Display AS path domain information for all routing instances.

**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 as-path domain on page 410**

### Output Fields

Table 16 on page 410 lists the output fields for the **show as-path domain** command. Output fields are listed in the approximate order in which they appear

Table 16: show as-path domain Output Fields

| Field Name | Field Description |
|---|---|
| **Domain** | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain. |
| **Primary** | Primary AS number. |
| **References** | Path reference count. |
| ***Number* Paths** | Number of known AS paths. |
| **Flags** | Information about the AS path:<br>• **ASLoop**—Path contains an AS loop.<br>• **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute.<br>• **Local**—Path was created by local aggregation.<br>• **Master**—Path was created by the master routing instance. |
| **Local AS** | AS number of the local routing device. |
| **Loops** | How many times this AS number can appear in an AS path. |

## Sample Output

### show as-path domain

`user@host>` **show as-path domain**

```
Domain: 1         Primary: 10458
    References:           3 Paths:      30383
    Flags: Master
    Local AS: 10458  Loops: 1
```

# show as-path summary

**List of Syntax**

Syntax on page 411
Syntax (EX Series Switches) on page 411

**Syntax**

```
show as-path summary
<logical-system (all | logical-system-name)>
```

**Syntax (EX Series Switches)**

```
show as-path 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 autonomous system (AS) path summary information.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

**Options**

**none**—(Optional) Display AS path summary information for all routing instances.

**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

show as-path | 403

**List of Sample Output**

**show as-path summary on page 412**

### Output Fields

Table 17 on page 412 lists the output fields for the **show as-path summary** command. Output fields are listed in the approximate order in which they appear.

**Table 17: show as-path summary Output Fields**

| Field Name | Field Description |
|---|---|
| **AS Paths** | Number of AS paths. |
| **Buckets** | Number of hash buckets in use. |
| **Max** | Maximum number of AS path entries per bucket. |
| **Min** | Minimum number of AS path entries per bucket. |
| **Avg** | Average number of AS path entries per bucket. |
| **Std deviation** | Standard deviation of AS path entries per bucket. |

## Sample Output

**show as-path summary**

`user@host>` **show as-path summary**

```
AS Paths  Buckets     Max    Min    Avg    Std deviation
   30425     1024      95     12     29         6.481419
```

# show chassis forwarding-options

### Syntax

```
show chassis forwarding-options
```

### Release Information

Command introduced in Junos OS Release 13.2
Support added to QFX5200 switches in Junos OS Release 15.1X53-D30

### Description

Display the configuration for the Unified Forwarding Table.

### Options

There are no options for this command.

> **NOTE:** Starting in Junos OS Releases 17.3R2, for QFX5200 Virtual Chassis, information about memory banks are displayed only for the Master, not for the other members. Values remain the same across all members. All configuration changes for the Unified Forwarding Table are made through the Master.

### Required Privilege Level

view

### RELATED DOCUMENTATION

*Configuring the Unified Forwarding Table on Switches*

*Example: Configuring a Unified Forwarding Table Custom Profile*

### List of Sample Output

**show chassis forwarding-options (l2-profile-three) on page 414**
**show chassis forwarding-options (custom-profile on QFX5200 Series switch) on page 415**
**show chassis forwarding-options (QFX5200 Virtual Chassis) on page 415**

### Output Fields

Table 18 on page 414 lists the output fields for the **show chassis forwarding-options** command. Output fields are listed in the approximate order in which they appear.

Table 18: show chassis forwarding-options Output Fields

| Field Name | Field Description |
|---|---|
| profile name | Name of profile configured:<br>• **custom-profile** (QFX5200 only)<br>• **l2-profile-one**<br>• **l2-profile-three** (default)<br>• **l2-profile-two**<br>• **l3-profile**<br>• **lpm-profile** |
| MAC | Maximum amount of memory allocated for Layer 2 entries. |
| L3-host | Maximum amount of memory allocated for Layer 3 host entries. |
| LPM | Maximum amount of memory allocated for longest match prefix (LPM) entries. |
| num-65-127-prefix | Maximum amount of memory allocated in LPM table for IP prefixes with lengths in the range /65 through /127. |
| Total scale(K) | (QFX5200 only) Maximum amount of memory allocated for each address type. This amount includes the amount configured plus the amount allocated through the dedicated hash table. |
| Bank details for various types of entries | (QFX5200 only) Maximum amount of memory configured by address type for each of the four shared memory banks and the dedicated hash table. |
| Entry type | (QFX5200 only) Type of forwarding-table entry: **L2(mac)**; **L3 (unicast and multicast)**; **Exact Match**; and **Longest Prefix Match (lpm)** |
| Dedicated bank size(K) | (QFX5200 only) Maximum amount of memory allocated for each address type in the dedicated hash table. |
| Shared bank size(K) | (QFX5200 only) Default Maximum amount of memory allocated for each address type in the shared memory banks. |

## Sample Output

**show chassis forwarding-options (l2-profile-three)**

`user@host>` **show chassis forwarding-options**

```
UFT Configuration:
l2-profile-three. (MAC: 160K L3-host: 144K LPM: 16K) (default)
num-65-127-prefix = none

{master:0}
```

**show chassis forwarding-options (custom-profile on QFX5200 Series switch)**

`user@host>` **show chasiss forwarding-options**

```
UFT Configuration:
custom-profile
Configured custom scale:
Entry type                 Total scale(K)
L2(mac)                         8
L3 (unicast & multicast)        72
Exact Match                     0
Longest Prefix Match (lpm)      80
num-65-127-prefix = 1K
-------------Bank details for various types of entries------------
Entry type                   Dedicated Bank Size(K)      Shared Bank Size(K)
L2 (mac)                     8                         32 * num shared banks
L3 (unicast & multicast      8                         32 * num shared banks
Exact match                  0                         16 * num shared banks
Longest Prefix match(lpm)    16                        32 * num shared banks
```

**show chassis forwarding-options (QFX5200 Virtual Chassis)**

`user@host>` **show chassis forwarding-options**

```
localre:
-
UFT Configuration:
l2-profile-three.(default)
num-65-127-prefix = 1K
-Bank details for various types of entries-
Entry type                         Dedicated Bank Size(K)     Shared Bank Size(K)
L2(mac)                           8                         32 * num shared banks
L3(unicast & multicast)            8                        32 * num shared banks
Exact Match                       0                         16 * num shared banks
Longest Prefix Match(lpm)          16                       32 * num shared banks

fpc1:
```

```
-
UFT Configurtion:
l2-profile-three.(default)
num-65-127-prefix = 1K
```

# show interfaces routing summary

### Syntax

```
show interfaces routing summary
<interface-name>
<logical-system (all | logical-system-name)>
```

### Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

### Description

Display a summary of the state of the router interfaces. Use this command for performing router diagnostics only, when you are determining whether the routing protocols and the Junos OS differ about the state of an interface.

### Options

**none**—Display summary information about the state of all router interfaces on all logical systems.

***interface-name***—(Optional) Name of a specific interface.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

### Additional Information

For information about how to configure routing protocols, see the *Junos OS Routing Protocols Library*. For information about related operational mode commands for routing instances and protocols, see the CLI Explorer.

### Required Privilege Level

view

### List of Sample Output

**show interfaces routing summary on page 418**
**show interfaces routing summary (TX Matrix Plus Router) on page 419**
**show interfaces routing summary (PTX5000 Packet Transport Routers) on page 419**

### Output Fields

Table 19 on page 418 lists the output fields for the **show interfaces routing summary** command. Output fields are listed in the approximate order in which they appear.

Table 19: show interfaces routing summary Output Fields

| Field Name | Field Description |
|---|---|
| ***n* physical interfaces** | Number of routing interfaces and number of interfaces in the **up** state. |
| ***n protocol* protocol interfaces** | Type and number of routing protocols and number of related interfaces in the **up** state. |
| **Interface** | Logical interface name. |
| **Index** | Logical interface index number, which reflects its initialization sequence. |
| **Metric** | Metric value for the interface. |
| **Trans** | Number of times the interface has transitioned from **Down** to **Up**. |
| **Status** | Interface status (**Up** or **Down**) and type. |

## Sample Output

### show interfaces routing summary

user@host> **show interfaces routing summary**

```
14 physical interfaces (12 up)
  11 INET protocol addresses (11 up)
  6 ISO protocol addresses (4 up)
  3 MPLS protocol addresses (3 up)
  3 CCC protocol addresses (3 up)
Interface  Index   Metric   Trans. Status
so-5/0/3.0    15        0        0 Broadcast PointToPoint Multicast
so-5/0/2.0    14        0        0 Up Broadcast PointToPoint Multicast
so-5/0/1.0    13        0        5 Up Broadcast PointToPoint Multicast
so-5/0/0.0    12        0        2 Up Broadcast PointToPoint Multicast
so-1/2/0.0    11        0        0 Broadcast PointToPoint Multicast
so-1/1/0.0    10        0        5 Up Broadcast PointToPoint Multicast
at-1/0/0.6     9        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.5     8        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.4     7        0        0 Up Broadcast PointToPoint Multicast
```

```
at-1/0/0.3     6        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.2     5        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.0     4        0        0 Up Broadcast PointToPoint Multicast
lo0.0          3        0        0 Up Broadcast Loopback Multicast
fxp1.0         2        0        1 Up Broadcast Multicast
fxp0.0         1        0        0 Up Broadcast Multicast
```

**show interfaces routing summary (TX Matrix Plus Router)**

`user@host>` **show interfaces routing summary**

```
9 physical interfaces (9 up)
  11 INET protocol addresses (11 up)
  6 MPLS protocol addresses (6 up)
  4 INET6 protocol addresses (4 up)

Interface                Index   Metric   Trans. Status
ge-23/0/8.0                 73        0        0 Up Broadcast Multicast
ge-23/0/7.0                 72        0        0 Up Broadcast Multicast
ge-23/0/6.0                 71        0        0 Up Broadcast Multicast
ge-7/0/9.0                  69        0        0 Up Broadcast Multicast
ge-15/0/9.0                 70        0        0 Up Broadcast Multicast
xe-6/1/1.0                  68        0        0 Up Broadcast Multicast
lo0.16385                   66        0        0 Up Broadcast Loopback Multicast
lo0.16384                   65        0        0 Up Broadcast Loopback Multicast
lo0.0                       64        0        0 Up Broadcast Loopback Multicast
ixgbe1.0                     5        0        0 Up Broadcast Multicast
ixgbe0.0                     4        0        0 Up Broadcast Multicast
em0.0                        3        0        0 Up Broadcast Multicast
```

**show interfaces routing summary (PTX5000 Packet Transport Routers)**

`user@host>` **show interfaces routing summary**

```
7 physical interfaces (68 up)
  7 INET protocol addresses (7 up)
  2 CCC protocol addresses (2 up)
  4 INET6 protocol addresses (4 up)

Interface                Index   Metric   Trans. Status
lo0.16385                   66        0        0 Up Broadcast Loopback Multicast
lo0.16384                   64        0        0 Up Broadcast Loopback Multicast
lo0.0                       65        0        0 Up Broadcast Loopback Multicast
```

```
ixgbe1.0                     5         0         0 Up Broadcast Multicast
ixgbe0.0                     4         0         0 Up Broadcast Multicast
et-5/0/5.32767              72         0         0 Up Broadcast Multicast
et-5/0/5.0                  68         0         0 Up Broadcast Multicast
et-5/0/0.32767              67         0         0 Up Broadcast Multicast
et-5/0/0.0                  71         0         0 Up Broadcast Multicast
em0.0                        3         0         0 Up Broadcast Multicast
```

# show route

### List of Syntax



### Syntax

```
show route
<all>
<destination-prefix>
<logical-system (all | logical-system-name)>
<private>
<te-ipv4-prefix-ip  te-ipv4-prefix-ip>
<te-ipv4-prefix-node-ip te-ipv4-prefix-node-ip>
<te-ipv4-prefix-node-iso te-ipv4-prefix-node-iso>
<rib-sharding (main | rib-shard-name)>
```

### 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.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.
Options **te-ipv4-prefix-ip**, **te-ipv4-prefix-node-ip**, and **te-ipv4-prefix-node-iso** introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.
**rib-sharding** option introduced in cRPD Release 20.1R1.

### 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.

**display-client-data** —(Optional) Display client id and cookie information for routes installed by the routing protocol process client applications.

**te-ipv4-prefix-ip *te-ipv4-prefix-ip***—(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.

**te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the **lsdist.0** table.

**te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the **lsdist.0** table.

**rib-sharding (main | *rib-shard-name*)**—(Optional) Display the rib shard name.

**Required Privilege Level**

view

## RELATED DOCUMENTATION

*Understanding IS-IS Configuration*

*Example: Configuring IS-IS*

*Examples: Configuring Internal BGP Peering*

*Examples: Configuring External BGP Peering*

*Examples: Configuring OSPF Routing Policy*

*Verifying and Managing Junos OS Enhanced Subscriber Management*

**List of Sample Output**

show route on page 426
show route (VPN) on page 427
show route (with Destination Prefix) on page 428
show route destination-prefix detail on page 428
show route extensive on page 428
show route extensive ( ECMP) on page 429

show route extensive (Multipath Resolution) on page 429
show route active-path extensive on page 430
show route (Enhanced Subscriber Management) on page 435
show route (IPv6 Flow Specification) on page 436
show route display-client-data detail on page 436
show route rib-sharding on page 437
show route te-ipv4-prefix-ip on page 438
show route te-ipv4-prefix-ip extensive on page 439
show route te-ipv4-prefix-node-iso on page 442
show route te-ipv4-prefix-node-iso extensive on page 443
show route te-ipv4-prefix-node-iso detail on page 446
show route rib-sharding junos-bgpshard14 on page 450

### Output Fields

Table 20 on page 423 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

**Table 20: 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:<br>• **active** (routes that are active).<br>• **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.<br>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.<br>If you have configured **uRPF-loose** mode, the holddown bit is most likely set because Kernel Routing Table (KRT) is using inactive route to build valid incoming interfaces. In this case, you can ignore the holddown state because nothing is wrong.<br>• **hidden** (routes that are not used because of a routing policy). |

Table 20: show route Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| ***destination-prefix*** | Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:<br>• ***MPLS-label*** (for example, 80001).<br>• ***interface-name*** (for example, ge-1/0/2).<br>• ***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):<br>  ◦ ***neighbor-address***—Address of the neighbor.<br>  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.<br>  ◦ ***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.<br>  ◦ ***vc-id***—Virtual circuit identifier.<br>  ◦ ***source***—Source of the advertisement: **Local** or **Remote**. |
| **[ *protocol*, *preference* ]** | Protocol from which the route was learned and the preference value for the route.<br>• **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br>• **-** —A hyphen indicates the last active route.<br>• *****—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.<br>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. |
| ***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. |

Table 20: show route Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **?**—Incomplete; typically, the AS path was aggregated.<br>When AS path numbers are included in the route, the format is as follows:<br>• **[ ]**—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.<br>• **{ }**—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.<br>• **( )**—Parentheses enclose a confederation.<br>• **( [ ] )**—Parentheses and brackets enclose a confederation set.<br>**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. |
| **encapsulated** | Extended next-hop encoding capability enabled for the specified BGP community for routing IPv4 traffic over IPv6 tunnels. When BGP receives routes without the tunnel community, IPv4-0ver IPv6 tunnels are not created and BGP routes are resolved without encapsulation. |
| **Route Labels** | Stack of labels carried in the BGP route update. |
| **validation-state** | (BGP-learned routes) Validation status of the route:<br>• **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.<br>• **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br>• **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.<br>• **Valid**—Indicates that the prefix and autonomous system pair are found in the database. |

Table 20: show route Output Fields *(continued)*

| Field Name | Field Description |
| --- | --- |
| **to** | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.<br><br>If the destination is **Discard**, traffic is dropped. |
| **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:<br><br>• **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.<br>• **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.<br>• ***lsp-path-name***—Name of the LSP used to reach the next hop.<br>• ***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). |
| **Private unicast** | (Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information. |
| balance | Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages. |

## 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 (VPN)

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 192.0.2.0**

```

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0 
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > 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 ?, validation-state: unverified
                    > 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 (with Destination Prefix)

user@host> **show route 192.168.0.0/12**

```

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

192.168.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

### show route destination-prefix detail

user@host> **show route 198.51.100.0 detail**

```

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
         Next hop type: Discard
         Address: 0x9041ae4
         Next-hop reference count: 2
         State: <NoReadvrt Int Ext AlwaysFlash>
         Inactive reason: Route Preference
         Local AS:   200
         Age: 4d 1:40:40
         Validation State: unverified
         Task: RT
         Announcement bits (1): 2-BGP_RT_Background
         AS path: 4 5 6 I
```

### 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 203.0.113.1

           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: 64510 Peer AS: 64511
            Age: 3  Metric2: 1
            Validation State: unverified
            Task: BGP_64510.10.0.0.30+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:64502:100 encapsulation:0L:14
            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 extensive ( ECMP)**

`user@host>` **show route extensive**

```
*IS-IS Preference: 15
                 Level: 1
                 Next hop type: Router, Next hop index: 1048577
                 Address: 0xXXXXXXXXXX
                 Next-hop reference count: YY
                 Next hop: 198.51.100.2 via ae1.0 balance 43%, selected
                 Session Id: 0x141
                 Next hop: 192.0.2.2 via ae0.0 balance 57%
```

**show route extensive (Multipath Resolution)**

`user@host>` **show route extensive**

```
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)
10.1.1.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.1.2/32 -> {indirect(1048574)}
        *Static Preference: 5
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d1b0
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 581
                Next hop: 10.1.1.2 via ge-2/0/1.0, selected
                Session Id: 0x144
                Next hop: 10.2.1.2 via ge-2/0/2.0, selected
                Session Id: 0x145
                Protocol next hop: 10.1.1.1
                Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
                State: <Active Int Ext>
                Age: 2:53  Metric2: 0
                Validation State: unverified
                Task: RT
                Announcement bits (2): 0-KRT 2-Resolve tree 1
                AS path: I
                Indirect next hops: 1
                        Protocol next hop: 10.1.1.1
                       Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143

                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 10.1.1.2 via ge-2/0/1.0
                                Session Id: 0x144
                                Next hop: 10.2.1.2 via ge-2/0/2.0
                                Session Id: 0x145
                     10.1.1.1/32 Originating RIB: inet.0
                        Node path count: 1
                          Node flags: 1
                        Forwarding nexthops: 2 (Merged)
                        Nexthop: 10.1.1.2 via ge-2/0/1.0
                         Nexthop: 10.2.1.2 via ge-2/0/2.0
```

**show route active-path extensive**

`user@host>` **show route active-path extensive**

```
user@host> show route 198.51.100.1 active-path extensive
```

```
inet.0: 1000061 destinations, 1000082 routes (1000061 active, 0 holddown, 0 hidden)
198.51.100.1/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.1/32 -> {indirect(1051215)}
unicast reverse-path: 0
[ae0.0 ae1.0]
Page 0 idx 0, (group Internet-IPv4 type External) Type 1 val 0xbb2e53d8 (adv_entry)
Advertised metrics:
Nexthop: Self
AS path: [500] 410 I
Communities:
Path 198.51.100.1 from 10.0.0.11 Vector len 4. Val: 0
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0x2e9aacdc
Next-hop reference count: 500000
Source: 10.0.0.11
Next hop type: Router, Next hop index: 0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 20(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 20: None;
Label element ptr: 0xb5dc1780
Label parent element ptr: 0x18d48080
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 22(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 22: None;
Label element ptr: 0xb5dc1700
Label parent element ptr: 0x18d41000
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 48(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 48: None;
Label element ptr: 0x18d40800
Label parent element ptr: 0x18d49780
```

```
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 49(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 49: None;
Label element ptr: 0xb5dc1680
Label parent element ptr: 0x18d48f00
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 21(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 21: None;
Label element ptr: 0xb5dc1600
Label parent element ptr: 0x18d44d80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 25(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 25: None;
Label element ptr: 0xb5dc1580
Label parent element ptr: 0x18d3da80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1, selected
Label operation: Push 3851, Push 24, Push 68(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 68: None;
Label element ptr: 0x18d41500
Label parent element ptr: 0x18d49000
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
```

```
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 24, Push 69(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 69: None;
Label element ptr: 0xb5dc1500
Label parent element ptr: 0x18d48300
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Protocol next hop: 10.0.0.11
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
State:
Local AS: 500 Peer AS: 500
Age: 1:40:03 Metric2: 2
Validation State: unverified
Task: BGP_500.10.0.0.11
Announcement bits (5): 0-KRT 8-KRT 9-BGP_RT_Background 10-Resolve tree 5 11-Resolve
 tree 8
AS path: 410 I
Accepted
Route Label: 3851
Localpref: 100
Router ID: 10.0.0.11
Indirect next hops: 1
Protocol next hop: 10.0.0.11 Metric: 2
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
Indirect path forwarding next hops (Merged): 8
Next hop type: Router
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
```

```
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
10.0.0.11/32 Originating RIB: inet.3
Metric: 1 Node path count: 4
Node flags: 1
Indirect nexthops: 4
Protocol Nexthop: 10.0.0.4 Metric: 1 Push 24
Indirect nexthop: 0x1880f200 1048597 INH Session ID: 0xb0c
Path forwarding nexthops link: 0x36120400
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.4/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.5 Metric: 1 Push 24
Indirect nexthop: 0x18810000 1048596 INH Session ID: 0xb0b
Path forwarding nexthops link: 0x1545be00
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.5/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.6 Metric: 1 Push 25
```

```
Indirect nexthop: 0x1880e600 1048588 INH Session ID: 0xb0a
Path forwarding nexthops link: 0x3611f440
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.6/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.7 Metric: 1 Push 25
Indirect nexthop: 0x1880dc00 1048586 INH Session ID: 0xb09
Path forwarding nexthops link: 0x15466d80
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.7/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
```

### show route (Enhanced Subscriber Management)

`user@host>` **show route**

```
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                > to #0 10.0.0.1.93.65 via demux0.1073741824
198.51.100.12/24    *[Access-internal/12] 00:00:08
                Private unicast
```

### show route (IPv6 Flow Specification)

`user@host>` **show route**

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

2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

### show route display-client-data detail

`user@host>` **show route 198.51.100.0/24 display-client-data detail**

```
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
```

```
          *BGP-Static Preference: 5/-101
                  Next hop type: Indirect, Next hop index: 0
                  Address: 0xa5c2af8
                  Next-hop reference count: 2
                  Next hop type: Router, Next hop index: 1641
                  Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                  Session Id: 0x160
                  Protocol next hop: 192.0.2.1
                  Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                  State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                  Age: 3:13       Metric2: 0
                  Validation State: unverified
                  Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
                  AS path: I
                  Client id: 1, Cookie: 1
```

#### show route rib-sharding

`user@host>` **show route rib-sharding *shard-name***

```
inet.0: 193295 destinations, 386345 routes (193295 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.0.4.0/22         *[BGP/170] 03:35:38, localpref 100, from 1.1.1.1
                    AS path: 69 10458 14203 4826 38803 56203 I, validation-state:
 unverified
                   >  to 10.205.191.254 via fxp0.0
                    [BGP/170] 03:35:38, localpref 100, from 1.1.1.2
                    AS path: 69 10458 14203 4826 38803 56203 I, validation-state:
 unverified
                   >  to 10.205.191.254 via fxp0.0
1.0.6.0/24         *[BGP/170] 03:35:38, localpref 100, from 1.1.1.1
                    AS path: 69 10458 14203 4826 38803 56203 I, validation-state:
 unverified
                   >  to 10.205.191.254 via fxp0.0
                    [BGP/170] 03:35:38, localpref 100, from 1.1.1.2
                    AS path: 69 10458 14203 4826 38803 56203 I, validation-state:
 unverified
                   >  to 10.205.191.254 via fxp0.0
1.0.64.0/18        *[BGP/170] 03:35:38, localpref 100, from 1.1.1.1
                      AS path: 69 10458 14203 2914 2497 7670 18144 I,
validation-state: unverified
                   >  to 10.205.191.254 via fxp0.0
                    [BGP/170] 03:35:38, localpref 100, from 1.1.1.2
```

```
                      AS path: 69 10458 14203 2914 2497 7670 18144 I,
validation-state: unverified
                    >  to 10.205.191.254 via fxp0.0
```

**show route te-ipv4-prefix-ip**

`user@host>` **show route te-ipv4-prefix-ip 10.10.10.10**

```
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
```

```
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                  *[IS-IS/18] 00:01:01
                     Fictitious
```

**show route te-ipv4-prefix-ip extensive**

`user@host>`**show route te-ipv4-prefix-ip 10.10.10.10 extensive**

```
lsdist.0: 298 destinations, 298 routes (298 active, 0 holddown, 0 hidden)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State:<Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
       State: <Active NotInstall
       Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0>

PREFIX { Node { AS:64496 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
```

```
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
       Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0
```

```
PREFIX { Node { AS:64496 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
```

```
                 Task: IS-IS
                 AS path: I
                 Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                 Level: 2
                 Next hop type: Fictitious, Next hop index: 0
                 Address: 0xa1a2ac4
                 Next-hop reference count: 298
                 Next hop:
                 State: <Active NotInstall>
                 Local AS:   64496
                 Age: 7:58
                 Validation State: unverified
                 Task: IS-IS
                 AS path: I
                 Prefix SID: 1000, Flags: 0x40, Algo: 0
```

**show route te-ipv4-prefix-node-iso**

user@host> **show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00**

```
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
```

```
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
```

**show route te-ipv4-prefix-node-iso extensive**

`user@host>` **show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 extensive**

```
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
```

```
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
```

```
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
```

```
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

**show route te-ipv4-prefix-node-iso detail**

`user@host>` **show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 detail**

```
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
```

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
```

```
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
```

```
                  Level: 2
                  Next hop type: Fictitious, Next hop index: 0
                  Address: 0xa1a2ac4
                  Next-hop reference count: 283
                  Next hop:
                  State: <Active NotInstall>
                  Local AS:   64496
                  Age: 6:54
                  Validation State: unverified
                  Task: IS-IS
                  AS path: I
                  Prefix SID: 1000, Flags: 0x40, Algo: 0
```

**show route rib-sharding junos-bgpshard14**

`user@host>` **show route rib-sharding junos-bgpshard14**

```
inet.0: 54 destinations, 54 routes (54 active, 0 holdown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.50.8/30        *[Direct/0] 00:29:15
                    >  via eth9.101
2.2.50.9/32        *[Local/0] 00:29:15
                       Local via eth9.101
2.2.51.8/30        *[Direct/0] 00:29:15
                    >  via eth9.102@
2.2.51.9/32        *[Local/0] 00:29:15
                       Local via eth9.102@
2.2.52.8/30        *[Direct/0] 00:29:15
                    >  via eth8.103
2.2.52.9/32        *[Local/0] 00:29:15
                       Local via eth8.103
2.2.54.8/30        *[Direct/0] 00:29:15
                    >  via eth8.104@
2.2.54.9/32        *[Local/0] 00:29:15
                       Local via eth8.104@
2.2.55.8/30        *[Direct/0] 00:29:15
                    >  via eth6.105
2.2.55.9/32        *[Local/0] 00:29:15
                       Local via eth6.105
2.2.56.8/30        *[Direct/0] 00:29:15
                    >  via eth6.106@
2.2.56.9/32        *[Local/0] 00:29:15
                       Local via eth6.106@
```

```
2.2.57.8/30        *[Direct/0] 00:29:15
                    >  via eth7.107
2.2.57.9/32        *[Local/0] 00:29:15
                       Local via eth7.107
2.2.58.8/30        *[Direct/0] 00:29:15
                    >  via eth7.108@
2.2.58.9/32        *[Local/0] 00:29:15
                       Local via eth7.108@
3.2.56.8/30        *[Direct/0] 00:23:38
                    >  via eth11.115
3.2.56.9/32        *[Local/0] 00:23:38
                       Local via eth11.115
3.2.57.8/30        *[Direct/0] 00:23:38
                    >  via eth11.116
3.2.57.9/32        *[Local/0] 00:23:38
                       Local via eth11.116
3.12.50.8/30       *[Direct/0] 00:23:38
                    >  via eth13.109
3.12.50.9/32       *[Local/0] 00:23:38
                       Local via eth13.109
3.12.51.8/30       *[Direct/0] 00:23:38
                    >  via eth13.110
3.12.51.9/32       *[Local/0] 00:23:38
                       Local via eth13.110
3.12.52.8/30       *[Direct/0] 00:23:38
                    >  via eth10.111
3.12.52.9/32       *[Local/0] 00:23:38
                       Local via eth10.111
3.12.53.8/30       *[Direct/0] 00:23:38
                    >  via eth10.112
3.12.53.9/32       *[Local/0] 00:23:38
                       Local via eth10.112
3.12.55.8/30       *[Direct/0] 00:23:38
                    >  via eth12.113
3.12.55.9/32       *[Local/0] 00:23:38
                       Local via eth12.113
3.12.56.8/30       *[Direct/0] 00:23:38
                    >  via eth12.114
3.12.56.9/32       *[Local/0] 00:23:38
                       Local via eth12.114
10.0.0.0/24        *[Direct/0] 1d 01:57:27
                    >  via eth5.100
10.0.0.1/32        *[Local/0] 1d 01:57:27
                       Local via eth5.100
```

```
10.216.160.0/21    *[Direct/0] 1d 01:57:27
                    >  via eth0,V
10.216.161.169/32  *[Local/0] 1d 01:57:27
                       Local via eth0,V
12.2.45.0/24       *[Direct/0] 16:49:27
                    >  via eth4,V
12.2.45.9/32       *[Local/0] 16:49:27
                       Local via eth4,V
12.2.46.0/24       *[Direct/0] 1d 01:57:27
                    >  via eth2,V
12.2.46.9/32       *[Local/0] 1d 01:57:27
                       Local via eth2,V
20.0.0.0/24        *[Direct/0] 1d 01:57:27
                    >  via eth5.200
20.0.0.1/32        *[Local/0] 1d 01:57:27
                       Local via eth5.200
20.255.255.10/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.11/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.12/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.13/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.14/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.15/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
20.255.255.16/32   *[Direct/0] 1d 01:57:27
                    >  via lo.0
30.0.0.0/24        *[Direct/0] 1d 01:57:27
                    >  via eth3,V
30.0.0.1/32        *[Local/0] 1d 01:57:27
                       Local via eth3,V
172.17.0.1/32      *[Local/0] 1d 01:57:27
                       Reject
172.18.0.1/32      *[Local/0] 1d 01:57:27
                       Reject
192.168.2.254/32   *[Local/0] 1d 01:57:27
                       Reject

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

```
47.0005.1921.6810.0012/72
                   *[Direct/0] 1d 01:57:27
                    >  via lo.0

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

fe80::42:ffff:fe7a:26ec/128
                   *[Local/0] 1d 01:57:27
                       Reject
fe80::ec4:7aff:feda:43c8/128
                   *[Local/0] 1d 01:57:27
                       Local via eth0,V
fe80::92e2:baff:fecc:3ed4/128
                   *[Local/0] 00:23:38
                       Local via eth12.114
fe80::92e2:baff:fecc:3ed5/128
                   *[Local/0] 00:23:38
                       Local via eth13.110
fe80::92e2:baff:fecd:a024/128
                   *[Local/0] 16:49:27
                       Local via eth4,V
fe80::92e2:baff:fecd:a025/128
                   *[Local/0] 1d 01:57:27
                       Local
fe80::92e2:baff:fecd:a030/128
                   *[Local/0] 00:29:15
                       Local via eth6.106@
fe80::92e2:baff:fecd:a031/128
                   *[Local/0] 00:29:15
                       Local via eth7.108@
fe80::92e2:baff:fecd:a038/128
                   *[Local/0] 1d 01:57:27
                       Local via eth2,V
fe80::92e2:baff:fecd:a039/128
                   *[Local/0] 1d 01:57:27
                       Local via eth3,V
fe80::92e2:baff:fecd:a050/128
                   *[Local/0] 00:29:15
                       Local via eth8.104@
fe80::92e2:baff:fecd:a051/128
                   *[Local/0] 00:29:15
                       Local via eth9.102@
fe80::92e2:baff:fecd:a054/128
```

```
                   *[Local/0] 00:23:38
                       Local via eth10.112
fe80::92e2:baff:fecd:a055/128
                   *[Local/0] 00:23:38
                       Local via eth11.116
```

# show route active-path

### List of Syntax



### 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

### 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 172.16.100.1 via so-2/1/3.0
172.16.100.1/24       *[Direct/0] 00:18:36
                    > via so-2/1/3.0
172.16.100.1/32       *[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 456.

### 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: 172.16.100.1 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

172.16.100.0/24 (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

172.16.100.1/32 (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 -> {172.16.100.1}
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: 172.16.100.1 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

172.16.100.1/24 (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

172.16.100.1/32 (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             >172.16.100.1.
* 172.16.100.0/24       D   0                       >so-2/1/3.0
* 172.16.100.2/32       L   0                        Local
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.70.19/32   L   0                        Local
```

# show route all

### List of Syntax

Syntax on page 461
Syntax (EX Series Switches) on page 461

### 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

RELATED DOCUMENTATION

show route brief | 471
show route detail | 475

### List of Sample Output

show route all on page 462

### 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 aspath-regex

### List of Syntax

Syntax on page 464
Syntax (EX Series Switches) on page 464

### Syntax

```
show route aspath-regex regular-expression
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show route aspath-regex regular-expression
```

### 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 match the specified autonomous system (AS) path regular expression.

### Options

***regular-expression***—Regular expression that matches an entire AS path.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

### Additional Information

You can specify a regular expression as:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table of AS path regular expression operators in the *Junos Policy Framework Configuration Guide*. The following list summarizes these operators:

- **{*m,n*}**—At least *m* and at most *n* repetitions of the AS path term.
- **{*m*}**—Exactly *m* repetitions of the AS path term.
- **{*m,*}**—*m* or more repetitions of the AS path term.

- *****—Zero or more repetitions of an AS path term.
- **+**—One or more repetitions of an AS path term.
- **?**—Zero or one repetition of an AS path term.
- ***aspath_term* | *aspath_term***—Match one of the two AS path terms.

  When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

  ```
  show route aspath-regex ".* 234 .*"
  ```

**Required Privilege Level**
view

RELATED DOCUMENTATION

*Example: Using AS Path Regular Expressions*

**List of Sample Output**


**Output Fields**
For information about output fields, see the output field table for the show route command.

## Sample Output

**show route aspath-regex (Matching a Specific AS Number)**

`user@host>` **show route aspath-regex 65477**

```
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                     AS Path: [65477] ({65548 65536}) IGP
                       to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25    *[IS-IS/15] 09:15:37, metric 37, tag 1
                       to 111.222.18.225 via fpa0.0(111.222.18.233)
```

```
                     [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                    AS Path: [65477] ({65548 65536}) IGP
                      to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

**show route aspath-regex (Matching Any Path with Two AS Numbers)**

`user@host>` **show route aspath-regex ".* 234 3561 .*"**

```

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 2685 2686 Incomplete
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 5696 7369 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 6327 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

# show route best

### List of Syntax

Syntax on page 467
Syntax (EX Series Switches) on page 467

### 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

RELATED DOCUMENTATION

show route brief | 471
show route detail | 475

**List of Sample Output**



**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 468.

#### 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

## List of Syntax

Syntax on page 471
Syntax (EX Series Switches) on page 471

## 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

### RELATED DOCUMENTATION

show route all | 461
show route best | 467

## List of Sample Output

**show route brief on page 472**

### 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
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

# show route cumulative

### Syntax

```
show route cumulative
<fabric>
<logical-system (all | logical-system-name)>
<vpn-family (inet.0 | inet6.0)>
```

### Release Information

Command introduced in Junos OS Release 13.3.

### Description

Shows the cumulative number of either IPv4 or IPv6 routes in the VRF table.

### Options

**fabric**— Internal fabric state.

**logical-system (all | *logical-system-name*)**— (Optional) Show cumulative routes on all logical systems or on a particular logical system.

**vpn-family (inet.0 | inet6.0)**— Enter **inet.0** for IPv4 routes or **inet6.0** for IPv6 routes.

### Required Privilege Level

view

RELATED DOCUMENTATION

show route summary | 645

### List of Sample Output

show route cumulative on page 474

### Output Fields

| Field Name | Field Description |
|---|---|
| **destinations** | Number of destinations for which there are VRF routes in the routing table. |

| Field Name | Field Description |
| --- | --- |
| **routes** | Number of VRF routes in the routing table:<br>• **active**—Number of routes that are active.<br>• **holddown**—Number of VRF routes that are in the hold-down state before being declared inactive.<br>• **hidden**—Number of VRF routes that are not used because of routing policy. |

## Sample Output

**show route cumulative**

user@host> **show route cumulative vpn-family inet.0**

```
Total VRF Routes: 720 destinations, 722 routes (720 active, 0 holddown, 0 hidden)
```

# show route detail

### List of Syntax



### 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.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
**DeletePending** flag added to the command output in Junos OS Release 19.4R1.

### 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 (Flexible VXLAN Tunnel Profile) on page 500

### Output Fields

Table 21 on page 476 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 22 on page 483 describes all possible values for the Next-hop Types output field.

Table 23 on page 485 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 24 on page 488 describes the possible values for the Communities output field.

**Table 21: 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:<br>• **active** (routes that are active)<br>• **holddown** (routes that are in the pending state before being declared inactive)<br>• **hidden** (routes that are not used because of a routing policy) |
| ***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:<br>• ***MPLS-label*** (for example, 80001).<br>• ***interface-name*** (for example, ge-1/0/2).<br>• ***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).<br>  ◦ ***neighbor-address***—Address of the neighbor.<br>  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.<br>  ◦ ***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.<br>  ◦ ***vc-id***—Virtual circuit identifier.<br>  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.<br>  ◦ ***source***—Source of the advertisement: **Local** or **Remote**. |

**Table 21: show route detail Output Fields *(continued)***

| Field Name | Field Description |
| --- | --- |
| **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.<br>• **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).<br>• 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). |

**Table 21: show route detail Output Fields *(continued)***

<table>
<tr><th>Field Name</th><th>Field Description</th></tr>
<tr><td>[protocol, preference]</td><td>
Protocol from which the route was learned and the preference value for the route.
<ul>
<li>+—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li>
<li>- —A hyphen indicates the last active route.</li>
<li>*—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.</li>
</ul>
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.<br><br>
Preference2 values are signed integers, that is, Preference2 values can be either positive or negative values. However, Junos OS evaluates Preference2 values as unsigned integers that are represented by positive values. Based on the Preference2 values, Junos OS evaluates a preferred route differently in the following scenarios:
<ul>
<li>Both Signed Preference2 values
<ul><li>Route A = -101</li><li>Route B = -156</li></ul>
Where both the Preference2 values are signed, Junos OS evaluates only the unsigned value of Preference2 and Route A, which has a lower Preference2 value is preferred.</li>
<li>Unsigned Preference2 values<br>
Now consider both unsigned Preference2 values:
<ul><li>Route A = 4294967096</li><li>Route B = 200</li></ul>
Here, Junos OS considers the lesser Preference2 value and Route B with a Preference2 value of 200 is preferred because it is less than 4294967096.</li>
<li>Combination of signed and unsigned Preference2 values<br>
When Preference2 values of two routes are compared, and for one route the Preference2 is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive Preference2 value over the negative Preference2 value. For example, consider the following signed and unsigned Preference2 values:
<ul><li>Route A = -200</li><li>Route B = 200</li></ul>
In this case, Route B with a Preference2 value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the Preference2 value.</li>
</ul>
</td></tr>
</table>

Table 21: show route detail Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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. |
| **PMSI** | Provider multicast service interface (MVPN routing table). |
| **Next-hop type** | Type of next hop. For a description of possible values for this field, see Table 22 on page 483. |
| **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:<br><br>• 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.<br>• 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. |

Table 21: show route detail Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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 23 on page 485. |
| **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. |
| **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.<br><br>For sample output, see **show route table**. |
| **Task** | Name of the protocol that has added the route. |
| **Announcement bits** | The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table.<br><br>• ***n***—An index used by Juniper Networks customer support only. |

**Table 21: show route detail Output Fields *(continued)***

| Field Name | Field Description |
| --- | --- |
| **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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **Recorded**—The AS path is recorded by the sample process (sampled).<br>• **?**—Incomplete; typically, the AS path was aggregated.<br>When AS path numbers are included in the route, the format is as follows:<br>• **[ ]**—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.<br>• **[ ]**—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.<br>• **{ }**—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.<br>• **( )**—Parentheses enclose a confederation.<br>• **( [ ] )**—Parentheses and brackets enclose a confederation set.<br>**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:<br>• **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.<br>• **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br>• **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.<br>• **Valid**—Indicates that the prefix and autonomous system pair are found in the database. |
| ORR Generation-ID | Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted. |

Table 21: show route detail Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| Primary Upstream | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks. |
| Label | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label. |
| weight | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. |
| **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 24 on page 488 for all possible values for this field. |
| **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. |

Table 21: show route detail Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **status vector** | Layer 2 VPN and VPLS network layer reachability information (NLRI). |
| **Accepted Multipath** | Current active path when BGP multipath is configured. |
| **Accepted LongLivedStale** | The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. |
| **Accepted LongLivedStaleImport** | The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.<br><br>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table |
| **ImportAccepted**<br><br>**LongLivedStaleImport**<br><br>**DeletePending** | Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table<br><br>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.<br><br>The DeletePending flag indicates that a BGP route needs to be processed due to a BGP peer down event. |
| **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 22: Next-hop Types Output Field Values

| Next-Hop Type | Description |
|---|---|
| **Broadcast (bcast)** | Broadcast next hop. |

Table 22: Next-hop Types Output Field Values *(continued)*

| Next-Hop Type | Description |
|---|---|
| **Deny** | Deny next hop. |
| **Discard** | Discard next hop. |
| Dynamic List | Dynamic list 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. |
| **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. |

Table 22: Next-hop Types Output Field Values *(continued)*

| Next-Hop Type | Description |
|---|---|
| **Router** | A specific node or set of nodes to which the routing device forwards packets that match the route prefix.<br><br>To qualify as next-hop type router, the route must meet the following criteria:<br><br>• Must not be a direct or local subnet for the routing device.<br>• Must have a next hop that is directly connected to the routing device. |
| **Software** | Next hop added to the Routing Engine forwarding table for remote IP addresses with prefix /32 for Junos OS Evolved only. |
| **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 23: 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. |

Table 23: State Output Field Values *(continued)*

| Value | Description |
|---|---|
| **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). |
| **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). |

Table 23: State Output Field Values *(continued)*

| Value | Description |
|---|---|
| **NotInstall** | Route not to be installed in the forwarding table. |
| **NSR-incapable** | Route added by non-NSR supported protocols. |
| **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. |
| **Programmed** | Route installed programatically by on-box or off-box applications using API. |
| **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:<br>• The route is damped.<br>• The route is rejected by an import policy.<br>• The route is unresolved. |
| **Update source** | Last tiebreaker is the lowest IP address value. |
| **ProtectionCand** | Indicates paths requesting protection. |
| **ProtectionPath** | Indicates the route entry that can be used as a protection path. |

**Table 24: 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
                ORR Generation-ID: 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
                ORR Generation-ID: 1
      Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (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

...

172.16.233.22/32 (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)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
```

```
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

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

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
1.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Router, Next hop index: 580
                Address: 0x9db6ed0
                Next-hop reference count: 8
                Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                Session Id: 0x18a
                State: <Active Int>
                Local AS:     2
```

```
                Age: 1:32       Metric: 10
                Validation State: unverified
                ORR Generation-ID: 1
                Task: IS-IS
                Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR
task
                AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
1.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.1 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

**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
               Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
```

```
                 Announcement bits (1): 0-KRT
                 AS path: 2 I
                 Accepted Multipath
                 Localpref: 100
                 Router ID: 172.16.1.2
          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
              Validation State: unverified
                 Task: BGP_2.10.1.1.6+58198
                 AS path: 2 I
                 Accepted MultipathContrib
                 Localpref: 100
                 Router ID: 172.16.1.3
```

**show route detail (with BGP, DeletePending)**

`user@host>` **show route detail**

```
2:1:10.1.1.12/30 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x95c4ee8
                Next-hop reference count: 6
                Source: 10.1.1.4
                Next hop type: Router, Next hop index: 809
                Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                Label operation: Push 299888, Push 299792(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 299888: None; Label 299792: None;
                Session Id: 0x142
                Protocol next hop: 10.1.1.4
                Label operation: Push 299888
                Label TTL action: prop-ttl
                Load balance label: Label 299888: None;
                Indirect next hop: 0x96f0110 1048574 INH Session ID: 0x14e
```

```
                State: <Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     2 Peer AS:     2
                Age: 2w1d 17:42:45      Metric2: 1
                Validation State: unverified
                Task: BGP_10.2.1.1.4+55190
                AS path: I
                Communities: target:2:1
                Import Accepted DeletePending
                VPN Label: 299888
                Localpref: 100
                Router ID: 10.1.1.4
                Secondary Tables: red.inet.0
```

**show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)**

`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
                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 label detail (Multipoint LDP with Multicast-Only Fast Reroute)**

`user@host>` **show route label 301568 detail**

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1,
src: 192.168.219.11
                Primary Upstream : 172.16.1.3:0--172.16.1.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                   RPF Nexthops :
                        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
                        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```

#### show route detail (Flexible VXLAN Tunnel Profile)

`user@host>` **show route 192.168.0.2 detail**

```
...
CUSTOMER_0001.inet.0: 5618 destinations, 6018 routes (5618 active, 0 holddown, 0
hidden)
```

```
192.168.0.2/32 (1 entry, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 74781
                Address: 0x5d9b03cc
                Next-hop reference count: 363
                Next hop: via fti0.6, selected
                Session Id: 0x24c8
                State: <Active Int NSR-incapable OpaqueData Programmed>
                Age: 1:25:53
                Validation State: unverified
                        Tag: 10000001   Tag2: 1
                Announcement bits (2): 1-KRT 3-Resolve tree 30
                AS path: I
                Flexible IPv6 VXLAN tunnel profile
                    Action: Encapsulate
                    Interface: fti0.6 (Index: 10921)
                    VNI: 10000001
                    Source Prefix: 2001:db8:255::2/128
                    Source UDP Port Range: 54614 - 60074
                    Destination Address: 2001:db8:80:1:1:1:0:1
                    Destination UDP Port: 4790
                    VXLAN Flags: 0x08
...
```

# show route exact

### List of Syntax



### 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



### 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

### List of Syntax

Syntax on page 505
Syntax (EX Series Switches) on page 505

### 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**



**Output Fields**

Table 25 on page 506 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

**Table 25: 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 |
| 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:<br>• **config auto-policy**—The policy was deduced from the configured IGP export policies.<br>• **cleanup**—Configuration information for this instance is no longer valid.<br>• **config**—The instance was explicitly configured. | **detail** |
| Options | (**instance** keyword only) Configured option displays the type of routing tables the feature handles:<br>• **unicast**—Indicates *instance*.inet.0.<br>• **multicast**—Indicates *instance*.inet.2.<br>• **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 export vrf-target

### Syntax

```
show route export vrf-target
<brief | detail>
<community community--regular-expression>
<logical-system (all | logical-system-name)>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### Description

Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).

### Options

**none**—Display standard information about all target communities.

**brief | detail**—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

**community *community-regular-expression***—(Optional) Display information about the specified community.

**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 export vrf-target on page 509
show route export vrf-target community on page 509
show route export vrf-target detail on page 510

### Output Fields

Table 26 on page 508 lists the output fields for the **show route export vrf-target** command. Output fields are listed in the approximate order in which they appear.

**Table 26: show route export vrf-target Output Fields**

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Route target** | Target communities for which auto-export is currently distributing routes. | **brief** none |

Table 26: show route export vrf-target Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Family** | Routing table entries for the specified family. | **brief** none |
| *type-of-routing-table(s)* | Type of routing tables the feature handles:<br>• **unicast**—Indicates ***instance*.inet.0**.<br>• **multicast**—Indicates ***instance*.inet.2**.<br>• **unicast multicast**—Indicates ***instance*.inet.0** and ***instance*.inet.2**. | **brief** none |
| **Import** | Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes. | **brief** none |
| **Export** | Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes. | **brief** none |
| **Target** | Target communities, family, and options for which auto-export is currently distributing routes. | **detail** |
| **Import table(s)** | Name of the routing tables that are importing a particular route target. | **detail** |
| **Export table(s)** | Name of the routing tables that are exporting a particular route target. | **detail** |

## Sample Output

### show route export vrf-target

`user@host>` **show route export vrf-target**

```
Route Target                Family                   Import      Export
69:1                        inet      unicast             2           2
69:2                        inet      unicast             2           2
```

### show route export vrf-target community

`user@host>` **show route export vrf-target community target:69:1**

```
Route Target                Family                   Import      Export
69:1                        inet      unicast             2           2
```

**show route export vrf-target detail**

user@host> **show route export vrf-target detail**

```
Target: 1:12                              inet     unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13                              inet     unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```

# show route extensive

### List of Syntax



### Syntax

```
show route extensive
<destination-prefix>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show route extensive
<destination-prefix>
```

### Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
**DeletePending** flag added to the command output in Junos OS Release 19.4R1.

### Description

Display extensive 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 label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 534
show route extensive (Flexible VXLAN Tunnel Profile) on page 535

### Output Fields

Table 27 on page 512 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

**Table 27: show route extensive 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:<br>• **active** (routes that are active).<br>• **holddown** (routes that are in the pending state before being declared inactive).<br>• **hidden** (routes that are not used because of a routing policy). |
| ***route-destination*** **(entry, announced)** | Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route 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:<br>• ***MPLS-label*** (for example, 80001 ).<br>• ***interface-name*** (for example, ge-1/0/2).<br>• ***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).<br>  ◦ ***neighbor-address***—Address of the neighbor.<br>  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.<br>  ◦ ***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.<br>  ◦ ***vc-id***—Virtual circuit identifier.<br>  ◦ ***source***—Source of the advertisement: **Local** or **Remote**. |
| **TSI** | Protocol header information. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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.<br>• **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).<br>• 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.<br>• **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br>• **-** —A hyphen indicates the last active route.<br>• *—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.<br>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 autonomous system (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. |
| **PMSI** | Provider multicast service interface (MVPN routing table). |
| **Next-hop type** | Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command. |
| **Next-hop reference count** | Number of references made to the next hop. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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:<br>• 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.<br>• 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). |
| **Offset** | Whether the metric has been increased or decreased by an offset value. |
| **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 recursively derive a forwarding 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). |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **Indirect next hops** | When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.<br><br>When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:<br><br>• 0x1 indicates active next hops.<br>• 0x4000 indicates passive next hops. |
| **State** | State of the route (a route can be in more than one state). See the Output Field table in the show route detail command. |
| **Session ID** | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA). |
| **Weight** | Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.<br><br>For sample output, see show route table. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **Inactive reason** | If the route is inactive, the reason for its current state is indicated. Typical reasons include:<br>• **Active preferred**—Currently active route was selected over this route.<br>• **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available.<br>• **AS path**—Shorter AS path is available.<br>• **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available.<br>• **Cluster list length**—Path with a shorter cluster list length is available.<br>• **Forwarding use only**—Path is only available for forwarding purposes.<br>• **IGP metric**—Path through the next hop with a lower IGP metric is available.<br>• **IGP metric type**—Path with a lower OSPF link-state advertisement type is available.<br>• **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available.<br>• **Local preference**—Path with a higher local preference value is available.<br>• **Next hop address**—Path with a lower metric next hop is available.<br>• **No difference**—Path from a neighbor with a lower IP address is available.<br>• **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).<br>• **Number of gateways**—Path with a higher number of next hops is available.<br>• **Origin**—Path with a lower origin code is available.<br>• **OSPF version**—Path does not support the indicated OSPF version.<br>• **RIB preference**—Route from a higher-numbered routing table is available.<br>• **Route destinguisher**—64-bit prefix added to IP subnets to make them unique.<br>• **Route metric or MED comparison**—Route with a lower metric or MED is available.<br>• **Route preference**—Route with a lower preference value is available.<br>• **Router ID**—Path through a neighbor with a lower ID is available.<br>• **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, or the route is unresolved.<br>• **Update source**—Last tiebreaker is the lowest IP address value. |
| **Local AS** | Autonomous system (AS) number of the local routing device. |
| **Age** | How long the route has been known. |
| **AIGP** | Accumulated interior gateway protocol (AIGP) BGP attribute. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **Metric** | 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. |
| **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.<br><br>For sample output, see show route table. |
| **Task** | Name of the protocol that has added the route. |
| **Announcement bits** | List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP).<br><br>The notation ***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. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **Recorded**—The AS path is recorded by the sample process (sampled).<br>• **?**—Incomplete; typically, the AS path was aggregated.<br>When AS path numbers are included in the route, the format is as follows:<br>• **[ ]**—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.<br>• **{ }**—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.<br>• **( )**—Parentheses enclose a confederation.<br>• **( [ ] )**—Parentheses and brackets enclose a confederation set.<br>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:<br>• **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.<br>• **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br>• **Unverified**—Indicates that origin validation is not enabled for the BGP peers.<br>• **Valid**—Indicates that the prefix and autonomous system pair are found in the database. |
| 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. |
| **AS path: I <Originator>** | (For route reflected output only) Originator ID attribute set by the route reflector. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **route status** | Indicates the status of a BGP route:<br>• **Accepted**—The specified BGP route is imported by the default BGP policy.<br>• **Import**—The route is imported into a Layer 3 VPN routing instance.<br>• **Import-Protect**—A remote instance egress that is protected.<br>• **Multipath**—A BGP multipath active route.<br>• **MultipathContrib**—The route is not active but contributes to the BGP multipath.<br>• **Protect**—An egress route that is protected.<br>• **Stale**—A route that is marked stale due to graceful restart. |
| Primary Upstream | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks. |
| Label | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label. |
| weight | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. |
| **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. |
| **Cluster list** | (For route reflected output only) Cluster ID sent by the route reflector. |
| **Originator ID** | (For route reflected output only) Address of router that originally sent the route to the route reflector. |
| **Prefixes bound to route** | Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP. |

Table 27: show route extensive Output Fields *(continued)*

| Field Name | Field Description |
| --- | --- |
| **Communities** | Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field. |
| **DeletePending** | The DeletePending flag indicates that a BGP route needs to be processed due to a BGP peer down event. See the show route detail command for example output. |
| **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). |
| **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. |
| **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. |
| **Node path count** | Number of nodes in the path. |
| **Forwarding nexthops** | 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 extensive

user@host> **show route extensive**

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {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:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (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:    64496
                Age: 1:32:40
                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:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (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:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
```

```
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                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)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (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:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                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)
```

```
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:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
```

```
                    Address: 0x92544f0
                    Next-hop reference count: 2
                    Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                    Label-switched-path R2-to-R200-p2mp
                    Label operation: Pop
                    Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                    Label operation: Pop
                    State: <Active Int>
                    Age: 1:29       Metric: 1
                    Task: RSVP
                    Announcement bits (1): 0-KRT
                    AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                  Next-hop reference count: 2
                  Next hop: via vt-3/2/0.32769, selected
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:31:53
                  Task: Common L2 VC
                  Announcement bits (1): 0-KRT
                  AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                  Next-hop reference count: 2
                  Next hop: 203.0.113.216 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: 203.0.113.103
                  Push 800012
                  Indirect next hop: 87272e4 1048574
                  State: <Active Int>
                  Age: 1:31:53    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
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                        Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::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:    64496
                Age: 1:34:07
                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:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
```

```
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                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: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
```

```
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+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: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    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

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    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)

TSI:

203.0.113.163: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: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
```

```
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
                Route Label: 300112
                Localpref: 100
                Router ID: 213.0.113.99
```

**show route extensive (Access Route)**

`user@host>` **show route 203.0.113.102 extensive**

```
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holdown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I
```

`user@host>` **show route 2001:db8:4641:1::/48 extensive**

```

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
```

```
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

#### show route extensive (BGP PIC Edge)

`user@host>` **show route 198.51.100.6 extensive**

```
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
          State: <CalcForwarding>
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
          #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
```

```
                    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                    Indirect next hop: weight 0x4000
                    State: <ForwardingOnly Int Ext)>
                    Inactive reason: Forwarding use only
                    Age: 25         Metric2: 15
                    Validation State: unverified
                    Task: RT
                    Announcement bits (1): 0-KRT
                    AS path: 3 I
                    Communities: target:2:1
```

**show route extensive (FRR and LFA)**

`user@host>` **show route 203.0.113.20 extensive**

```
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
                Address: 0xb9d78c4
                Next-hop reference count: 7
```

```
                Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                Session Id: 0x201
                State: Int
                Inactive reason: Route Preference
                Local AS:   64500
                Age: 5:35  Metric: 3
                Area: 0.0.0.0
                Task: OSPF
                AS path: I
```

**show route extensive (IS-IS)**

user@host> **show route extensive**

```
IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 203.0.113.22 via ae0.0 balance 57%
```

**show route extensive (Route Reflector)**

user@host> **show route extensive**

```
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
```

```
                  Localpref: 100
                  Router ID: 203.0.113.4
                  Indirect next hops: 1
                          Protocol next hop: 203.0.113.192 Metric: 0
                          Indirect next hop: 84ac908 40
                          Indirect path forwarding next hops: 0
                                  Next hop type: Discard
```

**show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)**

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
                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:  64511
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.113.1,
 src 192.168.142.2
```

**show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)**

user@host> **show route label 301568 detail**

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
```

```
                  Address: 0x2735208
                  Next-hop reference count: 3
                  Next hop type: Router, Next hop index: 1397
                  Address: 0x2735d2c
                  Next-hop reference count: 3
                  Next hop: 203.0.113.82 via ge-1/2/22.0
                  Label operation: Pop
                  Load balance label:  None;
                  Next hop type: Router, Next hop index: 1395
                  Address: 0x2736290
                  Next-hop reference count: 3
                  Next hop: 203.0.113.2 via ge-1/2/18.0
                  Label operation: Pop
                  Load balance label:  None;
                  State: <Active Int AckRequest MulticastRPF>
                  Local AS:    64500
                  Age: 54:05      Metric: 1
                  Validation State: unverified
                  Task: LDP
                  Announcement bits (1): 0-KRT
                  AS path: I
                FECs bound to route: P2MP root-addr 198.51.100.1, grp: 203.0.113.1,
src: 192.168.219.11
                  Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                     RPF Nexthops  :
                          ge-1/2/15.0, 10.2.94.1, Label: 301568, weight: 0x1
                          ge-1/2/14.0, 10.2.3.1, Label: 301568, weight: 0x1
                  Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                     RPF Nexthops :
                          ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                          ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```

**show route extensive (Flexible VXLAN Tunnel Profile)**

`user@host>` **show route 192.168.0.2 extensive**

```
...
CUSTOMER_0001.inet.0: 5618 destinations, 6018 routes (5618 active, 0 holddown, 0
hidden)

192.168.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.2/32 -> {fti0.6 Flags NSR-incapable}
```

```
Opaque data client: FLEX-TNL
Address: 0xd00eee8
Opaque-data reference count: 2
Opaque data: Flexible IPv6 VXLAN tunnel profile
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 74781
                Address: 0x5d9b03cc
                Next-hop reference count: 363
                Next hop: via fti0.6, selected
                Session Id: 0x24c8
                State: <Active Int NSR-incapable OpaqueData Programmed>
                Age: 1:34:00
                Validation State: unverified
                        Tag: 10000001   Tag2: 1
                Announcement bits (2): 1-KRT 3-Resolve tree 30
                AS path: I
                Flexible IPv6 VXLAN tunnel profile
                    Action: Encapsulate
                    Interface: fti0.6 (Index: 10921)
                    VNI: 10000001
                    Source Prefix: 2001:db8:255::2/128
                    Source UDP Port Range: 54614 - 60074
                    Destination Address: 2001:db8:80:1:1:1:0:1
                    Destination UDP Port: 4790
                    VXLAN Flags: 0x08
...
```

# show route export vrf-target

## Syntax

```
show route export vrf-target
<brief | detail>
<community community--regular-expression>
<logical-system (all | logical-system-name)>
```

## Release Information

Command introduced before Junos OS Release 7.4.

## Description

Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).

## Options

**none**—Display standard information about all target communities.

**brief | detail**—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

**community *community-regular-expression***—(Optional) Display information about the specified community.

**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 export vrf-target on page 538
show route export vrf-target community on page 538
show route export vrf-target detail on page 539

## Output Fields

Table 26 on page 508 lists the output fields for the **show route export vrf-target** command. Output fields are listed in the approximate order in which they appear.

**Table 28: show route export vrf-target Output Fields**

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Route target** | Target communities for which auto-export is currently distributing routes. | **brief** none |

Table 28: show route export vrf-target Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Family** | Routing table entries for the specified family. | **brief** none |
| *type-of-routing-table(s)* | Type of routing tables the feature handles:<br>• **unicast**—Indicates ***instance*.inet.0**.<br>• **multicast**—Indicates ***instance*.inet.2**.<br>• **unicast multicast**—Indicates ***instance*.inet.0** and ***instance*.inet.2**. | **brief** none |
| **Import** | Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes. | **brief** none |
| **Export** | Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes. | **brief** none |
| **Target** | Target communities, family, and options for which auto-export is currently distributing routes. | **detail** |
| **Import table(s)** | Name of the routing tables that are importing a particular route target. | **detail** |
| **Export table(s)** | Name of the routing tables that are exporting a particular route target. | **detail** |

## Sample Output

### show route export vrf-target

user@host> **show route export vrf-target**

```
Route Target                 Family                    Import      Export
69:1                         inet      unicast              2           2
69:2                         inet      unicast              2           2
```

### show route export vrf-target community

user@host> **show route export vrf-target community target:69:1**

```
Route Target                 Family                    Import      Export
69:1                         inet      unicast              2           2
```

**show route export vrf-target detail**

`user@host>` **show route export vrf-target detail**

```
Target: 1:12                              inet     unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13                              inet     unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```

# show route forwarding-table

### List of Syntax

Syntax on page 540
Syntax (MX Series Routers) on page 540
Syntax (TX Matrix and TX Matrix Plus Routers) on page 540

### Syntax

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all | vlan-name)>
<vpn vpn>
```

### Syntax (MX Series Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | domain-name)>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<learning-vlan-id learning-vlan-id>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all |  vlan-name)>
<vpn vpn>
```

### Syntax (TX Matrix and TX Matrix Plus Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<matching matching>
<label name>
<lcc number>
<multicast>
<table routing-instance-name>
<vpn vpn>
```

**Release Information**

Command introduced before Junos OS Release 7.4.
Option **bridge-domain** introduced in Junos OS Release 7.5
Option **learning-vlan-id** introduced in Junos OS Release 8.4
Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**

Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

> **NOTE:** The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

**Options**

**none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

**detail | extensive | summary**—(Optional) Display the specified level of output.

**all**—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

**bridge-domain (all | *bridge-domain-name*)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

**ccc *interface-name***—(Optional) Display route entries for the specified circuit cross-connect interface.

**destination *destination-prefix***—(Optional) Destination prefix.

**family *family***—(Optional) Display routing table entries for the specified family: **bridge (ccc | destination | detail | extensive | interface-name | label | learning-vlan-id | matching | multicast | summary | table | vlan | vpn)**, **ethernet-switching**, **evpn**, **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mcsnoop-inet**, **mcsnoop-inet6**, **mpls**, **satellite-inet**, **satellite-inet6**, **satellite-vpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

**interface-name *interface-name***—(Optional) Display routing table entries for the specified interface.

**label *name***—(Optional) Display route entries for the specified label.

**lcc *number***—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**learning-vlan-id *learning-vlan-id***—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

**matching *matching***—(Optional) Display routing table entries matching the specified prefix or prefix length.

**multicast**—(Optional) Display routing table entries for multicast routes.

**table** —(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the **show route instance** command.

**vlan (all | *vlan-name*)**—(Optional) Display information for all VLANs or for the specified VLAN.

**vpn *vpn***—(Optional) Display routing table entries for a specified VPN.

### Required Privilege Level

view

### List of Sample Output



### Output Fields

Table 29 on page 543 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 29: show route forwarding-table Output Fields**

| Field Name | Field Description | Level of Output |
|---|---|---|
| Logical system | Name of the logical system. This field is displayed if you specify the **table *logical-system-name/routing-instance-name*** option on a device that is configured for and supports logical systems. | All levels |
| Routing table | Name of the routing table (for example, inet, inet6, mpls). | All levels |

**Table 29: show route forwarding-table Output Fields *(continued)***

| Field Name | Field Description | Level of Output |
|---|---|---|
| Enabled protocols | | All levels |

**Table 29: show route forwarding-table Output Fields *(continued)***

| Field Name | Field Description | Level of Output |
|---|---|---|
| | The features and protocols that have been enabled for a given routing table. This field can contain the following values:<br>• BUM hashing—BUM hashing is enabled.<br>• MAC Stats—Mac Statistics is enabled.<br>• Bridging—Routing instance is a normal layer 2 bridge.<br>• No VLAN—No VLANs are associated with the bridge domain.<br>• All VLANs—The **vlan-id all** statement has been enabled for this bridge domain.<br>• Single VLAN—Single VLAN ID is associated with the bridge domain.<br>• MAC action drop—New MACs will be dropped when the MAC address limit is reached.<br>• Dual VLAN—Dual VLAN tags are associated with the bridge domain<br>• No local switching—No local switching is enabled for this routing instance..<br>• Learning disabled—Layer 2 learning is disabled for this routing instance.<br>• MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached.<br>• VPLS—The VPLS protocol is enabled.<br>• No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance.<br>• ACKed by all peers—All peers have acknowledged this routing instance.<br>• BUM Pruning—BUM pruning is enabled on the VPLS instance.<br>• Def BD VXLAN—VXLAN is enabled for the default bridge domain.<br>• EVPN—EVPN protocol is enabled for this routing instance.<br>• Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain.<br>• Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain.<br>• L2 backhaul—Layer 2 backhaul is enabled.<br>• FRR optimize—Fast reroute optimization<br>• MAC pinning—MAC pinning is enabled for this bridge domain.<br>• MAC Aging Timer—The MAC table aging time is set per routing instance.<br>• EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation.<br>• PBBN—This routing instance is configured as a provider backbone bridged network. | |

Table 29: show route forwarding-table Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| | • PBN—This routing instance is configured as a provider bridge network.<br>• ETREE—The ETREE protocol is enabled on this EVPN routing instance.<br>• ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance.<br>• Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain.<br>• MPLS control word—Control word is enabled for this MPLS routing instance. | |
| Address family | Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**, and **VPLS**). | All levels |
| Destination | Destination of the route. | **detail extensive** |
| Route Type (Type) | How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses):<br>• **cloned** (**clon**)—(TCP or multicast only) Cloned route.<br>• **destination** (**dest**)—Remote addresses directly reachable through an interface.<br>• **destination down** (**iddn**)—Destination route for which the interface is unreachable.<br>• **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable.<br>• **route down** (**ifdn**)—Interface route for which the interface is unreachable.<br>• **ignore** (**ignr**)—Ignore this route.<br>• **interface** (**intf**)—Installed as a result of configuring an interface.<br>• **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized.<br>• **user**—Routes installed by the routing protocol process or as a result of the configuration. | All levels |
| Route Reference (RtRef) | Number of routes to reference. | **detail extensive** |

Table 29: show route forwarding-table Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Flags | Route type flags:<br>• **none**—No flags are enabled.<br>• **accounting**—Route has accounting enabled.<br>• **cached**—Cache route.<br>• **incoming-iface *interface-number***—Check against incoming interface.<br>• **prefix load balance**—Load balancing is enabled for this prefix.<br>• **rt nh decoupled**—Route has been decoupled from the next hop to the destination.<br>• **sent to PFE**—Route has been sent to the Packet Forwarding Engine.<br>• **static**—Static route. | **extensive** |
| Next hop | IP address of the next hop to the destination.<br>NOTE: For static routes that use point-to-point (P2P) outgoing interfaces, the next-hop address is not displayed in the output. | **detail extensive** |
| Next hop Type (Type) | Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):<br>• **broadcast (bcst)**—Broadcast.<br>• **deny**—Deny.<br>• **discard (dscd)** —Discard.<br>• **hold**—Next hop is waiting to be resolved into a unicast or multicast type.<br>• **indexed (idxd)**—Indexed next hop.<br>• **indirect (indr)**—Indirect next hop.<br>• **local (locl)**—Local address on an interface.<br>• **routed multicast (mcrt)**—Regular multicast next hop.<br>• **multicast (mcst)**—Wire multicast next hop (limited to the LAN).<br>• **multicast discard (mdsc)**—Multicast discard.<br>• **multicast group (mgrp)**—Multicast group member.<br>• **receive (recv)**—Receive.<br>• **reject (rjct)**—Discard. An ICMP unreachable message was sent.<br>• **resolve (rslv)**—Resolving the next hop.<br>• **unicast (ucst)**—Unicast.<br>• **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. | **detail extensive** |

Table 29: show route forwarding-table Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Index | Software index of the next hop that is used to route the traffic for a given prefix. | **detail extensive** none |
| Route interface-index | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned. | **extensive** |
| Reference (NhRef) | Number of routes that refer to this next hop. | **detail extensive** none |
| Next-hop interface (Netif) | Interface used to reach the next hop. | **detail extensive** none |
| 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 (see the **Balance** field description). | **extensive** |
| Balance | Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. | **extensive** |
| RPF interface | List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when **rpf-check** is configured on the interface. | **extensive** |

## Sample Output

### show route forwarding-table

`user@host>` **show route forwarding-table**

```
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
```

```
default             perm     0                     rjct    46     4
0.0.0.0/32          perm     0                     dscd    44     1
172.16.1.0/24          ifdn     0                     rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0            recv   606     1 ge-2/0/1.0
172.16.1.1/32          user     0                     rjct    46     4
172.16.1.1/32          intf     0 172.16.1.1            locl   607     2
172.16.1.1/32          iddn     0 172.16.1.1            locl   607     2
172.16.1.255/32        iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24         intf     0                     rslv   616     1 ge-2/0/0.0
10.0.0.0/32         dest     0 10.0.0.0            recv   614     1 ge-2/0/0.0
10.0.0.1/32         intf     0 10.0.0.1            locl   615     2
10.0.0.1/32         dest     0 10.0.0.1            locl   615     2
10.0.0.255/32       dest     0 10.0.0.255          bcst   613     1 ge-2/0/0.0
10.1.1.0/24         ifdn     0                     rslv   612     1 ge-2/0/1.0
10.1.1.0/32         iddn     0 10.1.1.0            recv   610     1 ge-2/0/1.0
10.1.1.1/32         user     0                     rjct    46     4
10.1.1.1/32         intf     0 10.1.1.1            locl   611     2
10.1.1.1/32         iddn     0 10.1.1.1            locl   611     2
10.1.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff   bcst   609     1 ge-2/0/1.0
10.206.0.0/16       user     0 10.209.63.254       ucst   419    20 fxp0.0
10.209.0.0/16       user     1 0:12:1e:ca:98:0     ucst   419    20 fxp0.0
10.209.0.0/18       intf     0                     rslv   418     1 fxp0.0
10.209.0.0/32       dest     0 10.209.0.0          recv   416     1 fxp0.0
10.209.2.131/32     intf     0 10.209.2.131        locl   417     2
10.209.2.131/32     dest     0 10.209.2.131        locl   417     2
10.209.17.55/32     dest     0 0:30:48:5b:78:d2    ucst   435     1 fxp0.0
10.209.63.42/32     dest     0 0:23:7d:58:92:ca    ucst   434     1 fxp0.0
10.209.63.254/32    dest     0 0:12:1e:ca:98:0     ucst   419    20 fxp0.0
10.209.63.255/32    dest     0 10.209.63.255       bcst   415     1 fxp0.0
10.227.0.0/16       user     0 10.209.63.254       ucst   419    20 fxp0.0

...

Routing table: iso
ISO:
Destination         Type RtRef Next hop           Type Index NhRef Netif
default             perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination         Type RtRef Next hop           Type Index NhRef Netif
default             perm     0                   rjct     6     1
```

```
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

**show route forwarding-table detail**

`user@host>` **show route forwarding-table detail**

```
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132      4 fxp0.0
default            perm     0                    rjct    14      1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322      1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324      1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321      1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323      1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326      1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328      1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325      1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327      1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320      1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132      4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132      4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46      1
10.0.0.0/8         intf     0                    rslv   136      1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134      1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135      2
10.0.0.4/32        dest     0 10.0.0.4           locl   135      2

...

Routing table: iso
```

```
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

### show route forwarding-table destination extensive (Weights and Balances)

`user@host>` **show route forwarding-table destination 3.4.2.1 extensive**

```
Routing table: inet [Index 0]
Internet:

Destination:  3.4.2.1/32
  Route type: user
  Route reference: 0                  Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: unilist              Index: 262143   Reference: 1
  Nexthop: 172.16.4.4
  Next-hop type: unicast              Index: 335      Reference: 2
  Next-hop interface: so-1/1/0.0      Weight: 22      Balance: 3
  Nexthop: 145.12.1.2
  Next-hop type: unicast              Index: 337      Reference: 2
  Next-hop interface: so-0/1/2.0      Weight: 33      Balance: 33
```

### show route forwarding-table extensive

`user@host>` **show route forwarding-table extensive**

```
Routing table: inet [Index 0]
Internet:
```

```
Destination:  default
  Route type: user
  Route reference: 2                     Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast                 Index: 132      Reference: 4
  Next-hop interface: fxp0.0

Destination:  default
  Route type: permanent
  Route reference: 0                     Route interface-index: 0
  Flags: none
  Next-hop type: reject                  Index: 14       Reference: 1

Destination:  127.0.0.1/32
  Route type: interface
  Route reference: 0                     Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local                   Index: 320      Reference: 1

...

Routing table: private1__.inet [Index 1]
Internet:

Destination:  default
  Route type: permanent
  Route reference: 0                     Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                  Index: 46       Reference: 1

Destination:  10.0.0.0/8
  Route type: interface
  Route reference: 0                     Route interface-index: 3
  Flags: sent to PFE
  Next-hop type: resolve                 Index: 136      Reference: 1
  Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:
```

```
Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 38       Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 22       Reference: 1

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard     Index: 21       Reference: 1

...

Routing table: private1__.inet6 [Index 1]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 54       Reference: 1

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                 Index: 75       Reference: 1

...
```

### show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

user@host> **show route forwarding-table extensive**

```
Routing table: inet [Index 0]
Internet:
...
...
Destination:  192.0.2.3/32
  Route type: destination
  Route reference: 0                   Route interface-index: 67
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  Next-hop type: broadcast             Index: 328      Reference: 1
  Next-hop interface: so-1/1/0.0
  RPF interface: so-1/1/0.0
```

### show route forwarding-table extensive (PIM using point-to-multipoint mode)

user@host> **show route forwarding-table extensive**

```
Destination:  198.51.100.0/24
  Route type: user
  Route reference: 0                   Route interface-index: 335
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: cached, check incoming interface , accounting, sent to PFE, rt nh
decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 4
  Nexthop:
  Next-hop type: composite             Index: 627      Reference: 1
```

```
  Next-hop type: unicast               Index: 1048574  Reference: 2
  Next-hop interface: st0.1, 192.0.2.0
```

**show route forwarding-table (dynamic list next hop)**

The **show route forwarding table** output shows the two next hop elements for a multihomed EVPN destination.

`user@host> show route forwarding-table label 299952 extensive`

```
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048574  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301632, Push 299776(top) Index: 600 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

After one of the PE router has been disabled in the EVPN multihomed network, the same **show route forwarding table** output command shows one next hop element and one empty next hop element.

`user@host> show route forwarding-table label 299952 extensive`

```
Routing table: default.mpls [Index 0]
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
```

```
Multicast RPF nh index: 0
P2mpidx: 0
Flags: sent to PFE, rt nh decoupled
Next-hop type: indirect              Index: 1048575  Reference: 2
Nexthop:
Next-hop type: composite             Index: 601      Reference: 2
Next-hop type: indirect              Index: 1048577  Reference: 3
Nexthop: 1.0.0.4
Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
Load Balance Label: None
Next-hop interface: ge-0/0/1.0
```

### show route forwarding-table family mpls

user@host> **show route forwarding-table family mpls**

```
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    19     1
0                  user     0                   recv    18     3
1                  user     0                   recv    18     3
2                  user     0                   recv    18     3
100000             user     0 10.31.1.6         swap  100001     fe-1/1/0.0
800002             user     0                  Pop                vt-0/3/0.32770
vt-0/3/0.32770 (VPLS)
                   user     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
```

### show route forwarding-table family mpls ccc ge-0/0/1.1004

user@host>**show route forwarding-table mpls ccc ge-0/0/1.1004**

```
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop          Type  Index    NhRef Netif
ge-0/0/1.1004  (CCC) user     0                     ulst   1048577     2
                                                    comp      754         3
                                                  comp      755         3
                                                comp      756         3

Routing table: __mpls-oam__.mpls
```

```
MPLS:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      556     1
```

**show route forwarding-table family vpls**

`user@host>` **show route forwarding-table family vpls**

```
Routing table: green.vpls
VPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            dynm     0                   flood   353     1
default            perm     0                   rjct    298     1
fe-0/1/0.0         dynm     0                   flood   355     1
00:00:5E:00:53:1f/48             <<<<<Remote CE

                   dynm     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
00:00:5E:00:53:1f/48             <<<<<<Local CE

                   dynm     0                   ucst   354     2 fe-0/1/0.0
```

**show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)**

`user@host>` **show route forwarding-table vpls**

```
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      519     1
lsi.1048832        intf     0                    indr  1048574     4
                              172.16.3.2           Push 262145      621     2
ge-3/0/0.0
00:00:5E:00:53:01/48 user     0                  ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                    comp      627     2
ge-2/3/9.0         intf     0                    ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                    ucst      619     4 ge-3/1/3.0
0x30002/51         user     0                    comp      600     2
0x30001/51         user     0                    comp      597     2
```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

`user@host>` **show route forwarding-table vpls**

```
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      519     1
lsi.1048834        intf     0                    indr  1048574     4
                              172.16.3.2           Push 262145      592     2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user     0                  ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                    comp      630     2
ge-2/3/9.0         intf     0                    ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                    ucst      591     4 ge-3/1/3.0
0x30002/51         user     0                    comp      627     2
0x30001/51         user     0                    comp      624     2
```

### show route forwarding-table family vpls extensive

`user@host>` **show route forwarding-table family vpls extensive**

```
Routing table: green.vpls [Index 2]
VPLS:

Destination:  default
  Route type: dynamic
  Route reference: 0                   Route interface-index: 72
  Flags: sent to PFE
  Next-hop type: flood                 Index: 289      Reference: 1
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: none
  Next-hop type: discard               Index: 341      Reference: 1

Destination:  fe-0/1/2.0
  Route type: dynamic
```

```
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE
  Next-hop type: flood                 Index: 293      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0

Destination:  fe-0/1/3.0
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE
  Next-hop type: flood                 Index: 292      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  00:00:5E:00:53:01/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Route used as destination:
    Packet count:         6640    Byte count:      675786
  Route used as source
    Packet count:         6894    Byte count:      696424

Destination:  00:00:5E:00:53:04/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 290      Reference: 3
```

```
  Next-hop interface: fe-0/1/2.0
  Route used as destination:
    Packet count:           96    Byte count:        8079
  Route used as source:
    Packet count:          296    Byte count:       24955

Destination:  00:00:5E:00:53:05/48
  Route type: dynamic
  Route reference: 0                     Route interface-index: 74
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect                Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
```

**show route forwarding-table table default**

`user@host>` **show route forwarding-table table default**

```
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
10.0.60.0/30       user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.60.12/30      intf     0                    rslv   688     1 fe-0/1/3.0
10.0.60.12/32      dest     0 10.0.60.12         recv   686     1 fe-0/1/3.0
10.0.60.13/32      dest     0 0:5:85:8b:bc:22    ucst   713     5 fe-0/1/3.0
10.0.60.14/32      intf     0 10.0.60.14         locl   687     2
10.0.60.14/32      dest     0 10.0.60.14         locl   687     2
10.0.60.15/32      dest     0 10.0.60.15         bcst   685     1 fe-0/1/3.0
10.0.67.12/30      user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.80.0/30       ifdn     0 ff.3.0.21          ucst   676     1 so-0/0/1.0
10.0.80.0/32       dest     0 10.0.80.0          recv   678     1 so-0/0/1.0
10.0.80.2/32       user     0                    rjct    36     2
10.0.80.2/32       intf     0 10.0.80.2          locl   675     1
10.0.80.3/32       dest     0 10.0.80.3          bcst   677     1 so-0/0/1.0
10.0.90.12/30      intf     0                    rslv   684     1 fe-0/1/0.0
10.0.90.12/32      dest     0 10.0.90.12         recv   682     1 fe-0/1/0.0
10.0.90.14/32      intf     0 10.0.90.14         locl   683     2
10.0.90.14/32      dest     0 10.0.90.14         locl   683     2
10.0.90.15/32      dest     0 10.0.90.15         bcst   681     1 fe-0/1/0.0
10.5.0.0/16        user     0 192.168.187.126    ucst   324    15 fxp0.0
10.10.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
```

```
10.13.10.0/23      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.84.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.150.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.157.64.0/19     user     0 192.168.187.126    ucst   324    15 fxp0.0
10.209.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0

...

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    60     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1
::/128             perm     0                    dscd    42     1
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
```

**show route forwarding-table table logical-system-name/routing-instance-name**

`user@host>` **show route forwarding-table table R4/vpn-red**

```
Logical system: R4
Routing table: vpn-red.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   563     1
0.0.0.0/32         perm     0                    dscd   561     2
172.16.0.1/32         user     0                    dscd   561     2
172.16.2.0/24         intf     0                    rslv   771     1 ge-1/2/0.3
172.16.2.0/32        dest     0 172.16.2.0           recv   769     1 ge-1/2/0.3
172.16.2.1/32         intf     0 172.16.2.1           locl   770     2
172.16.2.1/32         dest     0 172.16.2.1           locl   770     2
172.16.2.2/32         dest     0 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0
                                                 ucst   789     1 ge-1/2/0.3
172.16.2.255/32      dest     0 172.16.2.255         bcst   768     1 ge-1/2/0.3
```

```
172.16.233.0/4        perm     1                    mdsc   562     1
172.16.233.1/32       perm     0 172.16.233.1          mcst   558     1
255.255.255.255/32 perm     0                    bcst   559     1

Logical system: R4
Routing table: vpn-red.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   608     1

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   708     1
::/128             perm     0                    dscd   706     1
ff00::/8           perm     0                    mdsc   707     1
ff02::1/128        perm     0 ff02::1            mcst   704     1

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   638
```

**show route forwarding-table vpn**

user@host> **show route forwarding-table vpn VPN-A**

```
      Routing table:: VPN-A.inet
Internet:
      Destination        Type RtRef Nexthop            Type Index NhRef Netif
      default            perm     0                    rjct     4     4
      10.39.10.20/30     intf     0 ff.3.0.21          ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21        locl    36     1
      10.255.14.172/32   user     0                    ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                    indr    81     3
                                                       Push 100004, Push
      100004(top) so-1/0/0.0
      172.16.233.0/4        perm     2                    mdsc     5     3
      172.16.233.1/32       perm     0 172.16.233.1          mcst     1     8
```

```
172.16.233.5/32        user      1 172.16.233.5         mcst      1      8
255.255.255.255/32 perm      0                     bcst      2      3
```

On QFX5200, the results for this command look like this:

`show route forwarding-table family mpls`

```
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
300400 user 0 ulst 131071 2
                              10.1.1.2 Pop 1713 1 xe-0/0/38.0
                              172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                              10.1.1.2 Pop 1715 1 xe-0/0/38.0
                              172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1
```

# show route forwarding-table interface-name

## Syntax

```
show route forwarding-table interface-name interface-name
<detail | extensive>
<all>
```

## Release Information

Command introduced in Junos OS Release 9.6.

## Description

Display the interfaces in the Routing Engine's forwarding table.

## Options

**none**—Display information for the specified interface.

**detail | extensive**—(Optional) Display the specified level of output.

**all**—(Optional) Display all interfaces in the routing table.

## Required Privilege Level

view

## List of Sample Output

show route forwarding-table interface-name fe-0/1/1 on page 565
show route forwarding-table interface-name all on page 565
show route forwarding-table interface-name all detail on page 566

## Output Fields

Table 30 on page 564 lists the output fields for the **show route forwarding-table interface-name** command. Output fields are listed in the approximate order in which they appear.

Table 30: show route forwarding-table interface-name Output Fields

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Name** | Name of the interface (for example **fe-0/1/1**, **lo0**, **ae0**, and so on). | All levels |
| **MTU** | Interface's maximum transmission unit (MTU). | All levels |
| **Afam** | Configured address family (for example **inet**, **tnp**, **inet6**, and so on). | **detail extensive** |

Table 30: show route forwarding-table interface-name Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **Network** | Network information: <br>• **<Link>**—Physical interface, not a logical interface. <br>• **<PtoP>**—Point-to-point network. <br>• **ipaddress**—Network address. | All levels |
| **Address** | Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on. | All levels |
| **IPkts** | Number of packets received on the interface. | All levels |
| **Ierr** | Number of packets received on the interface with errors. | All levels |
| **Opkts** | Number of packets transmitted or sent from the interface. | All levels |
| **Oerr** | Number of packets transmitted or sent from the interface with errors. | All levels |
| **Coll** | Number of packets that experienced collisions on the interface. | All levels |

## Sample Output

### show route forwarding-table interface-name fe-0/1/1

`user@host>` **show route forwarding-table interface-name fe-0/1/1**

```
Name            Mtu Network       Address                 Ipkts Ierr     Opkts Oerr   Coll
fe-0/1/1       1514 <Link>      00.05.85.88.cc.20             0    0         0    0      0
```

### show route forwarding-table interface-name all

`user@host>` **show route forwarding-table interface-name all**

```
Name            Mtu Network       Address                 Ipkts Ierr     Opkts Oerr   Coll
fxp0           1514 <Link>      00.a0.a5.56.03.83      180965    0     39907    0      0
 unit 0        1500 192.168.187.0/ 192.168.187.10
fxp1           1514 <Link>      02.00.00.00.00.04    33010497    0 30110800     0      0
 unit 0        1500 10.0.0.0/8      10.0.0.1
                 10.0.0.0/8      10.0.0.4
```

```
               128.0.0.0/2    128.0.0.1
               128.0.0.0/2    128.0.0.4
             1500 fe80::/64      fe80::200:ff:fe0
               fec0::/64      fec0::a:0:0:4
             1500                4
lsi          1496 <Link>
dsc           max <Link>                              0     0         0     0     0
lo0           max <Link>                           8980     0      8980     0     0
 unit 0       max 127.0.0.1/8    127.0.0.1
               192.168.0.1/8  192.168.0.1
 unit 16384   max 127.0.0.1/8    127.0.0.1
 unit 16385   max
gre           max <Link>
ipip          max <Link>
tap           max <Link>
pime          max <Link>
pimd          max <Link>
mtun          max <Link>
so-0/0/0     4474 <Link>                        1679900     0  1068611     0     0
 unit 0      4470 <PtoP>         10.0.60.2              0     0         0     0
0
so-0/0/1     4474 <Link>                              0     0         0     0     0
 unit 0      4470 <PtoP>         10.0.80.2              0     0         0     0
0
so-0/0/2     4474 <Link>                              0     0         0     0     0
so-0/0/3     4474 <Link>                              0     0         0     0     0
fe-0/1/0     1514 <Link>     00.05.85.88.cc.1f     523120     0   623044     0     0
 unit 0      1500 10.0.90.12/30  10.0.90.14             0     0         0     0
0
fe-0/1/1     1514 <Link>     00.05.85.88.cc.20        0     0         0     0     0
fe-0/1/2     1514 <Link>     00.05.85.88.cc.21        0     0         0     0     0

...
```

**show route forwarding-table interface-name all detail**

`user@host>` **show route forwarding-table interface-name all detail**

```
Name         Mtu AFam    Network      Address            Ipkts Ierr    Opkts Oerr
  Coll
fxp0         1514        <Link>     00.a0.a5.56.03.83    181005    0    39948
0     0
 unit 0      1500 inet    192.168.187.0/ 192.168.187.10
fxp1         1514        <Link>     02.00.00.00.00.04  33012676    0 30112468
```

```
0     0
 unit 0     1500 inet    10.0.0.0/8     10.0.0.1
                      10.0.0.0/8     10.0.0.4
                      128.0.0.0/2    128.0.0.1
                      128.0.0.0/2    128.0.0.4
            1500 inet6   fe80::/64      fe80::200:ff:fe0
                      fec0::/64      fec0::a:0:0:4
            1500 tnp                    4
lsi         1496         <Link>
dsc          max         <Link>                             0    0        0
0     0
lo0          max         <Link>                          8980    0     8980
0     0
 unit 0      max inet    127.0.0.1/8    127.0.0.1
                      192.168.0.1/8  192.168.0.1
 unit 16384  max inet    127.0.0.1/8    127.0.0.1
 unit 16385  max inet
gre          max         <Link>
ipip         max         <Link>
tap          max         <Link>
pime         max         <Link>
pimd         max         <Link>
mtun         max         <Link>
so-0/0/0    4474         <Link>                       1679980    0  1068661
0     0
 unit 0     4470 inet    <PtoP>         10.0.60.2              0    0        0
  0     0

...
```

# 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



### 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 569.

#### 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

**List of Syntax**



**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

RELATED DOCUMENTATION

show route active-path | 455

**List of Sample Output**

show route inactive-path detail on page 574
show route inactive-path extensive on page 575
show route inactive-path terse on page 575

### 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 574.

### 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 inactive-prefix

### List of Syntax

Syntax on page 577
Syntax (EX Series Switches) on page 577

### Syntax

```
show route inactive-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show route inactive-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 inactive route destinations in each routing table.

### Options

**none**—Display all inactive route destination.

**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-prefix on page 578**
**show route inactive-prefix detail on page 578**
**show route inactive-prefix extensive on page 578**
**show route inactive-prefix terse on page 578**

### 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-prefix

user@host> **show route inactive-prefix**

```
inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 00:04:54
                    > via lo0.0
```

### show route inactive-prefix detail

user@host> **show route inactive-prefix detail**

```
inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
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>
                Age: 4:51
                Task: IF
                AS path: I00:04:54
                    > via lo0.0
```

### show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-prefix detail on page 578.

### show route inactive-prefix terse

user@host> **show route inactive-prefix terse**

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = 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
```

# show route instance

### List of Syntax

Syntax on page 580
Syntax (EX Series Switches and QFX Series) on page 580

### 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.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX 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

## RELATED DOCUMENTATION

*Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling*

*Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

**List of Sample Output**



**Output Fields**

Table 31 on page 581 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

**Table 31: show route instance Output Fields**

| Field Name | Field Description | Level of Output |
|---|---|---|
| Instance or *instance-name* | Name of the routing instance. | All levels |
| 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 |

Table 31: show route instance Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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** |
| Vrf-edge-protection-id | Context identifier configured for edge-protection. | **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:<br>• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.<br>• **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 ]
  Vrf-edge-protection-id:  166.1.3.1  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

**List of Syntax**

Syntax on page 590
Syntax (EX Series Switches) on page 590

**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 591**

**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 localization

### Syntax

```
show route localization
```

### Release Information

Command introduced in Junos OS Release 11.4 for T-Series routers.
Command introduced in Junos OS Release 12.3 for MX Series routers.

### Description

(T320, T640, and T1600 routers only) Display route localization details.

### Options

**detail**—Display detailed output.

### Required Privilege Level

view

RELATED DOCUMENTATION

*Example: Configuring Packet Forwarding Engine FIB Localization*

### Output Fields

Table 32 on page 592 lists the output fields for the **show route localization** command. Output fields are listed in the approximate order in which they appear.

Table 32: show route localization Output Fields

| Field Name | Field Description |
|---|---|
| **FIB-local** | FPCs configured as FIB-local. |
| **FIB-remote** | FPCs configured as FIB-remote. |
| **Normal** | FPCs neither configured as **FIB-local** or **FIB-remote**. |
| **Protocols** | IPv4 (inet) or IPv6 (inet6) traffic configured for route localization. |

## Sample Output

`user@R0>` **show route localization**

```
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
```

`user@R0>` **show route localization detail**

```
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:  inet, inet6
  FIB-local:  FPC2
  FIB-remote: FPC0, FPC1
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5
  FPC3: 6
  FPC4: 8
  FPC5: 11
  FPC6: 13
  FPC7: 15
```

# show route martians

### List of Syntax

Syntax on page 594
Syntax (EX Series Switches) on page 594

### 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: Removing the Class E Prefix on Martian Addresses | 159
Understanding Martian Addresses | 157

**List of Sample Output**

show route martians on page 595

**Output Fields**

Table 33 on page 595 lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

**Table 33: 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

### List of Syntax

Syntax on page 597
Syntax (EX Series Switches) on page 597

### 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 598**
**show route next-hop detail on page 598**
**show route next-hop extensive on page 601**
**show route next-hop terse on page 603**

### 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

**List of Syntax**

Syntax on page 605
Syntax (EX Series Switches) on page 605

### 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.
**ospf2** and **ospf3** options introduced in Junos OS Release 9.2.
**ospf2** 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**



**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 detail (Labeled Unicast)**

user@host> **show route protocol bgp 1.1.1.8/32 detail**

```
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
1.1.1.8/32 (2 entries, 2 announced)
State:
*BGP Preference: 1/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xc007f30
Next-hop reference count: 2
Source: 1.1.1.1
Next hop type: Router, Next hop index: 614
Next hop: 20.1.1.2 via ge-0/0/1.0, selected
Label-switched-path lsp1
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123, Push 
299872(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None; Label 299872: None;
Label element ptr: 0xc007860
Label parent element ptr: 0xc0089a0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x140
Protocol next hop: 1.1.1.4
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None;
Indirect next hop: 0xae8d300 1048576 INH Session ID: 0x142
State:
Local AS: 5 Peer AS: 5
Age: 22:43 Metric2: 2
Validation State: unverified
Task: BGP_5.1.1.1.1
Announcement bits (2): 0-KRT 7-Resolve tree 2
AS path: I
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Localpref: 100
Router ID: 1.1.1.1
```

**show route protocol bgp detail (Aggregate Extended Community Bandwidth)**

`user@host>` **show route 10.0.2.0 protocol bgp detail**

```
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb618990
                Next-hop reference count: 3
                Source: 10.0.1.1
                Next hop: 10.0.0.2 via ge-0/0/0.0 balance 40%
                Session Id: 0x0
                Next hop: 10.0.1.1 via ge-0/0/1.0 balance 60%, selected
                Session Id: 0x0
                State: <Active Ext>
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
                Validation State: unverified
                Task: BGP_65001.10.0.1.1
                Announcement bits (3): 0-KRT 2-BGP_Listen.0.0.0.0+179
3-BGP_RT_Background
                AS path: 65001 I
                Communities: bandwidth:65000:60000000
          Accepted Multipath
                Localpref: 100
                Router ID: 128.49.121.137
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 595
                Address: 0xb7a1330
                Next-hop reference count: 9
                Source: 10.0.0.2
                Next hop: 10.0.0.2 via ge-0/0/0.0, selected
                Session Id: 0x141
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
                Validation State: unverified
                Task: BGP_65001.10.0.0.2
                AS path: 65001 I
                Communities: bandwidth:65000:40000000
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 128.49.121.132
```

**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            >172.16.100.1    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

172.16.8.0/24          *[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
172.16.30.0/24      *[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

2001:db8::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-0/0/0.0.0      /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: 64500
                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: 64500
                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: 64500
                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: 64500
                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: 64500
                  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
172.16.233.5/32       *[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
172.16.233.5/32       *[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
172.16.233.9/32       *[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 range

### List of Syntax

Syntax on page 623
Syntax (EX Series Switches) on page 623

### Syntax

```
show route range
<brief | detail | extensive | terse>
<destination-prefix>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show route range
<brief | detail | extensive | terse>
<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 routing table entries using a prefix range.

### Options

**none**—Display standard information about all routing table entries using a prefix range.

**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***—Destination and prefix mask for the range.

**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 range on page 624**
**show route range destination-prefix on page 624**
**show route range detail on page 625**

show route range extensive on page 626
show route range terse on page 627

### 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 range

user@host> **show route range**

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.255.71.14/32    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

### show route range destination-prefix

user@host> **show route range 192.168.0.0/16**

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.0.0/16     *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

**show route range detail**

user@host> **show route range detail**

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (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>
                Age: 30:05
                Task: IF
```

```
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

**show route range extensive**

`user@host>` **show route range extensive**

```

inet.0: 11 destinations, 11 routes (10 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>
                Age: 30:17
                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>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

```
10.255.71.14/32 (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>
                Age: 30:17
                Task: IF
                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>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

**show route range terse**

user@host> **show route range terse**

```

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = 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
* 10.255.71.14/32    D   0                       >lo0.0
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.71.14/32   L   0                        Local
* 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

```

```
__juniper_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
* 10.0.0.4/32        L   0                        Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                    D   0                       >lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  abcd::10:255:71:14/128
*                    D   0                       >lo0.0
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.0

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.16385
```

# show route resolution

### List of Syntax

Syntax on page 629
Syntax (EX Series Switches) on page 629

### 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 631
show route resolution (Multipath Resolution) on page 632
show route resolution summary on page 633
show route resolution unresolved on page 633

### Output Fields

Table 34 on page 630 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

**Table 34: 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. |

Table 34: show route resolution Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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.<br><br>Merged—Merged next hops when recursive resolution of multipath is configured. |

## 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 (Multipath Resolution)

user@host> **show route resolution detail**

```
user@host> show route resolution detail 10.1.1.2
Tree Index: 1, Nodes 36, Reference Count 3
Contributing routing tables: inet.0 inet.3
Policy: [ abc ]
10.1.1.2/32 Originating RIB: inet.0
  Node path count: 1
  Next hop subtype: INDIRECT
  Indirect next hops: 2
        Protocol next hop: 10.1.1.1
        Inode flags: 0x206 path flags: 0x08
        Path fnh link: 0xc9321c0 path inh link: 0x0
        Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 12.1.1.2 via ge-2/0/1.0
                Session Id: 0x144
                Next hop: 13.1.1.2 via ge-2/0/2.0
                Session Id: 0x145

 10.1.1.1/32 Originating RIB: inet.0
   Node path count: 1
  Node flags: 1
   Forwarding nexthops: 1 (Merged)
  Nexthop: 12.1.1.2 via ge-2/0/1.0

  Nexthop: 13.1.1.2 via ge-2/0/2.0
```

user@host> **show route resolution summary**

```
Tree Index: 1, Nodes 7, Reference Count 2
Contributing routing tables: inet.3
Tree Index: 2, Nodes 7, Reference Count 8213
Contributing routing tables: inet.3
Policy: [ RRwM ]
Tree Index: 3, Nodes 7, Reference Count 2
```

```
Contributing routing tables: inet6.3
Tree Index: 4, Nodes 1, Reference Count 1
Contributing routing tables: iso.0
Tree Index: 5, Nodes 1000061, Reference Count 13
Contributing routing tables: inet.0 inet.3
Policy: [ Community-RRwM ]
Tree Index: 6, Nodes 2013, Reference Count 6
Contributing routing tables: inet6.0 inet6.3
Policy: [ RRwM ]
Tree Index: 7, Nodes 7, Reference Count 1501
Contributing routing tables: inet6.3
Policy: [ RRwM ]
Tree Index: 8, Nodes 1000061, Reference Count 2
Contributing routing tables: inet.0 inet.3
Policy: [ RRwM ]
```

### 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 snooping

## Syntax

```
show route snooping
<brief | detail | extensive | terse>
<all>
<best address/prefix>
<exact address>
<logical-system logical-system-name>
<range prefix-range>
<summary>
<table table-name>
```

## Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 9.0 for EX Series switches.

## Description

Display the entries in the routing table that were learned from snooping.

## Options

**none**—Display the entries in the routing table that were learned from snooping.

**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**.

**all**—(Optional) Display all entries, including hidden entries.

**best *address/prefix***—(Optional) Display the longest match for the provided address and optional prefix.

**exact *address/prefix***—(Optional) Display exact matches for the provided address and optional prefix.

**logical-system *logical-system-name***—(Optional) Display information about a particular logical system, or type 'all'.

**range *prefix-range***—(Optional) Display information for the provided address range.

**summary**—(Optional) Display route snooping summary statisitics.

**table *table-name***—(Optional) Display information for the named table.

## Required Privilege Level

view

## List of Sample Output

show route snooping detail on page 635

show route snooping logical-system all on page 636

### 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 snooping detail

user@host> **show route snooping detail**

```
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4), Next hop index: 1048584
                Next-hop reference count: 4
                State: <Active Int>
                Age: 2:24
                Task: MC
                Announcement bits (1): 0-KRT
```

```
                 AS path: I

<snip>
```

**show route snooping logical-system all**

`user@host>` **show route snooping logical-system all**

```

logical-system: default

inet.1: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Unsupported
+ = Active Route, - = Last Active, * = Both

0.0,0.1,0.0,232.1.1.65,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.66,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.67,100.1.1.2/112*[Multicast/180] 00:07:36

<snip>

default-switch.inet.1: 237 dest, 237 rts (237 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.15,0.1,0.0,0.0.0.0,0.0.0.0,2/120*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite
0.15,0.1,0.0,0.0.0.0,0.0.0.0,2,17/128*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite

<snip>

```

# show route source-gateway

### List of Syntax

Syntax on page 637
Syntax (EX Series Switches) on page 637

### Syntax

```
show route source-gateway address
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches)

```
show route source-gateway address
<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 were learned from a particular address. The **Source** field in the **show route detail** command output lists the source for each route, if known.

### 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**.

***address***—IP address of the system.

**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 source-gateway on page 638**
**show route source-gateway detail on page 639**
**show route source-gateway extensive on page 641**

### 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 source-gateway**

**user@host> show route source-gateway 10.255.70.103**

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

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)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
```

```
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3
```

**show route source-gateway detail**

user@host> **show route source-gateway 10.255.70.103 detail**

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

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)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
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: 12:14:00   Metric2: 1
                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

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                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: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 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: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                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
                Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
```

### show route source-gateway extensive

`user@host>` **show route source-gateway 10.255.70.103 extensive**

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
Restart Complete

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)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
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: 12:15:24   Metric2: 1
                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

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                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: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 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: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                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
                Secondary Tables: green.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                         Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0
```

```
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0
```

# show route summary

**List of Syntax**

Syntax on page 645
Syntax (EX Series Switches) on page 645

**Syntax**

```
show route summary
<logical-system (all | logical-system-name)>
<table routing-table-name>
<rib-sharding (main | rib-shard-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.
**rib-sharding** option introduced for cRPD Release 20.1R1.

**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.

**rib-sharding (main | *rib-shard-name*)**—(Optional) Display name of the rib shard.

**Required Privilege Level**

view

**List of Sample Output**

show route summary on page 648
show route summary table (with rib-sharding configured) on page 649
show route summary table on page 650
show route summary table (with Route Limits Configured for the Routing Table) on page 651
show route summary rib-sharding on page 651

**Output Fields**

Table 35 on page 646 lists the output fields for the **show route summary** command. Output fields are listed in the approximate order in which they appear.

**Table 35: 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 example:<br>• **inet.0**— List all routes for IPv4.<br>• **junos-main::inet.0**—Lists all active routes for Ipv4.<br>• **junos-bgpshard0::inet.0**—Lists IPv4 routes present in bgpshard0 thread.<br>• **__raass__inet.inet.0**—Lists IPv4 routes stored in RaaS Server that are received from the clients.<br>• **__raasc__inet.inet.0**—Lists IPv4 routes stored in RaaS Clients that are used for local resolution.<br>• **inet.6**— List all routes for IPv6.<br>• **junos-main::inet6.0**— Lists all active routes for IPv6.<br>• **junos-bgpshard0::inet6.0**—Lists IPv6 routes present in bgpshard0 thread.<br>• **__raass__inet.inet6.0**—Lists IPv6 routes stored in RaaS Server that are received from the clients.<br>• **__raasc__inet.inet6.0**—Lists IPv6 routes stored in RaaS Clients that are used for local resolution. |
| **destinations** | Number of destinations for which there are routes in the routing table. |

Table 35: show route summary Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **routes** | Number of routes in the routing table:<br>• **active**—Number of routes that are active.<br>• **holddown**—Number of routes that are in the hold-down state before being declared inactive.<br>• **hidden**—Number of routes that are not used because of routing policy. |
| Restart complete | All protocols have restarted for this routing table.<br>Restart state:<br>• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.<br>• **Complete**—All protocols have restarted for this routing table.<br>For example, if the output shows-<br>• `LDP.inet.0: 5 routes (4 active, 1 holddown, 0 hidden)`<br>`    Restart Pending: OSPF LDP VPN`<br>This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for **LDP.inet.0** routing table.<br>• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`<br>`Restart Complete`<br>This indicates that all protocols have restarted for **vpls_1.l2vpn.0** routing table. |
| **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.<br>• **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.<br>• **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. |

Table 35: show route summary Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| **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 (with rib-sharding configured)**

The following command shows the route summary information in junos-bgpshard0 thread only:

`user@host>` **show route summary rib-sharding junos-bgpshard0**

```
Autonomous system number: 1
Router ID: 2.2.2.1

inet.0: 774078 destinations, 1547602 routes (774078 active, 0 holddown, 0 hidden)

              Direct:    108 routes,    108 active
               Local:      7 routes,      7 active
                OSPF:    408 routes,    408 active
                 BGP: 1547048 routes, 773524 active
              Static:     31 routes,     31 active

junos-main::inet.0: 774078 destinations, 774078 routes (774078 active, 0 holddown,
 0 hidden)
              Direct:    108 routes,    108 active
               Local:      7 routes,      7 active
                OSPF:    408 routes,    408 active
                 BGP: 773524 routes, 773524 active
              Static:     31 routes,     31 active

junos-bgpshard0::inet.0: 258448 destinations, 516635 routes (258448 active, 0
holddown, 0 hidden)
```

```
                 BGP: 516374 routes, 258187 active

junos-bgpshard1::inet.0: 258129 destinations, 515997 routes (258129 active, 0
holddown, 0 hidden)
                 BGP: 515736 routes, 257868 active

junos-bgpshard2::inet.0: 257731 destinations, 515200 routes (257731 active, 0
holddown, 0 hidden)
                 BGP: 514938 routes, 257469 active

__raass__inet.inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
                RaaS:      1 routes,      1 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
              Direct:      2 routes,      2 active
               INET6:      1 routes,      1 active

junos-main::inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
              Direct:      2 routes,      2 active
               INET6:      1 routes,      1 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 summary rib-sharding

user@host> **show route summary rib-sharding junos-bgpshard14**

```
Autonomous system number: 100
Router ID: 20.255.255.10

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
              Direct:     29 routes,     29 active
               Local:     25 routes,     25 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
               Local:     14 routes,     14 active
```

# show route table

### List of Syntax

Syntax on page 652
Syntax (EX Series Switches, QFX Series Switches) on page 652

### Syntax

```
show route table routing-table-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

### Syntax (EX Series Switches, QFX 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.
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX 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 names begin 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 | 645

### List of Sample Output

### Output Fields

Table 20 on page 423 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

**Table 36: show route table Output Fields**

| Field Name | Field Description |
|---|---|
| *routing-table-name* | Name of the routing table (for example, inet.0). |
| Restart complete | All protocols have restarted for this routing table.<br><br>Restart state:<br><br>• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.<br>• **Complete**—All protocols have restarted for this routing table.<br><br>For example, if the output shows-<br><br>• `LDP.inet.0              : 5 routes (4 active, 1 holddown, 0 hidden)`<br><br>`    Restart Pending: OSPF LDP VPN`<br><br>This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for the **LDP.inet.0** routing table.<br>• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`<br>`Restart Complete`<br><br>This indicates that all protocols have restarted for the **vpls_1.l2vpn.0** routing table. |
| *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:<br><br>• **active** (routes that are active)<br>• **holddown** (routes that are in the pending state before being declared inactive)<br>• **hidden** (routes that are not used because of a routing policy) |

**Table 36: show route table 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:<br>• ***MPLS-label*** (for example, 80001).<br>• ***interface-name*** (for example, ge-1/0/2).<br>• ***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).<br>  ◦ ***neighbor-address***—Address of the neighbor.<br>  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.<br>  ◦ ***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.<br>  ◦ ***vc-id***—Virtual circuit identifier.<br>  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.<br>• ***inclusive multicast Ethernet tag route***—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):<br>  ◦ *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.<br>  ◦ *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.<br>  ◦ *IP address length*—(1 octet) Length of IP address in bits.<br>  ◦ *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address. |
| 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.<br>• **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).<br>• 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). |

Table 36: show route table Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| [*protocol*, *preference*] | Protocol from which the route was learned and the preference value for the route.<br>• **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br>• **-** —A hyphen indicates the last active route.<br>• *—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.<br>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. |
| PMSI | Provider multicast service interface (MVPN routing table). |
| Next-hop type | Type of next hop. For a description of possible values for this field, see Table 22 on page 483. |
| 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. |

**Table 36: show route table 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 name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information:<br>• 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.<br>• 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 23 on page 485. |
| Local AS | AS number of the local routing devices. |
| 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. |
| MED-plus-IGP | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added. |

**Table 36: show route table Output Fields *(continued)***

| Field Name | Field Description |
|---|---|
| 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. |
| Task | Name of the protocol that has added the route. |
| Announcement bits | The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table.<br>• ***n***—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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **Recorded**—The AS path is recorded by the sample process (sampled).<br>• **?**—Incomplete; typically, the AS path was aggregated.<br>When AS path numbers are included in the route, the format is as follows:<br>• **[ ]**—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.<br>• **[ ]**—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.<br>• **{ }**—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.<br>• **( )**—Parentheses enclose a confederation.<br>• **( [ ] )**—Parentheses and brackets enclose a confederation set.<br>**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. |

**Table 36: show route table Output Fields *(continued)***

| Field Name | Field Description |
|---|---|
| validation-state | (BGP-learned routes) Validation status of the route:<br>• **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.<br>• **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br>• **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.<br>• **Valid**—Indicates that the prefix and autonomous system pair are found in the database. |
| FECs bound to route | Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured. |
| Primary Upstream | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops | When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks. |
| Label | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label. |
| weight | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. |
| 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. |

Table 36: show route table Output Fields *(continued)*

| Field Name | Field Description |
|---|---|
| Communities | Community path attribute for the route. See Table 24 on page 488 for all possible values for this field. |
| 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 LongLivedStale | The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. |
| Accepted LongLivedStaleImport | The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.<br><br>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table |
| ImportAccepted LongLivedStaleImport | Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table<br><br>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. |
| 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. |

**Table 36: show route table Output Fields *(continued)***

| Field Name | Field Description |
|---|---|
| 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 22 on page 483 describes all possible values for the Next-hop Types output field.

**Table 37: 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. |
| 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. |

Table 37: Next-hop Types Output Field Values *(continued)*

| Next-Hop Type | Description |
|---|---|
| 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.<br>To qualify as a next-hop type router, the route must meet the following criteria:<br>• Must not be a direct or local subnet for the routing device.<br>• 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 23 on page 485 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 38: 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. |

Table 38: State Output Field Values *(continued)*

| Value | Description |
| --- | --- |
| 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). |
| MartianOK | Route exempt from martian filtering. |

**Table 38: State Output Field Values *(continued)***

| Value | Description |
|---|---|
| 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:<br>• The route is damped.<br>• The route is rejected by an import policy.<br>• The route is unresolved. |
| Update source | Last tiebreaker is the lowest IP address value. |

Table 24 on page 488 describes the possible values for the Communities output field.

**Table 39: 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. |

Table 39: Communities Output Field Values *(continued)*

| Value | Description |
|---|---|
| **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. |
| evpn-mcast-flags | Identifies the value in the multicast flags extended community and whether snooping is enabled. A value of 0x1 indicates that the route supports IGMP proxy. |
| evpn-l2-info | Identifies whether Multihomed Proxy MAC and IP Address Route Advertisement is enabled. A value of 0x20 indicates that the proxy bit is set. .<br><br>Use the **show bridge mac-ip-table extensive** statement to determine whether the MAC and IP address route was learned locally or from a PE device. |

## Sample Output

### show route table bgp.l2vpn.0

user@host> **show route table bgp.l2vpn.0**

```
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)

10.255.245.12:1:172.16.4.0/8 (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 Target Filtering Is Configured)

user@host> **show route table bgp.rtarget.0**

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

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 bgp.evpn.0

user@host> **show route table bgp.evpn.0**

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

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```

```
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

**show route table evpna.evpn.0**

`user@host>` **show route table evpna.evpn.0**

```
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
                   *[EVPN/170] 01:37:09
                      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[EVPN/170] 01:37:12
                      Indirect
```

**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 172.16.5.254 via fxp0.0
10.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
10.0.0.2/32         *[Local/0] 00:51:58
                     Local
10.12.12.21/32     *[Local/0] 00:51:57
                     Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                     Local
10.13.13.21/32     *[Local/0] 00:51:58
```

```
                     Local
10.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
10.222.5.0/24      *[Direct/0] 00:51:58
                    > via fxp0.0
10.222.5.81/32     *[Local/0] 00:51:58
                     Local
```

**show route table inet.3**

`user@host>` **show route table inet.3**

```
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 10.2.94.2 via lt-1/2/0.49
                    > to 10.2.3.2 via lt-1/2/0.23
```

**show route table inet.3 protocol ospf**

`user@host>` **show route table inet.3 protocol ospf**

```
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32         [L-OSPF/10] 1d 00:00:56, metric 2
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                      to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32         [L-OSPF/10] 1d 00:01:01, metric 3
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                      to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32         [L-OSPF/10] 1d 00:01:01, metric 4
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                      to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32         [L-OSPF/10] 1d 00:01:01, metric 5
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                      to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32         [L-OSPF/10] 1d 00:01:01, metric 6
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                      to 10.0.6.60 via lt-1/2/0.12, Pop
```

**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: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 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: 64502
                 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 inetflow.0 extensive (BGP Flowspec Redirect to IP)**

user@host> **show route table inetflow.0 extensive**

```
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
2.2.2.2,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xb209500 (adv_entry)
Advertised metrics:
Nexthop: 21.1.4.5
Localpref: 100
AS path: [100] I
Communities: redirect-to-ip:21.1.4.5:0
Action(s): accept,count
*Flow Preference: 5
Next hop type: Indirect, Next hop index: 0
Address: 0xa2b931c
Next-hop reference count: 1Next hop:
State: <Active> L
ocal AS: 69
Age: 2
Validation State: unverified
Task: RT Flow
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-ip:21.1.4.5:0
```

`user@host>` **show route table inetflow.0 extensive**

```
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
2.2.2.2,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xb209500 (adv_entry)
Advertised metrics:
Nexthop: 21.1.4.5
Localpref: 100
AS path: [100] I
Communities: redirect-to-nexthop
Action(s): accept,count
*Flow Preference: 5
Next hop type: Indirect, Next hop index: 0
Address: 0xa2b931c
Next-hop reference count: 1
Next hop:
State: <Active>
Local AS: 69
Age: 2
Validation State: unverified
Task: RT Flow
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-nexthop
regress@10.102.178.210> show route table inetflow.0 extensive
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
4.4.4.4,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Action(s): accept,count
*BGP Preference: 170/-101
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5e3c30
Next-hop reference count: 3
Next hop: 21.1.4.5
State: <Active Int Ext>
Local AS: 100 Peer AS: 100
Age: 10
Validation State: unverified
Task: BGP_100.1.1.1.1+179
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-nexthop
```

```
Accepted
Localpref: 100
Router ID: 1.1.1.1
```

**show route table lsdist.0 extensive**

`user@host>` **show route table lsdist.0 extensive**

```
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }/1152
(1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f378 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     Localpref: 100
     AS path: [4170512532] I
     Communities:
Path NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }
Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                IPv4 Router-ids:
                   128.220.11.197
                Area membership:
                   47 00 05 80 ff f8 00 00 00 01 08 00 01
               SPRING-Capabilities:                  - SRGB block [Start: 800000,
 Range: 256, Flags: 0xc0]
                SPRING-Algorithms:
                   - Algo: 0
 LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.3300.5067)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f3cc (adv_entry)
```

```
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.33000
    *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                Color: 32768
                Maximum bandwidth: 1000Mbps
                Reservable bandwidth: 1000Mbps
                Unreserved bandwidth by priority:
                   0   1000Mbps
                   1   1000Mbps
                   2   1000Mbps
                   3   1000Mbps
                   4   1000Mbps
                   5   1000Mbps
                   6   1000Mbps
                   7   1000Mbps
                 Metric: 10
                 TE Metric: 10
                 LAN IPV4 Adj-SID -  Label: 299776, Flags: 0x30,
Weight: 0, Nbr: 10.220.1.83

PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
IPv4:128.220.11.197/32 } ISIS-L1:0 }/1152 (1 entry, 1 announced) TSI: Page 0 idx
0, (group ibgp type Internal) Type 1 val 0xa62f43c (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
```

```
Path PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
 IPv4:128.220.11.197/32 } ISIS-L1:0 } Vector len 4.  Val: 0
         *IS-IS  Preference: 15
                 Level: 1
                 Next hop type: Fictitious, Next hop index: 0
                 Address: 0x95dfc64
                 Next-hop reference count: 9
                 State:<Active NotInstall>
                 Local AS: 4170512532
                 Age: 6:05
                 Validation State: unverified
                 Task: IS-IS
                 Announcement bits (1): 0-BGP_RT_Background
                 AS path: I
                               Prefix SID: 67, Flags: 0x40, Algo: 0
```

**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 lsdist.0**

`user@host>` **show route table lsdist.0**

```
lsdist.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:4.4.4.4 } Remote { AS:4
BGP-LS ID:100 IPv4:7.7.7.7 }.{ IPv4:7.7.7.7 } Undefined:0 }/1152
                   *[BGP-LS-EPE/170] 00:20:56
                      Fictitious
LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:4.4.4.4 IfIndex:339 }
Remote { AS:4 BGP-LS ID:100 IPv4:7.7.7.7 }.{ IPv4:7.7.7.7 } Undefined:0 }/1152

                   *[BGP-LS-EPE/170] 00:20:56
                      Fictitious
LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:50.1.1.1 } Remote { AS:4
BGP-LS ID:100 IPv4:5.5.5.5 }.{ IPv4:50.1.1.2 } Undefined:0 }/1152
                   *[BGP-LS-EPE/170] 00:20:56
                      Fictitious
```

**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: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 11:39:56, metric 1
                      to table inet.0
0(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
1                  *[MPLS/0] 11:39:56, metric 1
                      Receive
2                  *[MPLS/0] 11:39:56, metric 1
                      to table inet6.0
2(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
13                 *[MPLS/0] 11:39:56, metric 1
                      Receive
303168             *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                      to table pbbn10.evpn-mac.0
303184             *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
                      to table pbbn10.evpn-mac.0
                    [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                      to table pbbn10.evpn-mac.0
303264             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 1000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 2000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

```
303344             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```

#### show route table mpls.0 detail (PTX Series)

`user@host>` **show route table mpls.0 detail**

```
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
```

```
                 Age: 21         Metric2: 1
                 Validation State: unverified
                 Task: Common L2 VC
                 Announcement bits (2): 0-KRT 2-Common L2 VC
                 AS path: I
```

### show route table mpls.0 ccc ge-0/0/1.1004 detail

`user@host>` **show route table mpls.0 ccc ge-0/0/1.1004 detail**

```
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
ge-0/0/1.1004 (1 entry, 1 announced)
        *EVPN   Preference: 7
                Next hop type: List, Next hop index: 1048577
                Address: 0xdc14770
                Next-hop reference count: 3
                Next hop: ELNH Address 0xd011e30
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd011e30
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.1
                    Label operation: Push 301952
                    Composite next hop: 0xd011dc0 754 INH Session ID: 0x146
                    Indirect next hop: 0xb69a890 1048615 INH Session ID: 0x146
                        Next hop type: Router, Next hop index: 735
                        Address: 0xd00e530
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe1
                        Label operation: Push 300320
                        Label TTL action: prop-ttl
                        Load balance label: Label 300320: None;
                        Label element ptr: 0xd00e580
                        Label parent element ptr: 0x0
                        Label element references: 18
                        Label element child references: 16
                        Label element lsp id: 5
                Next hop: ELNH Address 0xd012070
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd012070
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.2
                    Label operation: Push 301888
                    Composite next hop: 0xd012000 755 INH Session ID: 0x143
```

```
            Indirect next hop: 0xb69a9a0 1048641 INH Session ID: 0x143
                Next hop type: Router, Next hop index: 716
                Address: 0xd00e710
                Next-hop reference count: 23
                Next hop: 100.46.1.2 via ge-0/0/5.0
                Label-switched-path pe4_to_pe2
                Label operation: Push 300304
                Label TTL action: prop-ttl
                Load balance label: Label 300304: None;
                Label element ptr: 0xd00e760
                Label parent element ptr: 0x0
                Label element references: 15
                Label element child references: 13
                Label element lsp id: 6
        Next hop: ELNH Address 0xd0121f0, selected
            Next hop type: Indirect, Next hop index: 0
            Address: 0xd0121f0
            Next-hop reference count: 3
            Protocol next hop: 100.100.100.3
            Label operation: Push 301984
            Composite next hop: 0xd012180 756 INH Session ID: 0x145
            Indirect next hop: 0xb69aab0 1048642 INH Session ID: 0x145
                Next hop type: Router, Next hop index: 801
                Address: 0xd010ed0
                Next-hop reference count: 32
                Next hop: 100.46.1.2 via ge-0/0/5.0
                Label-switched-path pe4_to_pe3
                Label operation: Push 300336
                Label TTL action: prop-ttl
                Load balance label: Label 300336: None;
                Label element ptr: 0xd0108c0
                Label parent element ptr: 0x0
                Label element references: 22
                Label element child references: 20
                Label element lsp id: 7
        State: < Active Int >
        Age: 2:06:50
        Validation State: unverified
        Task: evpn global task
        Announcement bits (1): 1-KRT
        AS path: I
```

**show route table mpls.0 protocol evpn**

`user@host>`**show route table mpls.0 protocol evpn**

```
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299872             *[EVPN/7] 02:30:58, routing-instance mhevpn, route-type
Ingress-IM, vlan-id 10
                      to table mhevpn.evpn-mac.0
300016           *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
 vlan-id 110
                      to table VS-1.evpn-mac.0
300032           *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
 vlan-id 120
                      to table VS-1.evpn-mac.0
300048           *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
 vlan-id 130
                      to table VS-1.evpn-mac.0
300064           *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
 vlan-id 210
                      to table VS-2.evpn-mac.0
300080           *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
 vlan-id 220
                      to table VS-2.evpn-mac.0
300096           *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
 vlan-id 230
                      to table VS-2.evpn-mac.0
300112             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300128             *[EVPN/7] 02:29:22, routing-instance mhevpn, route-type
Ingress-Aliasing
                      to table mhevpn.evpn-mac.0
300144           *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
 ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300160             *[EVPN/7] 02:29:22, routing-instance VS-1, route-type
Ingress-Aliasing
                      to table VS-1.evpn-mac.0
300176           *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
 ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300192             *[EVPN/7] 02:29:22, routing-instance VS-2, route-type
Ingress-Aliasing
                      to table VS-2.evpn-mac.0
300208             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-IM, vlan-id 120
```

```
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300224             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300240             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300256             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300272             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300288             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300304             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance 
VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300320           *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
 ESI 00:11:11:11:11:11:11:11:11:11
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300336           *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
 ESI 00:33:33:33:33:33:33:33:33:33
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300368           *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
 ESI 00:33:33:33:33:33:33:33:33:33
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300384           *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
 ESI 00:11:11:11:11:11:11:11:11:11
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300416             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
```

```
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300432             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300480             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300496             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300560             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300592             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300608             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001, Pop
300624             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001, Pop
301232             *[EVPN/7] 02:29:17
                    > via ge-0/0/1.1002, Pop
301296             *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003, Pop
301312             *[EVPN/7] 02:27:06
                    > via ae10.2003, Pop
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301360             *[EVPN/7] 02:29:01
                    > via ge-0/0/1.1004, Pop
301408             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301456             *[EVPN/7] 02:27:06
                    > via ae10.1010, Pop
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301552           *[EVPN/7] 02:27:07, routing-instance VS-1, route-type Egress-MAC,
 ESI 00:22:22:22:22:22:22:22:22:22
```

```
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301568           *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
 ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301648             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301664             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301680             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301696             *[EVPN/7] 02:27:07, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301712             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301728             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301744             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301760             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301776             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301792             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301808             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301824             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301840             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1002, route-type Egress, vlan-id 2002
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301856             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
```

```
vpws1003, route-type Egress, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301872             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1003, route-type Egress Protection, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301888             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1010, route-type Egress Protection, vlan-id 1010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301904             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301920             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301936             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301952             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301968             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301984             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302000             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302016             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302032             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302048             *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302064             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302080             *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

```
302096            *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302112            *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302128            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302144            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302160            *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302176            *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302192            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302208            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302224            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302240            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302256            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302272            *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302288            *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302304            *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302320            *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
mhevpn, route-type Egress-MAC
```

```
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302336             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance 
mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302352             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance 
vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302368             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance 
mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302384             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance 
mhevpn, route-type Egress-SH, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302400             *[EVPN/7] 02:26:21
                    > via ge-0/0/1.3001, Pop
302432             *[EVPN/7] 02:26:21, remote-pe 100.100.100.3, routing-instance 
vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302448             *[EVPN/7] 02:26:21, remote-pe 100.100.100.1, routing-instance 
vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302464             *[EVPN/7] 02:26:20, remote-pe 100.100.100.2, routing-instance 
vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302480             *[EVPN/7] 02:26:14
                    > via ge-0/0/1.3016, Pop
302512             *[EVPN/7] 02:26:14, remote-pe 100.100.100.1, routing-instance 
vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302528             *[EVPN/7] 02:26:14, remote-pe 100.100.100.2, routing-instance 
vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302560             *[EVPN/7] 02:26:06
                    > via ae10.3011, Pop
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302592             *[EVPN/7] 02:26:07, remote-pe 100.100.100.1, routing-instance 
vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302608             *[EVPN/7] 02:26:07, remote-pe 100.100.100.2, routing-instance 
vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302624             *[EVPN/7] 02:26:07, remote-pe 100.100.100.3, routing-instance 
vpws3011, route-type Egress Protection, vlan-id 301100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

```
302656              *[EVPN/7] 02:25:59
                     > via ae10.3006, Pop
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302688              *[EVPN/7] 02:26:00, remote-pe 100.100.100.2, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302704              *[EVPN/7] 02:26:00, remote-pe 100.100.100.1, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302720              *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302736              *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
vpws3006, route-type Egress Protection, vlan-id 300600
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.1001       *[EVPN/7] 02:29:23
                     > via ge-0/0/1.2001
ge-0/0/1.2001       *[EVPN/7] 02:29:23
                     > via ge-0/0/1.1001
ge-0/0/1.1002       *[EVPN/7] 02:27:06
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.2003           *[EVPN/7] 02:29:10
                     > via ge-0/0/1.1003
ge-0/0/1.1003       *[EVPN/7] 02:27:06
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

                     > via ae10.2003
ge-0/0/1.1004       *[EVPN/7] 02:27:06
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.1010           *[EVPN/7] 02:27:06
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ge-0/0/1.3001       *[EVPN/7] 02:26:20
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.3016       *[EVPN/7] 02:26:13
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3011           *[EVPN/7] 02:26:06
```

```
                     > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3006            *[EVPN/7] 02:25:59
                     > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                       to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                       to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

### show route table mpls.0 protocol ospf

`user@host>` **show route table mpls.0 protocol ospf**

```
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952              *[L-OSPF/10] 23:59:42, metric 0
                     > to 10.0.10.70 via lt-1/2/0.14, Pop
                       to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)         *[L-OSPF/10] 23:59:42, metric 0
                     > to 10.0.10.70 via lt-1/2/0.14, Pop
                       to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968              *[L-OSPF/10] 23:59:48, metric 0
                     > to 10.0.6.60 via lt-1/2/0.12, Pop
```

### show route table mpls.0 extensive (PTX Series)

`user@host>` **show route table mpls.0 extensive**

```
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
```

```
                  Label TTL action: no-prop-ttl
                  Load balance label: Label 299872:Flow label PUSH;
                  Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                  Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                  State: <Active Int>
                  Age: 47         Metric2: 1
                  Validation State: unverified
                  Task: Common L2 VC
                  Announcement bits (2): 0-KRT 2-Common L2 VC
                  AS path: I
                  Composite next hops: 1
                          Protocol next hop: 10.255.255.1 Metric: 1
                          Label operation: Push 299872 Offset: 252
                          Label TTL action: no-prop-ttl
                          Load balance label: Label 299872:Flow label PUSH;
                          Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                          Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                          Indirect path forwarding next hops: 1
                                  Next hop type: Router
                                  Next hop: 10.0.0.1 via ge-0/0/1.0
                                  Session Id: 0x1
                          10.255.255.1/32 Originating RIB: inet.3
                            Metric: 1                       Node path count: 1
                            Forwarding nexthops: 1
                                  Nexthop: 10.0.0.1 via ge-0/0/1.0
```

**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 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
```

```
300352(S=0)        *[RSVP/7/1] 00:08:00, metric 1
                    > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384             *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)        *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 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

172.16.1.11:1000:1:1/96 (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

1:10.255.2.202:65536:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65536:128:::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
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

**show route table vrf1.mvpn.0 extensive**

`user@host>` **show route table vrf1.mvpn.0 extensive**

```
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
        *MVPN   Preference: 70
                PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
                Next hop type: Indirect
                Address: 0xbb2c944
                Next-hop reference count: 360
                Protocol next hop: 10.255.50.77
                Indirect next hop: 0x0 - INH Session ID: 0x0
                State: <Active Int Ext>
                Age: 53:03      Metric2: 1
                Validation State: unverified
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export
                AS path: I
```

**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: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 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: 64502
                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@host> **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

10.1.1.2:100:10.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:10.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
```

```
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
                   *[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)
10.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 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: 10.2.0.0
                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_172.16.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 10.2.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 10.2.0.0 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: 10.3.0.0
                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: 10.3.0.0
                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_172.16.3.0.0+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: 10.3.0.0
                  Primary Routing Table bgp.l3vpn.0
                  Composite next hops: 1
                          Protocol next hop: 10.3.0.0 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
                          10.3.0.0/32 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: 10.2.0.0
                  Push 16
                  Composite next hop: 0x25805988 - INH Session ID: 0x193c
                  Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                  Protocol next hop: 10.3.0.0
                  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 table bgp.evpn.0 extensive | no-more (EVPN)

`user@host>` **show route table bgp.evpn.0 extensive | no-more**

```
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17  Peer AS:17  Age:21:12  Metric2:1  Validation State:
unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 100
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        1.2.3.4/32 Originating RIB: inet.0
```

```
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:19:43  Metric2:1  Validation
State:unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:2222:22 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 200
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
```

```
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:17:21  Metric2:1  Validation State:
 unverified  Task: BGP 17,1,2,3,4+50756
          AS path: I
                Communities: target:3333:33 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 300
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17 Age:37:01  Metric2:1  Validation State:
unverified  Task: BGP 17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
```

```
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS: 17 Peer AS: 17   Age:35:22  Metric2:1  Validation
State:unverified  Task: BGP 17.1.2.3.4+50756
           AS path:I  Communities: target:2222:22 encapsulation):0:0:0:0:3

Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
```

```
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17 Peer AS:  17  Age 35:22 Metric2:1  Validation State:
 unverified   Task: BGP 17.1.2.3.4+5075
          6  AS path: I   Communities: target:3333:33  encapsulation0:0:0:0:3
Import Accepted   Localpref:100
          Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0
```

### show route table default-switch.evpn.0 extensive

The following shows the partial output listing for the EVPN VNI table.

`user@host>` **show route table default-switch.evpn.0 extensive**

```
3:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.0.1:00
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect, Next hop index: 0
                Address: 0xcebfad0
                Next-hop reference count: 26
                Source: 10.255.0.1
                Protocol next hop: 10.255.0.1
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                State: <Secondary Active Int Ext>
              Local AS:   100 Peer AS:   100
                Age: 1:35:30    Metric2: 2
                Validation State: unverified
```

```
                Task: BGP_100.10.255.0.1
                Announcement bits (1): 0-default-switch-evpn
                AS path: I
                Communities: target:100:100 encapsulation:vxlan (0x8)
evpn-mcast-flags:0x1:snooping-enabled
               . . .
```

**show route table evpn1.evpn-mcsn**

The following shows the output listing for the multicast information used by the rpd and mcsnoopd.

`user@host>` **show route table default-switch.evpn-mcsn.1**

```
default-switch.evpn-mcsn.1: 9 destinations, 9 routes (9 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

0.14,0.0,0.0/48    *[Multicast/180] 00:01:02
                       to 1.1.1.1 via vtep.32770
                       to 1.2.2.2 via vtep.32771
                       to 1.6.6.6 via vtep.32769
                       to 1.3.3.3 via vtep.32772
0.14,0.0,0.0,224.0.0.0/52*[Multicast/180] 00:01:02
                       to 1.1.1.1 via vtep.32770
                       to 1.2.2.2 via vtep.32771
                       to 1.6.6.6 via vtep.32769
0.14,0.0,0.0,225.1.1.1/80*[Multicast/180] 00:00:06
                       to 1.1.1.1 via vtep.32770
                       to 1.2.2.2 via vtep.32771
                       to 1.6.6.6 via vtep.32769
                       to 1.3.3.3 via vtep.32772
```

**show route table evpn1 (Multihomed Proxy MAC and IP Address)**

The following shows a partial output listing for an EVPN instance. This indicates when Multihomed Proxy MAC and IP Address Route Advertisement is enabled.

`user@host>` **show route table evpn-1**

```
2:666:11010003::1002::00:00:00:00:00:02::102.1.1.2/304 MAC/IP (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group vteps type Internal) Type 1 val 0xb20eb10 (adv_entry)
   Advertised metrics:
```

```
          Nexthop: 103.1.1.1
          Localpref: 100
          AS path: [666] I
Communities: target:666:1002 evpn-l2-info:0x20:proxy (mtu 0)
Path 2:666:11010003::1002::00:00:00:00:00:02::102.1.1.2 Vector len 4.  Val: 0
                *EVPN   Preference: 170
                        Next hop type: Indirect, Next hop index: 0
                        Address: 0xc3a9cf0
                        Next-hop reference count: 36
                        Protocol next hop: 103.1.1.1
                        Indirect next hop: 0x0 - INH Session ID: 0x0
                        State: <Active Int Ext>
```

# show route terse

### List of Syntax

Syntax on page 708
Syntax (EX Series Switches) on page 708

### 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 711**

### Output Fields

Table 40 on page 709 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

**Table 40: 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:<br>• **active** (routes that are active)<br>• **holddown** (routes that are in the pending state before being declared inactive)<br>• **hidden** (routes that are not used because of a routing policy) |
| *route key* | Key for the state of the route:<br>• **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br>• **-** —A hyphen indicates the last active route.<br>• *****—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:<br>• **?**—Not evaluated. Indicates that the route was not learned through BGP.<br>• **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.<br>• **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br>• **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database. |
| Destination | Destination of the route. |

**Table 40: show route terse Output Fields *(continued)***

| Field Name | Field Description |
|---|---|
| P | Protocol through which the route was learned:<br>• **A**—Aggregate<br>• **B**—BGP<br>• **C**—CCC<br>• **D**—Direct<br>• **G**—GMPLS<br>• **I**—IS-IS<br>• **L**—L2CKT, L2VPN, LDP, Local<br>• **K**—Kernel<br>• **M**—MPLS, MSDP<br>• **O**—OSPF<br>• **P**—PIM<br>• **R**—RIP, RIPng<br>• **S**—Static<br>• **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. |
| 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:<br>• **I**—IGP.<br>• **E**—EGP.<br>• **?**—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
* ? 172.16.1.1/32         O  10          1            >10.0.0.2
  ?                    B 170        100                               I
  unverified                                       >10.0.0.2
* ? 172.16.1.1/32         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
* ? 172.16.233.5/32       O  10          1             MultiRecv
```