# Junos® OS

# Protocol-Independent Routing Properties User Guide

Published
2023-10-16

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

## 4　Route Aggregation

## 5　Martian Addresses

## 6　Packet Forwarding

## 7   Troubleshooting

## 8   Configuration Statements

9 **Operational Commands**

# About This Guide

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

# 1
**CHAPTER**

## Overview

# 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

# 2
**CHAPTER**

# Junos OS Routing Tables

# Configuring Junos OS Routing Tables

## 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 13.

  > **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.

## 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 1: 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 | – | – |

**Table 1: Routing Table Route Properties** *(Continued)*

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

**Table 1: Routing Table Route Properties** *(Continued)*

| Description | Static | Aggregate | Generated |
|---|---|---|---|
| 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 | – | – |

## 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 ];
    }
 }
```

**SEE ALSO**

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

## 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

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

**Procedure**

**Step-by-Step Procedure**

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

- Checking the Routing Table | 13

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.

## 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 | **14**
- Configuring Specific Route Export Between Routing Tables | **15**

**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 Junos OS 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.

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.

# 3
**CHAPTER**

## Static Routes

# Configure Static Routes

## Understand Basic Static Routing

Static routing is often used when the complexity of a dynamic routing protocol is not desired. A route that does not frequently change, and for which there is only one (or very few) paths to the destination, is a good candidate for static routing. The classic use case for static routing is a single-homed customer attaching to an upstream provider. This type of attachments creates a stub network.

Static routes are defined manually. The route consist of a destination prefix and a next-hop forwarding address. The static route is activated in the routing table and inserted into the forwarding table when the next-hop address is reachable. Traffic that matches the static route is forwarded to the specified next-hop address.

You can specify options that define additional information about static routes. These attributes, for example a community tag or a route metric, are included with the route when it's installed in the routing table. These additional route attributes are not required for basic static routing.

## Example: Configure IPv4 Static Routing for a Stub Network

NOTE: Our content testing team has validated and updated this example.

This example shows how to configure basic static routing for IPv4.

## Requirements

Two devices running Junos OS with a shared network link. No special configuration beyond basic device initialization (management interface, remote access, user login accounts, and so on), is required before you configure this example.

## IPv4 Static Routing Overview

**IN THIS SECTION**

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks. Stub networks have a single point of entry and egress, making them well suited to the simplicity of a static route. In Junos OS, static routes have a global preference (administrative distance) of 5. This value makes them preferred over routes learned from dynamic protocols like OSPF or BGP.

**IPv4 Static Routing Topology**

shows the example topology.

In this example, you configure the static route 192.168.47.0/24 on the provider device (R1), using a next-hop address of 172.16.1.2. This route allows the provider device to reach the remote networks at the customer site. You also configure a static default route of 0.0.0.0/0 on the customer device (R2), using a next-hop address of 172.16.1.1. The default route ensures the customer can reach all nonlocal networks by forwarding this traffic to the provider network.

Multiple loopback addresses are configured on both devices. These loopback addresses provide remote destinations to ping, so you can verify the IPv4 static routing works properly.

**Figure 1: IPv4 Stub Network Connected to a Service Provider**
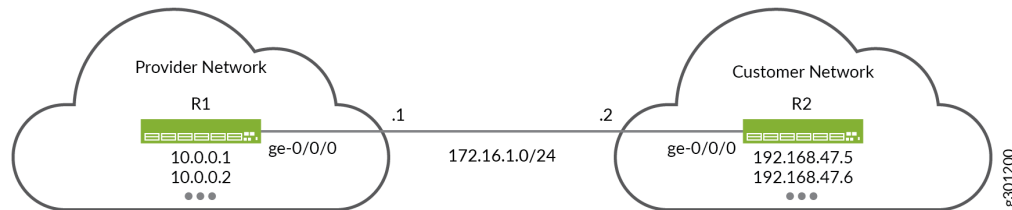


## IPv4 Static Route Configuration

**IN THIS SECTION**

**CLI Quick Configuration**

To quickly configure basic IPv4 static routing on the R1 and R2 devices, edit the following commands as needed and paste them into the CLI at the [edit] hierarchy level. Be sure to issue a commit from configuration mode to activate the changes.

**R1 Device (Provider)**

```
set system host-name R1
set interfaces ge-0/0/0 unit 0 description "Link from R1 to R2"
set interfaces ge-0/0/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
```

**R2 Device (Customer)**

```
set system host-name R2
set interfaces ge-0/0/0 unit 0 description "Link from R2 to R1"
set interfaces ge-0/0/0 unit 0 family inet address 172.16.1.2/24
```

```
set interfaces lo0 unit 0 family inet address 192.168.47.5/32
set interfaces lo0 unit 0 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
```

**Configure the R1 and R2 Devices**

**Step-by-Step Procedure**

This 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 Junos OS CLI User Guide.

To configure basic static routes:

1. Configure the hostname on the R1 (provider) device.

   ```
   [edit ]
   user@R1# set system host-name R1
   ```

2. Configure the interfaces on the R1 (provider) device.

   ```
   [edit interfaces]
   user@R1# set ge-0/0/0 unit 0 description "Link from R1 to R2"
   user@R1# set ge-0/0/0 unit 0 family inet address 172.16.1.1/24
   user@R1# set lo0 unit 0 family inet address 10.0.0.1/32
   user@R1# set lo0 unit 0 family inet address 10.0.0.2/32
   ```

3. Define the static route to the customer's prefix on the R1 device. Be sure to specify the R2 end of the point-to-point link as the next hop for the static route.

   The static route ensures the provider network can route to all remote destinations in the customer network by forwarding traffic through the R2 device.

   ```
   [edit routing-options]
   user@R1# set static route 192.168.47.0/24 next-hop 172.16.1.2
   ```

4. Commit your changes on the R1 device.

```
[edit ]
user@R1# commit
```

5. Configure the hostname on the R2 (customer) device.

```
[edit ]
user@R2# set system host-name R2
```

6. Configure the interfaces on the R2 (customer) device.

```
[edit interfaces]
user@R2# set ge-0/0/0 unit 0 description "Link from R2 to R1"
user@R2# set ge-0/0/0 unit 0 family inet address 172.16.1.2/24
user@R2# set lo0 unit 0 family inet address 192.168.47.5/32
user@R2# set lo0 unit 0 family inet address 192.168.47.6/32
```

7. Define the IPv4 static default route on the R2 device. Be sure to specify the R1 end of the point-to-point link as the next hop for the static route.

   The IPv4 default route ensures the customer can route to all nonlocal destinations by forwarding traffic to the R1 device in the provider network.

```
[edit routing-options]
user@R2# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

8. Commit your changes on the R2 device.

```
[edit]
user@R2# 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.

**R1 Device**

```
user@R1# show interfaces
ge-0/0/0 {
    unit 0 {
        description "Link from R1 to R2";
        family inet {
            address 172.16.1.1/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
            address 10.0.0.2/32;
        }
    }
}
```

```
user@R1# show routing-options
static {
    route 192.168.47.0/24 next-hop 172.16.1.2;
}
```

**R2 Device**

```
user@R2# show interfaces
ge-0/0/0 {
    unit 0 {
        description "Link from R2 to R1";
        family inet {
            address 172.16.1.2/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.47.5/32;
```

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

```
user@R2# show routing-options
static {
    route 0.0.0.0/0 next-hop 172.16.1.1;
}
```

## Verification

**IN THIS SECTION**

Confirm your IPv4 static routing is working properly.

**Check the Routing Tables**

### Purpose

Confirm the IPv4 static routes are listed as active in the routing tables of both devices.

### Action

```
user@R1> 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 lo0.0
10.0.0.2/32        *[Direct/0] 00:29:43
                    > via lo0.0
172.16.1.0/24      *[Direct/0] 00:34:40
```

```
                          > via ge-0/0/0.0
172.16.1.1/32        *[Local/0] 00:34:40
                          Local via ge-0/0/0.0
192.168.47.0/24      *[Static/5] 00:31:23
                       > to 172.16.1.2 via ge-0/0/0.0
```

```
user@R2> 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-0/0/0.0
172.16.1.2/32        *[Local/0] 00:35:21
                          Local via ge-0/0/0.0
192.168.47.5/32      *[Direct/0] 00:35:22
                       > via lo0.0
192.168.47.6/32      *[Direct/0] 00:35:21
                       > via lo0.0
```

### Meaning

The output confirms the static routes are present in the routing tables of both devices. The * symbol indicates the routes are active. The next hop for the static routes correctly point to the IP address assigned to the remote end of the link.

**Ping the Remote Loopback Addresses**

### Purpose

Verify that the IPv4 static routes provide connectivity between the loopback addresses of both devices. It's a good idea to source your test traffic from a loopback address on the local device using the source option. This approach validates forwarding between the loopback addresses of both devices in a single command.

From the R1 device, ping a loopback interface address on the R2 device.

From the R2 device, ping a loopback interface address on the R1 device.

**Action**

```
user@R1> ping 192.168.47.5 count 2 source 10.0.0.1
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=1.344 ms
64 bytes from 192.168.47.5: icmp_seq=1 ttl=64 time=1.279 ms

--- 192.168.47.5 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.279/1.312/1.344/0.032 ms
```

```
user@R2> ping 10.0.0.1 count 2 source 192.168.47.5
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=1.939 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=2.139 ms

--- 10.0.0.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.939/2.039/2.139/0.100 ms
```

**Meaning**

The output confirms the static routes allow traffic to be forwarded between the provider and customer networks.

# Example: Configure IPv6 Static Routing for a Stub Network

**IN THIS SECTION**

**NOTE**: Our content testing team has validated and updated this example.

This example shows how to configure basic static routes for IPv6.

## Requirements

Two devices running Junos OS with a shared network link. No special configuration beyond basic device initialization (management interface, remote access, user login accounts, and so on), is required before you configure this example.

## IPv6 Static Routing Overview

**IN THIS SECTION**

- IPv6 Static Routing Topology | **31**

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 when the specified next hop is reachable.

You can specify options that define additional information about static IPv6 routes. These attributes, for example a community tag or route metric, are included with the route when it's installed in the routing table. These additional route attributes are not required for basic IPv6 static routing.

**IPv6 Static Routing Topology**

provides the IPv6 static routing topology.

In this example the provider and customer networks have been allocated the IPv6 prefixes 2001:db8:1::/48 and 2001:db8:2::/48, respectively. Both networks are free to allocate longer prefixes (subnetworks) from their assigned prefix block. The point-to-point link is numbered from the provider's address space using a /126 prefix length. Each device has two loopback addresses allocated from their assigned prefix using a /128 prefix length.

You configure a static route to the customer prefix (2001:db8:2::/48) on the provider (R1) network device, using a next hop of 2001:db8:1:1::2. This route provides reachability from the provider device to the remote networks at the customer site. On the customer device (R2), you configure a static default

route of ::/0, using a next-hop address 2001:db8:1:1::1. The default route provides the customer with reachability to all nonlocal prefixes through the provider's network.

Multiple loopback addresses are configured on both devices. These loopback addresses provide remote destinations to ping, allowing you to verify the IPv6 static routing works properly.

**Figure 2: IPv6 Stub Network Connected to a Service Provider**
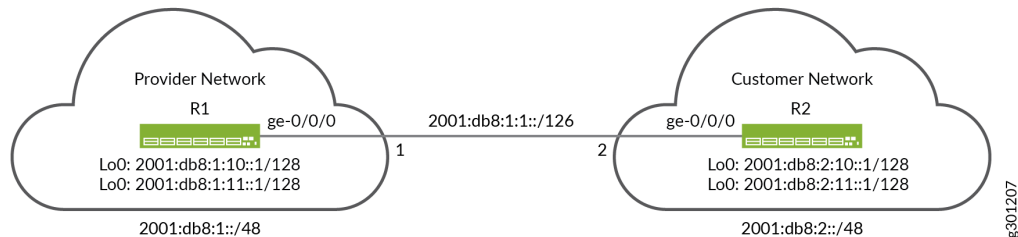


## IPv6 Static Route Configuration

**IN THIS SECTION**

**CLI Quick Configuration**

To quickly configure basic IPv6 static routing on the R1 and R2 devices, edit the following commands as needed and paste them into the CLI at the `[edit]` hierarchy level. Be sure to issue a `commit` from configuration mode to activate the changes.

**R1 Device (Provider)**

```
set system host-name R1
set interfaces ge-0/0/0 description "Link from R1 to R2"
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:db8:1:1::1/126
set interfaces lo0 unit 0 family inet6 address 2001:db8:1:10::1/128
```

```
set interfaces lo0 unit 0 family inet6 address 2001:db8:1:11::1/128
set routing-options rib inet6.0 static route 2001:db8:2::/48 next-hop 2001:db8:1:1::2
```

### R2 Device (Customer)

```
set system host-name R2
set interfaces ge-0/0/0 description "Link from R2 to R1"
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:db8:1:1::2/126
set interfaces lo0 unit 0 family inet6 address 2001:db8:2:10::1/128
set interfaces lo0 unit 0 family inet6 address 2001:db8:2:11::1/128
set routing-options rib inet6.0 static route ::/0 next-hop 2001:db8:1:1::1
```

**Configure the R1 and R2 Devices**

**Step-by-Step Procedure**

This example 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 Junos OS CLI User Guide.

Follow these steps to configure basic IPv6 static routes:

1. Configure the hostname on the R1 (provider) device.

   ```
   [edit ]
   user@R1# set system host-name R1
   ```

2. Configure the interfaces on the R1 (provider) device.

   ```
   [edit interfaces]
   user@R1# set ge-0/0/0 description "Link from R1 to R2"
   user@R1# set ge-0/0/0 unit 0 family inet6 address 2001:db8:1:1::1/126
   user@R1# set lo0 unit 0 family inet6 address 2001:db8:1:10::1/128
   user@R1# set lo0 unit 0 family inet6 address 2001:db8:1:11::1/128
   ```

3. Define the static route to the customer's IPv6 prefix on the R1 device. Be sure to set the next-hop address to the customer end of the point-to-point link.

The use of a /48 bit prefix length ensures that the R1 device can reach all possible remote destinations in the customer network by forwarding through the R2 device.

```
[edit routing-options]
user@R1# set rib inet6.0 static route 2001:db8:2::/48 next-hop 2001:db8:1:1::2
```

4. Commit your changes on the R1 device.

```
[edit ]
user@R1# commit
```

5. Configure the hostname on the R2 (customer) device.

```
[edit ]
user@R2# set system host-name R2
```

6. Configure the interfaces on the R2 (customer) device.

```
[edit interfaces]
user@R2# set ge-0/0/0 description "Link from R2 to R1"
user@R2# set ge-0/0/0 unit 0 family inet6 address 2001:db8:1:1::2/126
user@R2# set lo0 unit 0 family inet6 address 2001:db8:2:10::1/128
user@R2# set lo0 unit 0 family inet6 address 2001:db8:2:10::2/128
```

7. Define the IPv6 static default route on the R2 device. Be sure to set the next-hop address to the provider end of the point-to-point link.

The IPv6 default route ensures that the R2 device can reach all nonlocal destinations by forwarding traffic through the R1 device in the provider network.

```
[edit routing-options]
user@R2# set rib inet6.0 static route ::/0 next-hop 2001:db8:1:1::1
```

8. Commit your changes on the R2 device.

```
[edit]
user@R2# 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.

### R1 Device

```
user@R1# show interfaces
ge-0/0/0 {
    description "Link from R1 to R2";
    unit 0 {
        family inet6 {
            address 2001:db8:1:1::1/126;
        }
    }
}
lo0 {
    unit 0 {
        family inet6 {
            address 2001:db8:1:10::1/128;
            address 2001:db8:1:11::1/128;
        }
    }
}
```

```
user@R1# show routing-options
rib inet6.0 {
    static {
        route 2001:db8:2::/48 next-hop 2001:db8:1:1::2;
    }
}
```

### R2 Device

```
user@R2# show interfaces
ge-0/0/0 {
    description "Link from R2 to R1";
    unit 0 {
        family inet6 {
```

```
            address 2001:db8:1:1::2/126;
        }
    }
}
lo0 {
    unit 0 {
        family inet6 {
            address 2001:db8:2:10::1/128;
            address 2001:db8:2:11::1/128;
        }
    }
}
```

```
user@R2# show routing-options
rib inet6.0 {
    static {
        route ::/0 next-hop 2001:db8:1:1::1;
    }
}
```

## Verification

**IN THIS SECTION**

- Checking the Routing Tables | 36
- Ping the Remote Loopback Addresses | 37

Confirm IPv6 static routing works properly.

**Checking the Routing Tables**

## Purpose

Verify the IPv6 static routes are active in the routing tables of both devices.

## Action

```
user@R1> show route protocol static
inet6.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


2001:db8:2::/48    *[Static/5] 02:07:11
                    >  to 2001:db8:1:1::2 via ge-0/0/0.0
```

```
user@R2> show route protocol static
inet6.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


::/0                   *[Static/5] 02:13:56
                    >  to 2001:db8:1:1::1 via ge-0/0/0.0
```

## Meaning

The output confirms the IPv6 static routes are present in the routing tables of both devices. The *
symbol indicates the routes are active. Both the static routes correctly point to the remote end of the
point-to-point link as the next hop for matching traffic.

**Ping the Remote Loopback Addresses**

## Purpose

Verify that the IPv6 static routes provide connectivity between the loopback addresses of both devices.
It's a good idea to source your test traffic from a loopback address on the local device using the source
option. This approach validates forwarding between the loopback addresses of both devices in a single
command.

From the R1 device, ping a loopback address on the R2 device.

From the R2 device, ping q loopback address on the R1 device.

## Action

```
user@R1> ping 2001:db8:2:10::1 source 2001:db8:1:10::1 count 2
PING6(56=40+8+8 bytes) 2001:db8:1:10::1 --> 2001:db8:2:10::1
```

```
16 bytes from 2001:db8:2:10::1, icmp_seq=0 hlim=64 time=2.770 ms
16 bytes from 2001:db8:2:10::1, icmp_seq=1 hlim=64 time=2.373 ms

--- 2001:db8:2:10::1 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 2.373/2.572/2.770/0.198 ms
```

```
user@R2> ping 2001:db8:1:10::1 source 2001:db8:2:10::1 count 2
PING6(56=40+8+8 bytes) 2001:db8:2:10::1 --> 2001:db8:1:10::1
16 bytes from 2001:db8:1:10::1, icmp_seq=0 hlim=64 time=1.985 ms
16 bytes from 2001:db8:1:10::1, icmp_seq=1 hlim=64 time=1.704 ms

--- 2001:db8:1:10::1 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.704/1.845/1.985/0.140 ms
```

**Meaning**

The output confirms the IPv6 static routes allow traffic to be forwarded between the provider and customer networks.

# Static Route Preferences and Qualified Next Hops

**IN THIS SECTION**

## 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.

## 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 SECTION**

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 3 on page 41.

**Figure 3: Controlling Static Route Selection**
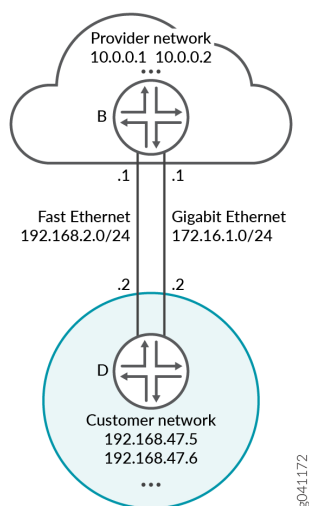


**Topology**

## Configuration

**IN THIS SECTION**

- Procedure | **41**

**Procedure**

## CLI Quick Configuration

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

### Device 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 Junos OS 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.

# 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. illustrates the issue.

**Figure 4: Inefficient Use of IP Address Space**
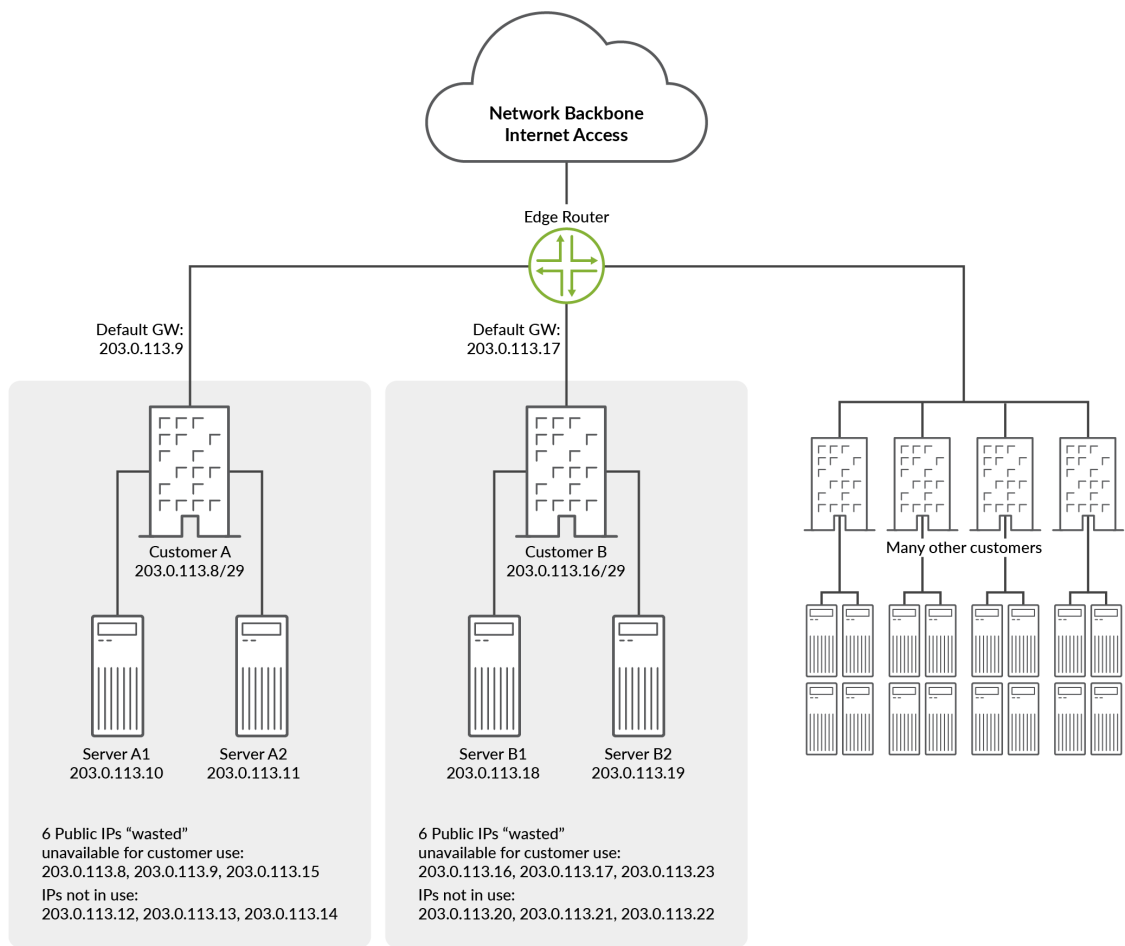


In this configuration, each customer is allocated a /29 block of address space. For each block, the network, broadcast, and gateway addresses are not available for server IP addressing, which results in three IP addresses being used inefficiently. In addition, the blocks consume unused IP addresses for 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 5 on page 50 shows the efficient use of IP address space.

**Figure 5: Configuration Using the Shared Address Space**
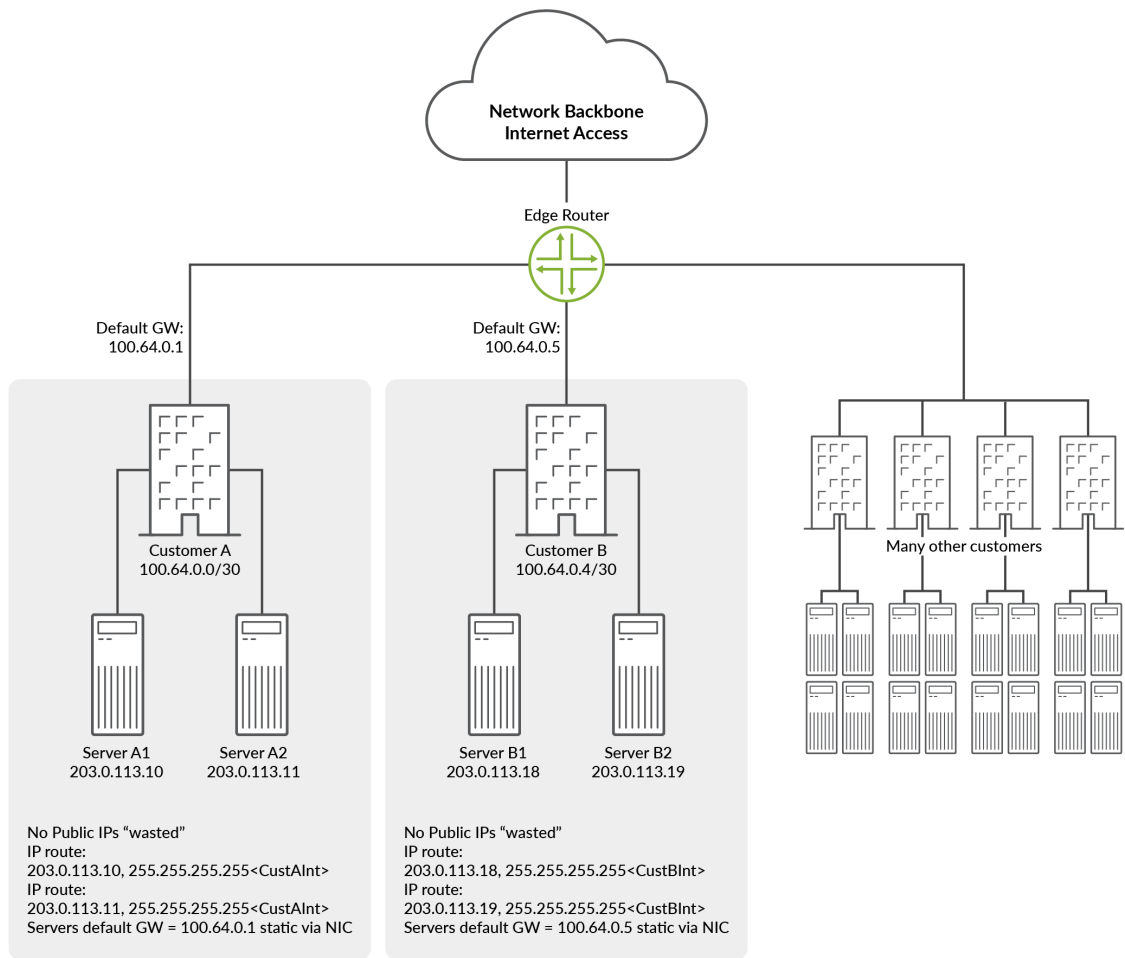


In this configuration, each customer gets allocated individual IP addresses per server. There is a static route that can be configured as a host route. The interface in the router gets allocated an IP address from the RFC 6598 space, so it does not consume publicly routable address space, and connectivity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

## 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.

## 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

**IN THIS SECTION**

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.

**Topology**

shows the sample network.

**Figure 6: Customer Routes Connected to a Service Provider**
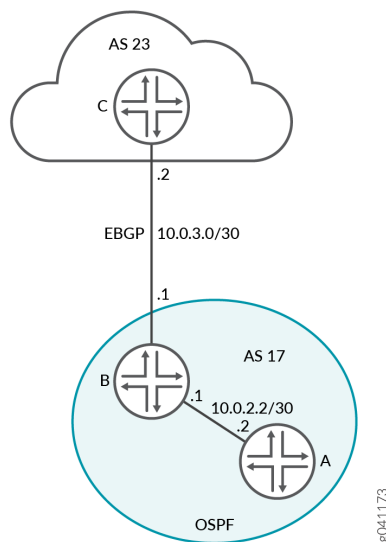


**Configuration**

**IN THIS SECTION**

- Procedure | **55**

**Procedure**

**CLI Quick Configuration**

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

**Device 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 Junos OS 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.

## Verifying the Static Route Configuration

**IN THIS SECTION**

## 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

**command-name**

```
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.

### SEE ALSO

*show route terse*

[CLI Explorer](CLI Explorer)

**Release History Table**

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

# Bidirectional Forwarding Detection for Static Routes

**IN THIS SECTION**

## 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 Firewalls. 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.

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**: 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.

### SEE ALSO

*Enabling Dedicated and Real-Time BFD*

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

**IN THIS SECTION**

- Requirements | **69**

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

**IN THIS SECTION**

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.

shows the sample network.

**Figure 7: Customer Routes Connected to a Service Provider**



**Topology**

## 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
```

**Procedure**

**Step-by-Step Procedure**

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

## Understanding BFD Authentication for Static Route Security

**IN THIS SECTION**

- BFD Authentication Algorithms | **79**
- Security Authentication Keychains | **80**
- Strict Versus Loose Authentication | **80**

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.

## 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

**IN THIS SECTION**

- Topology | **82**

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
> ```

shows the sample network.

**Figure 8: Customer Routes Connected to a Service Provider**



**Topology**

## 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"
```

**Procedure**

**Step-by-Step Procedure**

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

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

## 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 SECTION**

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 9: BFD Enabled on Qualified Next Hops**
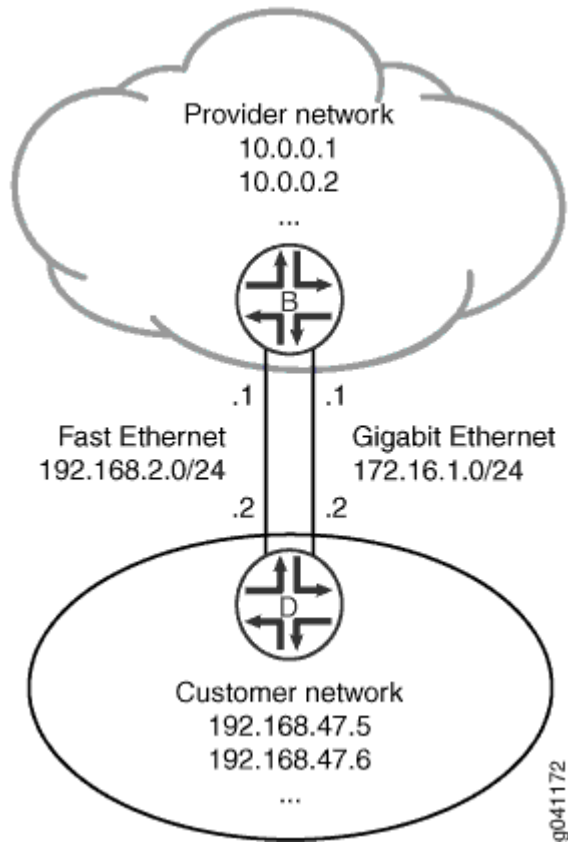


**Topology**

## Configuration

**IN THIS SECTION**

- Procedure | **92**

**Procedure**

**CLI Quick Configuration**

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

**Device 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 Junos OS 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.

**Release History Table**

| Release | Description |
| --- | --- |
| 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 Firewalls. This field is applicable only for the static routes. |

# Static Routes for CLNS

## Understanding Static Routes for CLNS

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

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

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

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

**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, 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
```

**Procedure**

**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.

CLNS Configuration Overview

# 4
**CHAPTER**

# Route Aggregation

# Configuring Route Aggregation

## 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;
}
```

## 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 R1 has a static default route to reach the ISP network (10.0.45.0).

- Device R2 has static routes configured to reach Device R1's customer networks (10.200.1.0/24 and 10.200.2.0/24).

- Device R2 also has a routing policy configured to advertise all static routes to its neighbor, Device R3.

- When Device R3 sends information about these routes (10.200.1.0/24 and 10.200.2.0/24) to Device ISP, the information is summarized as a single aggregate route (10.200.0.0/16).

- Device R2 and Device R3 share an IBGP session and have OSPF as the IGP.

- Device ISP injects a default route into AS 64501.

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

**Topology**

shows the sample network.

**Figure 10: Aggregate Route Advertised to an ISP**
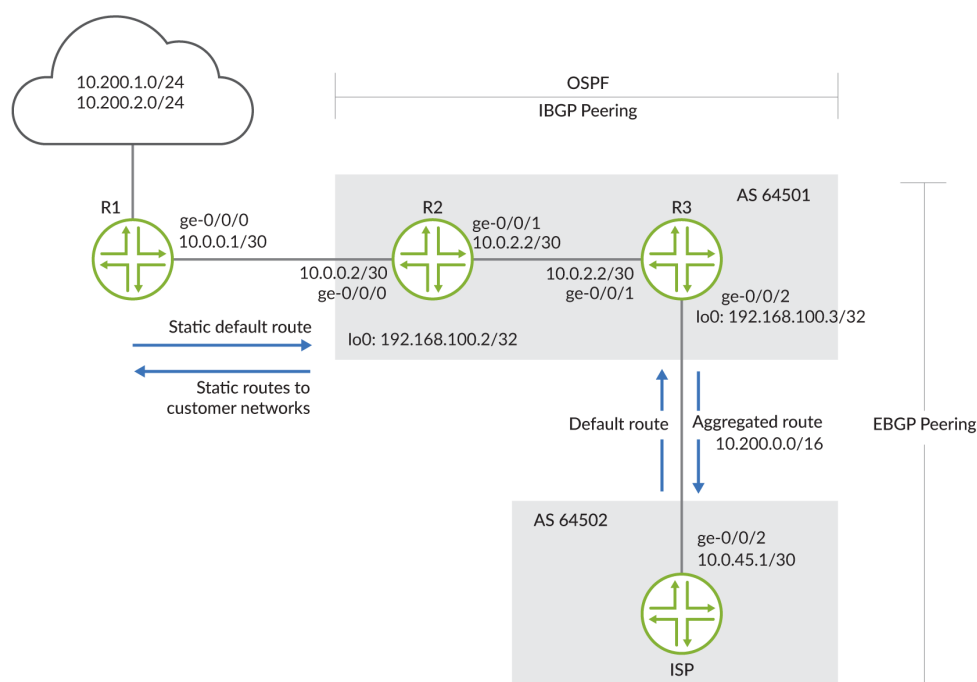


## Configuration

**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-0/0/0 description R1-to-R2
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/30
```

```
set interfaces lo0 unit 0 family inet address 10.200.1.1/24
set interfaces lo0 unit 0 family inet address 10.200.2.2/24
set routing-options static route 0.0.0.0/0 next-hop 10.0.0.2
```

**Device R2**

```
set interfaces ge-0/0/0 description R2-to-R1
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-0/0/1 description R2-to-R3
set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.2/30
set interfaces lo0 unit 0 family inet address 192.168.100.2/32
set policy-options policy-statement send-customer-routes from protocol static
set policy-options policy-statement send-customer-routes then accept
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.100.2
set protocols bgp group internal export send-customer-routes
set protocols bgp group internal neighbor 192.168.100.3
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options autonomous-system 64501
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
```

**Device R3**

```
set interfaces ge-0/0/1 description R3-to-R2
set interfaces ge-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces ge-0/0/2 description R3-to-ISP
set interfaces ge-0/0/2 unit 0 family inet address 10.0.45.2/30
set interfaces lo0 unit 0 family inet address 192.168.100.3/32
set policy-options policy-statement next-hop-self term 1 from protocol bgp
set policy-options policy-statement next-hop-self term 1 then next-hop self
set policy-options policy-statement next-hop-self term 1 then accept
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 protocols bgp group external type external
set protocols bgp group external export send-aggregate
```

```
set protocols bgp group external peer-as 64502
set protocols bgp group external neighbor 10.0.45.1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.100.3
set protocols bgp group internal export next-hop-self
set protocols bgp group internal neighbor 192.168.100.2
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options autonomous-system 64501
set routing-options aggregate route 10.200.0.0/16
```

**Device ISP**

```
set interfaces ge-0/0/2 description ISP-to-R3
set interfaces ge-0/0/2 unit 0 family inet address 10.0.45.1/30
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 protocols bgp group external type external
set protocols bgp group external export advertise-default
set protocols bgp group external peer-as 64501
set protocols bgp group external neighbor 10.0.45.2
set routing-options autonomous-system 64502
set routing-options static route 0.0.0.0/0 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 Junos OS CLI User Guide.

To configure Device R3:

1. Configure the device interfaces.

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

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

```
user@R3# set lo0 unit 0 family inet address 192.168.100.3/32
```

2. Configure the AS number.

```
[edit routing-options]
user@R3# set autonomous-system 64501
```

3. Configure an EBGP session with the ISP device.

```
[edit protocols]
user@R3# set bgp group external type external
user@R3# set bgp group external peer-as 64502
user@R3# bgp group external neighbor 10.0.45.1
```

4. Configure an IBGP session with Device R2.

```
[edit protocols]
user@R3# set bgp group internal type internal
user@R3# set bgp group internal local-address 192.168.100.3
user@R3# set bgp group internal neighbor 192.168.100.2
```

5. Configure OSPF as the IGP.

```
[edit protocols]
user@R3# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@R3# set ospf area 0.0.0.0 interface lo0.0 passive
```

6. Configure the aggregate route for the customer network routes.

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

7. Configure a 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]
user@R3# set policy-statement send-aggregate term 1 from protocol aggregate
user@R3# set policy-statement send-aggregate term 1 then accept

user@R3# set policy-statement send-aggregate term suppress-specific-routes from route-
filter 10.200.0.0/16 longer
user@R3# set policy-statement send-aggregate term suppress-specific-routes then reject
```

8. Configure a routing policy to report Device R3 as the next hop as a result of participating in the EBGP session with Device ISP.

```
[edit policy-options]
user@R3# set policy-statement next-hop-self term 1 from protocol bgp
user@R3# set policy-statement next-hop-self term 1 then next-hop self
user@R3# set policy-statement next-hop-self term 1 then accept
```

9. Apply the aggregate route policy to the EBGP session with Device ISP.

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

10. Apply the next-hop-self policy to the IBGP session with Device R2.

```
[edit protocols]
user@R3# set bgp group internal export next-hop-self
```

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-0/0/1 {
    description R3-to-R2;
    unit 0 {
        family inet {
            address 10.0.2.1/30;
        }
    }
}
ge-0/0/2 {
    description R3-to-ISP;
    unit 0 {
        family inet {
            address 10.0.45.2/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.100.3/32;
        }
    }
}
user@R3# show protocols
bgp {
    group external {
        type external;
        export send-aggregate;
        peer-as 64502;
        neighbor 10.0.45.1;
    }
    group internal {
        type internal;
        local-address 192.168.100.3;
        export next-hop-self;
```

```
            neighbor 192.168.100.2;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface ge-0/0/1.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
user@R3# show policy-options
policy-statement next-hop-self {
    term 1 {
        from protocol bgp;
        then {
            next-hop self;
            accept;
        }
    }
}
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;
    }
}
user@R3# show routing-options
autonomous-system 64501;
aggregate {
    route 10.200.0.0/16;
}
```

## Verification

Confirm that the configuration is working properly.

**Verifying That Device R3 Has the Expected Routes**

### Purpose

Confirm that Device R3 has the advertised static routes from Device R2.

### 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      64502 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 learned the static routes configured on Device R2 to reach Device R1's customer networks (10.200.1.0/24 and 10.200.2.0/24) through IBGP peering.

**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: 20 destinations, 20 routes (20 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.

**Verifying End-to-End Connection**

**Purpose**

Confirm end-to-end connection from the customer network on Device R1 to Device ISP.

**Action**

```
user@R1>ping 10.0.45.2 source 10.200.1.1
PING 10.0.45.2 (10.0.45.2): 56 data bytes
64 bytes from 10.0.45.2: icmp_seq=0 ttl=63 time=3.953 ms
64 bytes from 10.0.45.2: icmp_seq=1 ttl=63 time=4.979 ms
64 bytes from 10.0.45.2: icmp_seq=2 ttl=63 time=3.789 ms
```

**Meaning**

The output shows a successful ping verifying the reachability to Device ISP from customer network 10.200.1.1.

**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

Understanding Conditionally Generated Routes

*Example: Configuring a Conditional Default Route Policy*

# 5

**CHAPTER**

## Martian Addresses

# Recognize Martian Addresses for Routing

**IN THIS SECTION**

## Understanding Martian Addresses

Martian addresses are host or network addresses about which all routing information is ignored. When received by the routing device, these routes are ignored. They commonly are sent by improperly configured systems on the network and have destination addresses that are obviously invalid.

In IPv6, the loopback address and the multicast resolve and discard routes are the default martian addresses.

In Junos OS Release 10.4R5 and later, the reserved IPv6 multicast address space (ff00::/8 and ff02::/16) is added to the list of martian addresses.

In Junos OS Release 9.6 and later, you can configure Class E addresses on interfaces. Class E addresses are treated like any other unicast address for the purpose of forwarding. To allow Class E addresses to be configured on interfaces, you must remove the Class E prefix from the list of martian addresses. To remove the Class E prefix from the list of martian addresses include the `martians 240/4 orlonger allow` statement at the `[edit routing-options]` hierarchy level.

To view the default and configured martian routes, run the `show route martians` command.

### IPv4 Martian Addresses

```
user@host> show route martians table inet.

inet.0:
          0.0.0.0/0 exact -- allowed
          0.0.0.0/8 orlonger -- disallowed
          127.0.0.0/8 orlonger -- disallowed
          192.0.0.0/24 orlonger -- disallowed
          240.0.0.0/4 orlonger -- disallowed
          224.0.0.0/4 exact -- disallowed
```

```
              224.0.0.0/24 exact -- disallowed


inet.1:

              0.0.0.0/0 exact -- allowed
              0.0.0.0/8 orlonger -- disallowed
              127.0.0.0/8 orlonger -- disallowed
              192.0.0.0/24 orlonger -- disallowed
              240.0.0.0/4 orlonger -- disallowed


inet.2:

              0.0.0.0/0 exact -- allowed
              0.0.0.0/8 orlonger -- disallowed
              127.0.0.0/8 orlonger -- disallowed
              192.0.0.0/24 orlonger -- disallowed
              240.0.0.0/4 orlonger -- disallowed
              224.0.0.0/4 exact -- disallowed
              224.0.0.0/24 exact -- disallowed


inet.3:

              0.0.0.0/0 exact -- allowed
              0.0.0.0/8 orlonger -- disallowed
              127.0.0.0/8 orlonger -- disallowed
              192.0.0.0/24 orlonger -- disallowed
              240.0.0.0/4 orlonger -- disallowed
              224.0.0.0/4 exact -- disallowed
              224.0.0.0/24 exact -- disallowed
```

## IPv6 Martian Addresses

```
user@host> show route martians table inet6
inet6.0:
              ::1/128 exact -- disallowed
              ff00::/8 exact -- disallowed
              ff02::/16 exact -- disallowed


inet6.1:
              ::1/128 exact -- disallowed


inet6.2:
              ::1/128 exact -- disallowed
              ff00::/8 exact -- disallowed
              ff02::/16 exact -- disallowed
```

```
inet6.3:

            ::1/128 exact -- disallowed
            ff00::/8 exact -- disallowed
            ff02::/16 exact -- disallowed
```

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

**IN THIS SECTION**

- Requirements | **130**
- Overview | **130**
- Configuration | **130**
- Verification | **133**

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

**IN THIS SECTION**

- CLI Quick Configuration | **131**

**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 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
```

**Procedure**

**Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the Junos OS 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. 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 inet.3 {
    martians {
        240.0.0.0/4 orlonger allow;
    }
}
rib inet.2 {
    martians {
        240.0.0.0/4 orlonger allow;
    }
}
martians {
```

```
        240.0.0.0/4 orlonger allow;
}
```

## Verification

Confirm that the configuration is working properly.

**Verifying That the 240.0.0.0/4 Routes Are Now Accepted**

### 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.

RELATED DOCUMENTATION

*Example: Creating an Interface on a Logical System*

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

# 6
**CHAPTER**

# Packet Forwarding

# Configuring Packet Forwarding Behavior

**IN THIS SECTION**

## 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 11 on page 136 illustrates the route to forwarding next-hop bindings with indirect next hop disabled.

**Figure 11: Route to Forwarding Next-Hop Bindings**
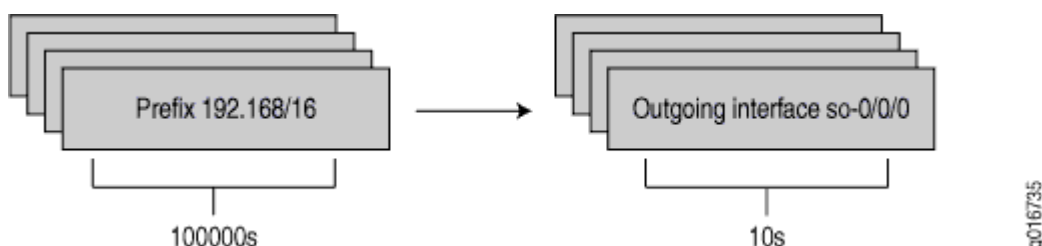


You can enable Junos OS to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table. As a result, fewer route to forwarding next-hop bindings need to be updated, which improves the route convergence time. Figure 12 on page 137 illustrates the route to forwarding next-hop bindings with indirect next hop enabled.

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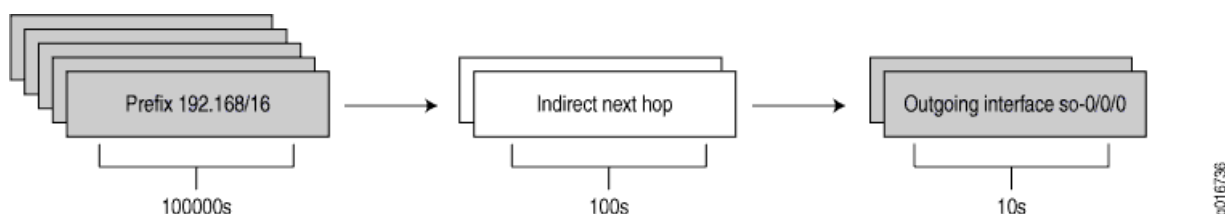


# Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine

**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 SECTION**

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.

**Topology**

Figure 13 on page 138 shows the sample network.

**Figure 13: Sample Topology for Indirect Next Hops**
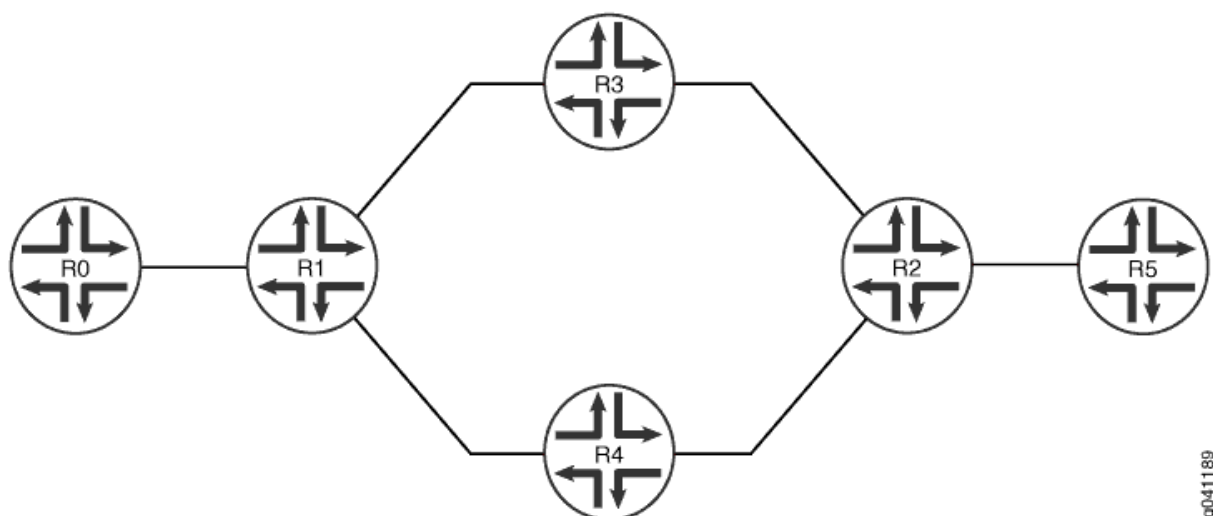


The section shows the full configuration on all of the devices in . Otherwise, the example focuses on Device R0, Device R1, and Device R2.

## 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 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 10.1.0.1/32
set interfaces lo0 unit 1 family inet address 10.1.0.2/32
set interfaces lo0 unit 1 family inet address 10.1.0.3/32
set interfaces lo0 unit 1 family inet address 10.1.0.4/32
set interfaces lo0 unit 1 family inet address 10.1.0.5/32
set interfaces lo0 unit 1 family inet address 10.1.0.6/32
set interfaces lo0 unit 1 family inet address 10.1.0.7/32
set interfaces lo0 unit 1 family inet address 10.1.0.8/32
set interfaces lo0 unit 1 family inet address 10.1.0.9/32
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 10.1.1.1/32
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 10.1.1.1
set protocols bgp group int neighbor 10.2.2.2
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 10.1.0.2/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.1/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.3/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.4/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.5/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.6/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.7/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.8/32 next-hop 10.0.0.1
set routing-options static route 10.1.0.9/32 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 10.2.2.2/32
set protocols bgp export send-local
set protocols bgp group int type internal
set protocols bgp group int local-address 10.2.2.2
set protocols bgp group int family inet unicast
set protocols bgp group int family inet-vpn unicast
set protocols bgp group int neighbor 10.1.1.1
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 10.3.3.3/32
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 10.4.4.4/32
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 10.5.5.5/32
```

**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 Junos OS 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 10.1.0.1/32
user@R0# set lo0 unit 1 family inet address 10.1.0.2/32
```

```
user@R0# set lo0 unit 1 family inet address 10.1.0.3/32
user@R0# set lo0 unit 1 family inet address 10.1.0.4/32
user@R0# set lo0 unit 1 family inet address 10.1.0.5/32
user@R0# set lo0 unit 1 family inet address 10.1.0.6/32
user@R0# set lo0 unit 1 family inet address 10.1.0.7/32
user@R0# set lo0 unit 1 family inet address 10.1.0.8/32
user@R0# set lo0 unit 1 family inet address 10.1.0.9/32
```

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 Junos OS 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 10.1.1.1/32
```

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 10.1.1.1
user@R1# set bgp group int neighbor 10.2.2.2
```

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 10.1.0.2/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.1/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.3/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.4/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.5/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.6/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.7/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.8/32 next-hop 10.0.0.1
user@R1# set routing-options static route 10.1.0.9/32 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 Junos OS 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 10.2.2.2/32
```

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 10.2.2.2
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 10.1.1.1
```

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 10.1.0.1/32;
            address 10.1.0.2/32;
            address 10.1.0.3/32;
            address 10.1.0.4/32;
            address 10.1.0.5/32;
            address 10.1.0.6/32;
            address 10.1.0.7/32;
            address 10.1.0.8/32;
            address 10.1.0.9/32;
        }
    }
}
```

```
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 10.1.1.1/32;
            }
        }
    }
}
```

```
user@R1# show protocols
bgp {
    export [ send-local send-static ];
    group int {
        type internal;
        local-address 10.1.1.1;
        neighbor 10.2.2.2;
    }
}
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 10.1.0.2/32 next-hop 10.0.0.1;
    route 10.1.0.1/32 next-hop 10.0.0.1;
    route 10.1.0.3/32 next-hop 10.0.0.1;
    route 10.1.0.4/32 next-hop 10.0.0.1;
    route 10.1.0.5/32 next-hop 10.0.0.1;
    route 10.1.0.6/32 next-hop 10.0.0.1;
    route 10.1.0.7/32 next-hop 10.0.0.1;
    route 10.1.0.8/32 next-hop 10.0.0.1;
    route 10.1.0.9/32 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 10.2.2.2/32;
                }
            }
        }
```

```
user@R2# show protocols
bgp {
    export send-local;
    group int {
        type internal;
        local-address 10.2.2.2;
        family inet {
            unicast;
        }
        family inet-vpn {
            unicast;
        }
        neighbor 10.1.1.1;
    }
}
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 139.

## Verification

**IN THIS SECTION**

- Verifying That the Routes Have the Expected Indirect-Next-Hop Flag | **150**

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: 10.50.244.9 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.

# 7
**CHAPTER**

# Troubleshooting

# Troubleshooting Network Issues

## 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 2: Checklist for Working with Problems on Your Network**

| Tasks | Command or Action |
|---|---|
| *Isolating a Broken Network Connection* | |
| 1. *Identifying the Symptoms of a Broken Network Connection* | ping (*ip-address* \| *hostname*) show route (*ip-address* \| *hostname*) traceroute (*ip-address* \| *hostname*) |
| 1. *Isolating the Causes of a Network Problem* | `show < configuration | interfaces | protocols | route >` |
| 1. *Taking Appropriate Action for Resolving the Network Problem* | `[edit]` `delete routing options static route` *destination-prefix* **commit and-quit show route** *destination-prefix* |
| 1. *Evaluating the Solution to Check Whether the Network Problem Is Resolved* | `show route (`*ip-address* \| *hostname*`) ping (`*ip-address* \| *hostname*`)` **count 3 traceroute (***ip-address* \| *hostname***)** |

## Isolating a Broken Network Connection

By applying the standard four-step process illustrated in , 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 14: Process for Diagnosing Problems in Your Network**
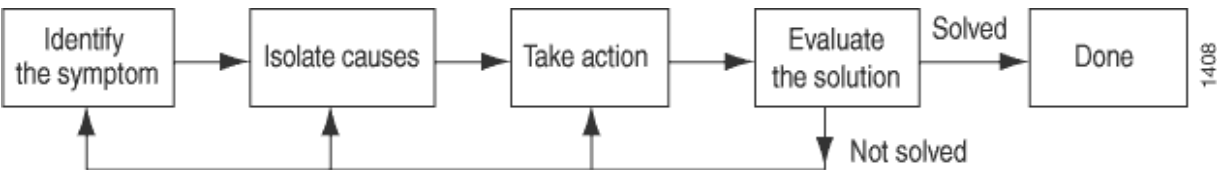
Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic 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 15 on page 155 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

**Figure 15: Network with a Problem**
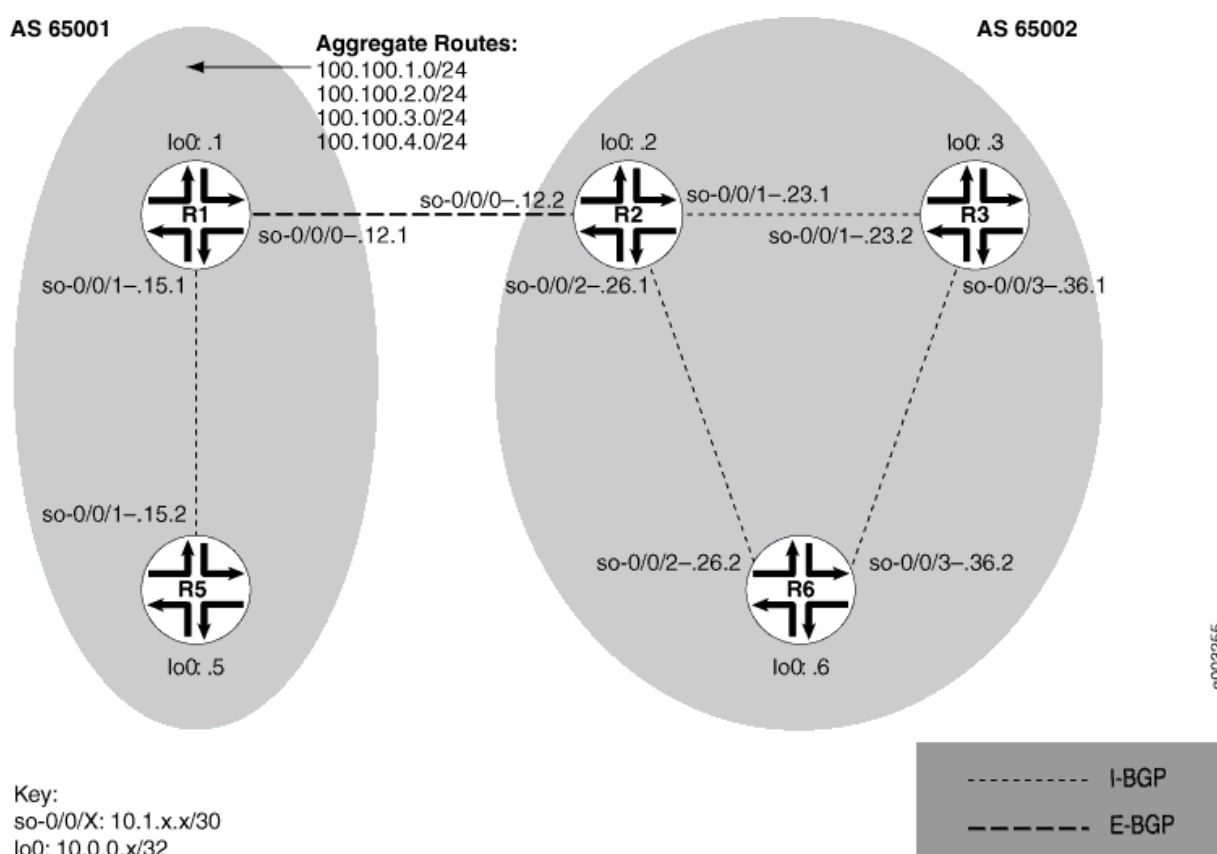


The network in Figure 15 on page 155 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (R1) in AS 65001 announces aggregated prefixes `100.100/24` to the AS 65002 network. The problem in this network is that R6 does not have 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*

- *Taking Appropriate Action for Resolving the Network Problem*

- *Taking Appropriate Action for Resolving the Network Problem*

- *Evaluating the Solution to Check Whether the Network Problem Is Resolved*

## Identifying the Symptoms of a Broken Network Connection

**IN THIS SECTION**

### Problem

### 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

**IN THIS SECTION**

- Problem | **160**
- Solution | **160**

## Problem

### Description

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

### Solution

- Check the local router's configuration and edit it if appropriate.

- Troubleshoot the intermediate router.

- Check the remote host configuration and edit it if appropriate.

- Troubleshoot routing protocols.

- Identify additional possible causes.

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

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

**Sample Output**

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

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

user@R2> show route 10.0.0.5

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

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

### Meaning

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

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

**IN THIS SECTION**

- Problem | **161**
- Solution | **162**

### Problem

### Description

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

You can address possible causes in any order. In relation to the network in *Isolating a Broken Network Connection*, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you have reason to believe that the problem is related to a known issue, such as a recent change in configuration.

## Solution

To evaluate the solution, enter the following Junos OS CLI commands:

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

## Sample Output

```
user@R6> show route 10.0.0.5


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


10.0.0.5/32        *[BGP/170]  00:01:35, MED 5, localpref 100, from 10.0.0.2
                     AS path: 65001 I
                   > to 10.1.26.1 via so-0/0/2.0


user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=253 time=0.866 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=253 time=0.837 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=253 time=0.796 ms
^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.796/0.833/0.866/0.029 ms


user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.629 ms  0.538 ms  0.497 ms
 2  10.1.12.1 (10.1.12.1)  0.534 ms  0.538 ms  0.510 ms
 3  10.0.0.5 (10.0.0.5)  0.776 ms  0.705 ms  0.672 ms
```

## Meaning

The sample output shows that there is now a connection between R6 and R5. The `show route` command shows that the BGP route to R5 is preferred, as indicated by the asterisk (*). The `ping` command is successful and the `traceroute` command shows that the path from R6 to R5 is through R2 (`10.1.26.1`), and then through R1 (`10.1.12.1`).

## Checklist for Tracking Error Conditions

### Problem

#### Description

Table 3 on page 163 provides links and commands for configuring routing protocol daemon tracing, Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS) protocol, and Open Shortest Path First (OSPF) protocol tracing to diagnose error conditions.

### Solution

**Table 3: Checklist for Tracking Error Conditions**

| Tasks | Command or Action |
|-------|-------------------|
| **Configure Routing Protocol Process Tracing** | |
| 1. *Configure Routing Protocol Process Tracing* | `[edit] edit routing-options traceoptions` *filename* `size` *size* **files** *number* **show co** log *filename* |
| 1. *Configure Routing Protocol Tracing for a Specific Routing Protocol* | `[edit] edit protocol` *protocol-name* **trace** *filename* **size** *size* **files** *number* **show con** log *filename* |
| 1. *Monitor Trace File Messages Written in Near-Real Time* | **monitor start** *filename* |
| 1. *Stop Trace File Monitoring* | **monitor stop** *filename* |

**Table 3: Checklist for Tracking Error Conditions** *(Continued)*

| Tasks | Command or Action |
|-------|-------------------|
| Configure BGP-Specific Options | |
| **1.** Display Detailed BGP Protocol Information | [edit] edit protocol bgp traceoptions se<br>detail show commit run show log *filename* |
| **1.** Display Sent or Received BGP Packets | [edit] edit protocol bgp traceoptions se<br>(send \| receive) show commit run show lo |
| **1.** Diagnose BGP Session Establishment Problems | [edit] edit protocol bgp set traceoption<br>detail show commit run show log *filename* |
| Configure IS-IS-Specific Options | |
| **1.** Displaying Detailed IS-IS Protocol Information | [edit] edit protocol isis traceoptions s<br>detail show commit run show log *filename* |
| **1.** Displaying Sent or Received IS-IS Protocol Packets | [edit] edit protocols isis traceoptions<br>(send \| receive) show commit run show lo |
| **1.** Analyzing IS-IS Link-State PDUs in Detail | [edit] edit protocols isis traceoptions<br>detail show commit run show log *filename* |
| Configure OSPF-Specific Options | |
| **1.** Diagnose OSPF Session Establishment Problems | [edit] edit protocols ospf traceoptions<br>detail show commit run show log *filename* |
| **1.** Analyze OSPF Link-State Advertisement Packets in Detail | [edit] edit protocols ospf traceoptions<br>update detail show commit run show log |

## Configure Routing Protocol Process Tracing

### Action

To configure routing protocol process (rpd) tracing, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit routing-options traceoptions
   ```

2. Configure the file, file size, number, and flags:

   ```
   [edit routing-options traceoptions]
   user@host#  set file filename size size file number
   [edit routing-options traceoptions]
   user@host#  set flag flag
   ```

   For example:

   ```
   [edit routing-options traceoptions]
   user@host# set file daemonlog size 10240 files 10
   [edit routing-options traceoptions]
   user@host# set flag general
   ```

3. Verify the configuration:

   ```
   user@host# show
   ```

For example:

```
[edit routing-options traceoptions]
user@host# show
file daemonlog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```

> **NOTE**: Some traceoptions flags generate an extensive amount of information. Tracing can also slow down the operation of routing protocols. Delete the traceoptions configuration if you no longer require it.

1. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit routing-options traceoptions]
user@pro4-a# run show log daemonlog
Sep 17 14:17:31 trace_on: Tracing to "/var/log/daemonlog" started
Sep 17 14:17:31 Tracing flags enabled: general
Sep 17 14:17:31 inet_routerid_notify: Router ID: 10.255.245.44
Sep 17 14:17:31 inet_routerid_notify: No Router ID assigned
Sep 17 14:17:31 Initializing LSI globals
Sep 17 14:17:31 LSI initialization complete
Sep 17 14:17:31 Initializing OSPF instances
Sep 17 14:17:31 Reinitializing OSPFv2 instance master
Sep 17 14:17:31 OSPFv2 instance master running
[...Output truncated...]
```

## Meaning

lists tracing flags and example output for Junos-supported routing protocol daemon tracing.

**Table 4: Routing Protocol Daemon Tracing Flags**

| Tracing Flag | Description | Example Output |
|---|---|---|
| **all** | All operations | Not available. |
| **general** | Normal operations and routing table change | Not available. |
| **normal** | Normal operations | Not available. |
| **policy** | Policy operations and actions | Nov 29 22:19:58 export: Dest 10.0.0.0 proto Static Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0 Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0 Nov 29 22:19:58 export: Dest 10.10.10.0 proto IS-IS |
| **route** | Routing table changes | Nov 29 22:23:59 Nov 29 22:23:59 rtlist_walker_job: rt_list walk for RIB inet.0 started with 42 entries Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) start Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) done Nov 29 22:23:59 rtlist_walker_job: rt_list walk for inet.0 ended with 42 entries Nov 29 22:23:59 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 CHANGE route/user af 2 addr 172.16.0.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 172.17.0.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 10.149.3.0 nhop-type unicast nhop 10.10.10.33 Nov 29 22:24:19 trace_on: Tracing to "/var/log/rpdlog" started Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 10.10.218.0 nhop-type unicast nhop 10.10.10.29 Nov 29 22:24:19 RELEASE 10.10.218.0 255.255.255.0 gw 10.10.10.29,10.10.10.33 BGP pref 170/-101 metric so-1/1/0.0,so-1/1/1.0 <Release Delete Int Ext> as 65401 Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 172.18.0.0 nhop-type unicast nhop 10.10.10.33 |
| state | State transitions | Not available. |

**Table 4: Routing Protocol Daemon Tracing Flags** *(Continued)*

| Tracing Flag | Description | Example Output |
|---|---|---|
| task | Interface transactions and processing | Nov 29 22:50:04 foreground dispatch running job task_collect for task Scheduler Nov 29 22:50:04 task_collect_job: freeing task MGMT_Listen (DELETED) Nov 29 22:50:04 foreground dispatch completed job task_collect for task Scheduler Nov 29 22:50:04 background dispatch running job rt_static_update for task RT Nov 29 22:50:04 task_job_delete: delete background job rt_static_update for task RT Nov 29 22:50:04 background dispatch completed job rt_static_update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 background dispatch returned job Flash update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT Nov 29 22:50:04 background dispatch completed job Flash update for task RT Nov 29 22:50:04 background dispatch running job Flash update for task RT Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT |
| timer | Timer usage | Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 1 timer Nov 29 22:52:07 task_timer_hiprio_dispatch: running high priority timer queue Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 1 timer Nov 29 22:52:07 main: running normal priority timer queue Nov 29 22:52:07 main: ran 2 timers |

## Configure Routing Protocol Tracing for a Specific Routing Protocol

**IN THIS SECTION**

- Action | **168**
- Meaning | **170**

### Action

To configure routing protocol tracing for a specific routing protocol, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
                 user@host# edit protocol protocol-name traceoptions
```

2. Configure the file, file size, number, and flags:

```
[edit protocols protocol name traceoptions]
                 user@host# set file filename size size files
number
[edit protocols protocol name traceoptions]
user@host# set flag flag
```

For example:

```
[edit protocols ospf traceoptions]
user@host# set file ospflog size 10240 files 10
[edit protocols ospf traceoptions]
user@host# set flag general
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospflog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```

**5.** View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit protocols ospf traceoptions]
user@pro4-a# run show log ospflog
Sep 17 14:23:10 trace_on: Tracing to "/var/log/ospflog" started
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) start
Sep 17 14:23:10 OSPF:  multicast address 224.0.0.5/32, route ignored
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) done
Sep 17 14:23:10 CHANGE    10.255.245.46/32   gw 10.10.208.67  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Delete Int>
Sep 17 14:23:10 CHANGE    10.255.245.46/32   gw 10.10.208.67  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 ADD       10.255.245.46/32   gw 10.10.208.67  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 CHANGE    10.255.245.48/32   gw 10.10.208.69  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Delete Int>
Sep 17 14:23:10 CHANGE    10.255.245.48/32   gw 10.10.208.69  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 ADD       10.255.245.48/32   gw 10.10.208.69  OSPF     pref 10/0 metric 1/0
fe-0/0/0.0 <Active Int>
Sep 17 14:23:10 rt_close: 4/4 routes proto OSPF
[...Output truncated...]
```

## Meaning

Table 5 on page 170 lists standard tracing options that are available globally or that can be applied to specific protocols. You can also configure tracing for a specific BGP peer or peer group. For more information, see the *Junos System Basics Configuration Guide*.

**Table 5: Standard Trace Options for Routing Protocols**

| Tracing Flag | Description |
| --- | --- |
| **all** | All operations |

**Table 5: Standard Trace Options for Routing Protocols** *(Continued)*

| Tracing Flag | Description |
|---|---|
| general | Normal operations and routing table changes |
| normal | Normal operations |
| policy | Policy operations and actions |
| route | Routing table changes |
| state | State transitions |
| task | Interface transactions and processing |
| timer | Timer usage |

## Monitor Trace File Messages Written in Near-Real Time

**IN THIS SECTION**

- Purpose | 171
- Action | 172

### Purpose

To monitor messages in near-real time as they are being written to a trace file.

## Action

To monitor messages in near-real time as they are being written to a trace file, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> monitor start  filename
```

## Sample Output

### command-name

```
user@host> monitor start isis
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
```

# Stop Trace File Monitoring

**IN THIS SECTION**

- Action | **173**
- Sample Output | **173**

## Action

To stop monitoring a trace file in near-real time, use the following Junos OS CLI operational mode command after you have started monitoring:

```
user@host                 monitor stop filename
```

## Sample Output

```
user@host> monitor start isis
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
monitor stop isis
user@host>
```

# Tracing Global Routing Protocol Operations

**IN THIS SECTION**

## 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.

### SEE ALSO

Junos OS Administration Library for Routing Devices

## Example: Tracing Global Routing Protocol Operations

This example shows how to list and view files that are created when you enable global routing trace operations.

### Requirements

You must have the **view** privilege.

### Overview

To configure global routing protocol tracing, include the `traceoptions` statement at the `[edit routing-options]` hierarchy level:

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

The flags in a `traceoptions flag` statement are identifiers. When you use the `set` command to configure a flag, any flags that might already be set are not modified. In the following example, setting the **timer** tracing flag has no effect on the already configured **task** flag. Use the `delete` command to delete a particular flag.

```
[edit routing-options traceoptions]
user@host# show
flag task;
user@host# set traceoptions flag timer
user@host# show
flag task;
```

```
  flag timer;
  user@host# delete traceoptions flag task
  user@host# show
  flag timer;
```

This example shows how to configure and view a trace file that tracks changes in the routing table. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.

**TIP**: To view a list of hierarchy levels that support tracing operations, enter the `help apropos traceoptions` command in configuration mode.

## Configuration

**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 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 Junos OS 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

**IN THIS SECTION**

- Verifying That the Trace Log File Is Operating | **180**

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
```

# 8
**CHAPTER**

# Configuration Statements

# access (Static Access Routes)

## Syntax

```
access {
    route  ip-prefix</prefix-length> {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        qualified-next-hop next-hop;
        tag tag-number
    }
```

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# access-internal (Static Access-Internal Routes)

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

# active

**IN THIS SECTION**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# aggregate (Routing)

**IN THIS SECTION**

## Syntax

```
aggregate {
    defaults {
```

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

## 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]
```

## 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*;

algorithm *number*—Specify an IGP algorithm for the aggregate route. IS-IS can advertise aggregate routes as SRv6 locators for locator summarization, In addition to the IPv6 aggregate prefix advertisement, IS-IS advertises the summary locator only if the node participates in the specified flexible algorithm.

- **Range**—Specify an algorithm from 0 through 255.

> **NOTE:**
> *IGP Algorithm Types*

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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation | 115

*How to Enable SRv6 Network Programming in IS-IS Networks*

# as-path (Routing Options)

**IN THIS SECTION**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# auto-export

**IN THIS SECTION**

## 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]
```

## 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.

## Default

`auto-export` is disabled by default. If you want to export routes between routing instances, include the `auto-export` CLI statement at the [`edit routing-options`] or at the [edit routing-instances *routing-instance-name* routing-options] hierarchy levels. You can use the `show route export` CLI command to verify the route export information.

## Options

**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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

**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.

# bfd

**IN THIS SECTION**

## Syntax

```
bfd {
      traceoptions {
          file filename <files number> <match regular-expression> <size size> <world-
readable | no-world-readable>;
          flag flag <flag-modifier> <disable>;
      }
      sbfd {
          pool value;
          local-discriminator value {
              minimum-receive-interval value;
              local-ip-address IP address;
          }
          local-discriminator-ip IP address;
```

```
        }
    }
```

## 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]
```

## Description

Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic. Configure seamless BFD (S-BFD) over non-colored static segment-routing LSPs with first-hop label resolution and use S-BFD as a fast mechanism to detect path failures and to trigger global convergence.

## Default

If you do not include this statement, no BFD tracing operations are performed.

## Options

**disable**—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file** *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the **/var/log** directory . We recommend that you place global routing protocol tracing output in the **routing-log** file.

**files** *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files

- **Default:** 2 files

**flag** *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. These are the BFD protocol tracing options:

- **adjacency**—Trace adjacency messages.

- **all**—Trace all options for BFD.

- **error**—Trace all errors.

- **event**—Trace all events.

- **issu**—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.

- **nsr-synchronization**—Trace NSR synchronization events.

- **packet**—Trace all packets.

- **pipe**—Trace pipe messages.

- **pipe-detail**—Trace pipe messages in detail.

- **ppm-packet**—Trace packet activity by periodic packet management (PPM).

- **state**—Trace state transitions.

- **timer**—Trace timer processing.

**match** *regular-expression*—(Optional) Regular expression for lines to be logged.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the trace file again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB

- **Range:** 10 KB through the maximum file size supported on your system

- **Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

**sbfd**—Seamless BFD parameters. In S-BFD, each responder might have multiple discriminators. The discriminators might be advertised by IGP to other routers, or they might be statically configured on these routers.

**local-discriminator** *value*—Local discriminator for seamless BFD responder.

**Range: 1 through 4294967295.**

- **local-ip-address** *IP address*—IPv4 source address.

- **minimum-receive-interval** *value*—Minimum receive interval for seamless BFD responder.

    **Range:** 10 through 255000 milliseconds.

**local-discriminator-ip** *IP address*—Local discriminator IP address for sBFD responder.

**pool** *value*—List of seamless BFD endpoints.

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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

# 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;
    }
    echo
        minimum-interval milliseconds
    echo-lite
        minimum-interval 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)]
```

## 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.

**echo**—(Optional) Enable BFD `echo` mode. This feature sends BFD echo packets back and forth from a neighboring device to ensure that a forwarding path is available. BFD `echo` mode only works if the neighboring device supports BFD.

**echo-lite**—(Optional) Enable BFD `echo-lite` mode. This feature performs the same funciton as BFD `echo` mode but it works even if the neighboring device does not support BFD.

**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.

## 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.

`echo` and `echo-lite` statements introduced in Junos OS Release 22.4R1.

### RELATED DOCUMENTATION

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

# 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

# color

**IN THIS SECTION**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0.

### RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation | **115**

aggregate (Routing) | **191**

generate | **255**

static (Routing Options) | **397**

preference (Routing Options) | **337**

# community (Routing Options)

**IN THIS SECTION**

- Syntax | **218**
- Hierarchy Level | **218**
- Description | **218**
- Default | **218**
- Options | **219**
- Required Privilege Level | **220**
- Release Information | **220**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

`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.

### RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation | **115**

aggregate (Routing) | **191**

generate | **255**

static (Routing Options) | **397**

# confederation

**IN THIS SECTION**

- Syntax | **221**
- Hierarchy Level | **221**
- Description | **221**

## Syntax

```
confederation confederation-autonomous-system members [ autonomous-systems ];
```

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

*Configuring GRE Tunnels for Layer 3 VPNs*

*Dynamic Tunnels Overview*

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]
```

## Description

Disable graceful restart.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Junos OS High Availability User Guide

# discard

**IN THIS SECTION**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# dynamic-tunnels

- Release Information | **231**

## Syntax

```
dynamic-tunnels tunnel-name {
    bgp-signal;
    destination-networks prefix;
    gre;
    ipip
    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 (BGP-SIGNAL | GRE | UDP | V4oV6);
    }
  v4ov6 ipv6-anycast-source-duplication;
}
```

## 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]
```

## 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

**bgp-signal**      Enable the creation of a tunnel signaled by BGP.

**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.

**ipip**      Enable dynamic IP in IP encapsulation type tunnel mode for IPv4

- Values:

  - `full-resolved-next-hop-based-tunnel`—Enable fully resolved 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

`bgp-signal` option introduced in Junos OS Release 17.4R1.

### 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Example: Load Balancing BGP Traffic

# export-rib

IN THIS SECTION

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table  |  **13**

*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  |  **263**

passive (Routing Options)  |  **329**

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

*Configuring Alternate Backup Paths Using Fate Sharing*

MPLS Applications User Guide

# filter

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

*Configuring Forwarding Table Filters*

*Applying Forwarding Table Filters*

# firewall-install-disable

**IN THIS SECTION**

- Syntax | **241**
- Hierarchy Level | **241**

## 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Releases 12.1X48 and 12.3.

### RELATED DOCUMENTATION

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

*Understanding BGP Flow Routes for Traffic Filtering*

# flow

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

`term-order` statement introduced in Junos OS Release 10.0

`firewall-install-disable` statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.

### RELATED DOCUMENTATION

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

*Understanding BGP Flow Routes for Traffic Filtering*

# forwarding-table

### IN THIS SECTION

- Syntax | **245**
- Hierarchy Level | **246**
- Description | **246**
- Options | **246**
- Required Privilege Level | **247**
- Release Information | **247**

## 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
    multicast-composite-delete-defer-timeout;
    route-origin-proto-tlv;
    srv6-chain-merge;
    unicast-reverse-path (active-paths | feasible-paths);
}
```

## Hierarchy Level

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

## Description

Configure information about the routing device's forwarding table.

The remaining statements are explained separately. See CLI Explorer.

## Options

| | |
|---|---|
| **route-origin-proto-tlv** | Configure this option to include the origin protocol information in route messages. The origin protocol is the protocol from which the route entry is learned. |
| **srv6-chain-merge** | Enable merging of srv6 chain nexthops. |
| **multicast-composite-delete-defer-timeout** | Use this command to delay the deletion of a composite next-hop from the FIB (in the event of dropped multicast packets following a downstream interface change ). You can view the set time by running the **show krt state** command.<br><br>Range: 1 to 3600 seconds |

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

`srv6-chain-merge` option introduced in Junos OS Release 21.3R1.

`multicast-composite-delete-defer-timeout` option introduced in Junos OS Release 22.1R2 and Junos OS Evolved Release 22.1R2

`route-origin-proto-tlv` option introduced in Junos OS Release 22.4R1 and Junos OS Evolved Release 22.4R1.

### RELATED DOCUMENTATION

Example: Load Balancing BGP Traffic

# forwarding-options

**IN THIS SECTION**

## 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;
```

```
            }
        }
    }
    storm-control {
        enhancedenhanced;
    }
    filter (VLANs) {
        input filter-name;
        output filter-name;
    }
    sflow {
        egress-multicast enable;
        egress-multicast max-replication-rate rate;
    }
    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)]
```

## 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 primary 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 primary 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 6 on page 252 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 6: 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.

| 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.

**egress-multicast enable**

Enable egress sFlow sampling of known multicast traffic.

**egress-multicast max-replication-rate** *rate*

Specify the maximum sFlow sample replication rate per FPC for known multicast traffic.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

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.

### RELATED DOCUMENTATION

Understanding the Unified Forwarding Table

Example: Configuring a Unified Forwarding Table Custom Profile

*Configuring Traffic Forwarding and Monitoring*

# full

**IN THIS SECTION**

- See | **254**

## See

# 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]
```

# 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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

IN THIS SECTION

- Syntax | 260
- Hierarchy Level | 260
- Description | 260

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

*Example: Configuring Route Resolution on PE Routers*

# import-policy

**IN THIS SECTION**

-
-
-
-
-
-

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# import-rib

**IN THIS SECTION**

- Syntax | **263**
- Hierarchy Level | **263**
- Description | **264**
- Options | **264**
- Required Privilege Level | **264**
- Release Information | **264**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

`no-attrset` option introduced in Junos OS Release 10.4.

*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

# indirect-next-hop

**IN THIS SECTION**

- Syntax | **267**
- Hierarchy Level | **267**
- Description | **268**
- Default | **268**
- Options | **268**
- Required Privilege Level | **268**
- Release Information | **269**

## 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 8.2.

# indirect-next-hop-change-acknowledgements

**IN THIS SECTION**

## 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]
```

## 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 PFEs 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 PFEs 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.

## Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine | **137**

krt-nexthop-ack-timeout | **288**

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# install (Routing Options)

**IN THIS SECTION**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# interface (Multicast Static Routes)

**IN THIS SECTION**

- Syntax | **280**
- Hierarchy Level | **281**
- Description | **281**
- Options | **281**
- Required Privilege Level | **282**
- Release Information | **282**

## 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 8.1.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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

RELATED DOCUMENTATION

Example: Populating a Routing Table Created by Virtual Router Configuration

*Example: Configuring Multiple Routing Instances of OSPF*

passive (Routing Options) | **329**

# 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]
```

## Description

Configure a Juniper Entropy Beacon (JEB) server. JEB can be used to feed high quality entropy over the network from a SRX345 Firewall 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

## Release Information

Statement introduced in Junos OS Release 19.1R1.

# krt-nexthop-ack

## Syntax

```
krt-nexthop-ack;
```

## Hierarchy Level

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

## Description

When you configure this option, the routing protocol process (rpd) waits for the next hop to get acknowledged by PFE before using it for a route.

Currently, only BGP-labeled routes and RSVP routes support this option. All other routes will ignore this option.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 22.1.

# krt-nexthop-ack-timeout

**IN THIS SECTION**

## Syntax

```
krt-nexthop-ack-timeout interval;
```

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 12.2.

# longest-match (Static Routes)

**IN THIS SECTION**

## 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],
```

## 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 to binary and compare them.

**Table 7: 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.

## Release Information

Statement introduced in Junos OS Release 15.1.

# lsp-next-hop (Static Routes)

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 8.0.

*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]
```

## 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, programmed, or interface routes.

> **NOTE**: Static, programmed [that is, programmable routing protocol process (prpd) configured], and interface routes contribute to the overall route total count used in the maximum-prefix calculation. The number of such routes affects the number of dynamic routes that are accepted in to the virtual routing instance. In the configuration example shown below, if VRF1 contains 98 static routes and 2 interface routes, then a maximum of 900 BGP routes (that is, 1000 - 98 - 2 = 900) are allowed, with a syslog warning shown at a total of 800 dynamic routes, which is 80%.
>
> ```
> set routing-instances VRF1 routing-options maximum-prefixes 1000
> set routing-instances VRF1 routing-options maximum-prefixes threshold 80
> ```
>
> If the number of static, programmed, or interface routes increases, then it can affect the maximum number of dynamic routes that are accepted. Note that the static, programmed, or interface routes are not subject to the configured maximum-prefixes limit and can have a negative or undesired effect on the dynamic routes.

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.

## Release Information

Statement introduced in Junos OS Release 8.0.

*Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

# med-igp-update-interval

**IN THIS SECTION**

## Syntax

```
med-igp-update-interval minutes;
```

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 9.0

# metric

**IN THIS SECTION**

## Syntax

```
metric route-cost;
```

## Hierarchy Level

```
[edit routing-options access route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# metric (Qualified Next Hop on Static Route)

**IN THIS SECTION**

- Syntax | **307**
- Hierarchy Level | **307**
- Description | **308**
- Options | **308**
- Required Privilege Level | **308**
- Release Information | **308**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

qualified-next-hop (Static Routes) | **347**

static (Routing Options) | **397**

# multicast

IN THIS SECTION

- Syntax | **309**

## Syntax

```
multicast {
    asm-override-ssm;
    backup-pe-group group-name {
        backups [ addresses ];
        local-address address;
    }
    cont-stats-collection-interval interval;
    flow-map flow-map-name {
        bandwidth (bps | adaptive);
        forwarding-cache {
            timeout (never non-discard-entry-only | minutes);
        }
        policy [ policy-names ];
        redundant-sources [ addresses ];
    }
    forwarding-cache {
        threshold suppress value <reuse value>;
        timeout minutes;
    }
    interface interface-name {
        enable;
        maximum-bandwidth bps;
        no-qos-adjust;
        reverse-oif-mapping {
            no-qos-adjust;
        }
        subscriber-leave-timer seconds;
    }
    local-address address
    omit-wildcard-address
```

```
    pim-to-igmp-proxy {
        upstream-interface [ interface-names ];
    }
    pim-to-mld-proxy {
        upstream-interface [ interface-names ];
    }
    rpf-check-policy [ policy-names ];
    scope scope-name {
        interface [ interface-names ];
        prefix destination-prefix;
    }
    scope-policy [ policy-names ];
    ssm-groups [ addresses ];
    ssm-map ssm-map-name {
        policy [ policy-names ];
        source [ addresses ];
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <disable>;
    }
}
```

## Hierarchy Level

```
[edit dynamic-profiles profile-name routing-options],
[edit dynamic-profiles profile-name routing-instances routing-instance-name routing-options],
[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]
```

## Description

Configure multicast routing options properties. Note that you cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the scope statement does apply individually 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

**interface** and **maximum-bandwidth** statements introduced in Junos OS Release 8.3.

**interface** and **maximum-bandwidth** statements introduced in Junos OS Release 9.0 for EX Series switches.

Statement added to **[edit dynamic-profiles routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options]** hierarchy levels in Junos OS Release 9.6.

RELATED DOCUMENTATION

*Examples: Configuring Administrative Scoping*

*Example: Configuring the Multicast Forwarding Cache*

*Example: Configuring a Multicast Flow Map*

*Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

indirect-next-hop

# next-hop (Access)

## Syntax

```
next-hop next-hop;
```

## Hierarchy Level

```
[edit routing-options access route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# next-hop (Access Internal)

**IN THIS SECTION**

## Syntax

```
next-hop next-hop;
```

## Hierarchy Level

```
[edit routing-options access-internal route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# no-delegate-processing

## Syntax

```
no-delegate-processing;
```

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# no-hierarchical-ecmp

## Syntax

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

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 15.1X53-D210.

# 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.

## Description

For routing platforms with two Routing Engines, configure a primary 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.

## Release Information

Statement introduced in Junos OS Release 8.4.

### 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 13.2.

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

### RELATED DOCUMENTATION

*Configuring the Unified Forwarding Table on Switches*

# options (Routing Options)

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

Junos OS Administration Library for Routing Devices

# per-route-accounting

## Syntax

```
per-route-accounting
```

## Hierarchy Level

```
[edit routing-options flow]
[edit routing-instance flow]
[edit routing-options-rib-inet6.0 flow]
[edit routing-intance-rib-inet6.0 flow]
```

## Description

Earlier to Junos OS Evolved Release 21.1R1, by default, in the usual BGP FlowSpec filter, an individual counter was added for each FlowSpec term configured. As part of optimization of FlowSpec, starting from Junos OS Evolved Release 21.1R1 this behavior is changed, and counter action is not added by default. This will create more room for filter terms compression. When the term counters are not programmed, the traffic matching against a specific flow term is not accounted. You can use the new configuration statement `per-route-accounting` under the `[edit routing-options flow]` hierarchy level to enable creation of term counters for flowspec at global level on PTX10001-36MR, PTX10004, and PTX10008.

## Default

The term counters are disabled by default.

## Required Privilege Level

routing-options flow-To enable the new configuration statement.

# 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 13.3.

# p2mp-lsp-next-hop

**IN THIS SECTION**

- Syntax | **328**
- Hierarchy Level | **328**
- Description | **328**
- Required Privilege Level | **328**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# passive (Routing Options)

**IN THIS SECTION**

- See | **329**

## See

# policy (Aggregate and Generated Routes)

**IN THIS SECTION**

- Syntax | **330**
- Hierarchy Level | **330**
- Description | **330**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

**RELATED DOCUMENTATION**

# ppm

**IN THIS SECTION**

## Syntax

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

## Syntax (QFX5110, QFX5120, QFX5200, and QFX5210)

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

## Hierarchy Level

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

## Description

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 spanning tree protocols work seamlessly on chassis or virtual chassis based devices during ISSU/NSSU, NSR and GRES. Recommended value is 120 seconds. |
| `inline-processing-enable` | Enable inline processing. Inline processing allows you to set BFD timers to less than one second. On platforms that support hardware-assisted inline BFD, use this option to enable it. |

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.4.

`inline-processing-enable` option introduced in Junos OS Release 19.4R1.

# precision-timers-max-period

**IN THIS SECTION**

## Syntax

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

## Hierarchy Level

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

## 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 primary 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 primary 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.

## Release Information

Statement introduced in Junos OS Release 16.1.

# preference (Access)

## Syntax

```
preference route-distance;
```

## Hierarchy Level

```
[edit routing-options access route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# preference (Routing Options)

**IN THIS SECTION**

-

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# prefix

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

**RELATED DOCUMENTATION**

*Examples: Configuring Administrative Scoping*

*Example: Creating a Named Scope for Multicast Scoping*

multicast | 308

# prefix-65-127-disable

## Syntax

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

## Hierarchy Level

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

## 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.

## Release Information

Statement introduced in Junos OS Release 13.2X51-D15.

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.

# qualified-next-hop (Access)

## Syntax

```
qualified-next-hop next-hop;
```

## Hierarchy Level

```
[edit routing-options access route ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# qualified-next-hop (Access-Internal)

## Syntax

```
qualified-next-hop next-hop;
```

## Hierarchy Level

```
[edit routing-options access-internal route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# qualified-next-hop (Static Routes)

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# resolution

## Syntax

```
resolution {
    (preserve-nexthop-hierarchy | no-preserve-nexthop-hierarchy);
    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 ];
        mpls-import [ policy-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]
```

## 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" on page 355 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

| | |
|---|---|
| **disable-auto-resolve-on-scheme-changes** | (Optional) Service routes are not re-resolved automatically when resolution schemes are created, deleted, or changed. However, you can manually trigger route re-resolution on resolution scheme changes. |
| **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. |
| **preserve-nexthop-hierarchy** | Preserve nexthop hierarchy when installing nexthop to forwarding plane. |

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

inet6-import and inet6-resolution-ribs options added in Junos OS Release 17.2R1.

**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

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# resolve

IN THIS SECTION

-
-
-
-
-
-

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

static (Routing Options) **| 397**

# restart-duration

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

Junos OS High Availability User Guide

# retain

**IN THIS SECTION**

-
-

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# rib (General)

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

`policy-multipath` statement introduced in Junos Release 19.1R1 for all platfroms.

# rib (Route Resolution)

**IN THIS SECTION**

- Syntax | **369**
- Hierarchy Level | **370**
- Description | **370**
- Required Privilege Level | **370**
- Release Information | **370**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# 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]
```

# 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

Example: Exporting Specific Routes from One Routing Table Into Another Routing Table | 13

rib-group (Routing Options) | 371

# route (Access)

- Release Information | **377**

## 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]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# route (Access-Internal)

**IN THIS SECTION**

## Syntax

```
route  ip-prefix</prefix-length> {
    next-hop next-hop;
    qualified-next-hop next-hop;
}
```

## Hierarchy Level

```
[edit routing-options access-internal]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# route-distinguisher-id

## Syntax

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

## Hierarchy Level

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

## 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, you must 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*Enabling Flow Aggregation*

Junos OS Services Interfaces Library for Routing Devices

# router-id

**IN THIS SECTION**

-

## 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]
```

## 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 OSPF3 or IPv6 based BGP peering in a routing instance, you must configure a router identifier (`router-id`) in the routing instance itself. In other words, the `router-id` used for the main routing instance is **not** inherited by other routing instances. Even if you run *only* OSPF3 or IPv6 BGP peering in a routing instance, a 32-bit `router-id` must be configured in the instance. This is because IPv6 routing protocols use the `router-id` for handshaking. **The router ID must be configured as a 4 octet (32 bit) unsigned non-zero integer value.**
> **It's often convenient to use an IPv4 address as the router ID. However, a valid IPv4 address is not required. The RID does not have to be a routable IPv4 address. You can configure any 32-bit value that is unique within the routing domain.** If you do not configure the `router-id` in an IPv6 OSPF or BGP routing instance the IPv6 protocols will use an invalid value for the router ID (`0.0.0.0`) and the adjacency and connections will fail

When you configure a router ID with an IP address that differs from the BGP local address for EVPN routing instances, you must also configure a policy that exports the local IP address for BGP routes. This ensures that EVPN traffic will be properly routed. For more information on exporting BGP routes, see *Basic BGP Routing Policies*.

## Options

*address*—IP address of the routing device.

- **Default:** Address of the first interface encountered by Junos OS

Starting in Junos OS Release 20.4R1, you can establish a BGP connection with the BGP identifier that is a 4-octet, unsigned, and non-zero integer and it needs to be unique only within the Autonomous System (AS) per *RFC6286, Autonomous-System-Wide Unique BGP Identifier for BGP-4*

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

# routing-options

**IN THIS SECTION**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

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

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*Configuring Multicast Snooping*

# scope-policy

### IN THIS SECTION

- Syntax | **389**
- Hierarchy Level | **389**

## 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.

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# spring-te-lsp-next-hop

## Syntax

```
spring-te-lsp-next-hop{
    <nexthop>
    lsp-source {
        bgp;
        dtm;
        pcep;
        static;
        }
}
```

## Hierarchy Level

```
[edit logical-systems name routing-options static route destination],
[edit logical-systems name routing-options rib rib name static route destination]
[edit routing-options static route destination],
[edit routing-options rib rib name static route destination]
```

## Description

The configuration allows you to resolve static route over SR-TE tunnels.

## Options

**nexthop**—Name of the SR-TE tunnel to reach destination address.

**lsp-source**—Source of tunnel creation and the following source types are currently supported. If the source is not configured, by default, it takes the next hop as static. So you must configure the source type if LSP is dtm, bgp, or pcep.

- **bgp**

- **dtm**

- **pcep**

- **static**

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 21.2R1.

# 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

# source-routing

**IN THIS SECTION**

## Syntax

```
source-routing {
    (ip | ipv6)
}
```

## Hierarchy Level

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

## 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.

## 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.

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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

*Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

# static (Routing Options)

**IN THIS SECTION**

- Syntax | **398**

## 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;
        spring-te-lsp-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 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]
```

## 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.

> **NOTE**: With the fix for PR 1383419, configuring a static route with a next-table action on QFX5xxx/EX4xxx switches results in software forwarding when the prefix length is less than 16 bits for IPv4 or 64 bits for IPv6. To avoid possible performance impact you should configure multiple static routes with prefixes equal to or longer than 16 and 64 bits for IPv6 and IPv6, respectively

- **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.

## Release Information

Statement introduced before Junos OS Release 7.4.

Support for BFD authentication introduced in Junos 9.6.

Support for BFD authentication introduced in Junos 9.6 for EX Series switches.

# tag (Access)

## Syntax

```
tag  tag-number;
```

## Hierarchy Level

```
[edit routing-options access route  ip-prefix</prefix-length>]
```

## 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.

## Release Information

Statement introduced in Junos OS Release 10.1.

# tag (Routing Options)

**IN THIS SECTION**

- Syntax | **406**

## 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

Example: Summarizing Static Routes Through Route Aggregation | **115**

aggregate (Routing) | **191**

generate | **255**

# 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)]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

*Examples: Configuring the Multicast Forwarding Cache*

*show multicast forwarding-cache statistics*

# traceoptions

**IN THIS SECTION**

## 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]
```

## Description

Define tracing operations that track all routing protocol functionality in the routing device. You can also trace clock synchronization information using this command.

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:** *x***k** to specify KB, *x***m** to specify MB, or *x***g** to specify GB

- **Range:** 10 KB through the maximum file size supported on your system

- **Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

## Required Privilege Level

routing and trace—To view this statement in the configuration.

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

## 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.

**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.

**nsr-synchronization** flag for MSDP added in Junos OS Release 12.1.

### RELATED DOCUMENTATION

Example: Tracing Global Routing Protocol Operations | **175**

# unicast-reverse-path

**IN THIS SECTION**

- Syntax | **414**
- Hierarchy Level | **415**
- Description | **415**
- Options | **415**
- Required Privilege Level | **415**
- Release Information | **415**

## 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]
```

## 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.

## Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances added in Junos OS Release 8.3.

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

RELATED DOCUMENTATION

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

*Enabling Unicast Reverse-Path Forwarding Check for VPNs*

# 9
**CHAPTER**

## Operational Commands

# clear bfd adaptation

## Syntax

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

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# clear bfd session

**IN THIS SECTION**

## 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>
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

*show bfd session*

# show bfd session

**IN THIS SECTION**

## 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>
```

## 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

## Output Fields

[Table 8 on page 425](#) describes the output fields for the `show bfd session` command. Output fields are listed in the approximate order in which they appear.

**Table 8: 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 |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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` |
| 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` |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| algo | BFD authentication algorithm being used for a specific client: `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, `meticulous-keyed-sha-1`, or `simple-password`.<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><br>• `CtlExpire`—Control detection time expired<br><br>• `EchoExpire`—Echo detection time expired<br><br>• `NbrSignal`—Neighbor signalled session down<br><br>• `FwdPlaneReset`—Forwarding plane reset<br><br>• `PathDown`—Path down<br><br>• `ConcatPathDown`—Concatenated path down<br><br>• `AdminDown`—Administratively down | `detail extensive` |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Remote diagnostic | Remote diagnostic information about failing BFD sessions.<br><br>Following are the expected values for Remote Diagnostic output field:<br><br>• `None`—No diagnostic<br><br>• `CtlExpire`—Control detection time expired<br><br>• `EchoExpire`—Echo detection time expired<br><br>• `NbrSignal`—Neighbor signalled session down<br><br>• `FwdPlaneReset`—Forwarding plane reset<br><br>• `PathDown`—Path down<br><br>• `ConcatPathDown`—Concatenated path down<br><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 the device has a backup Routing Engine installed and the BFD session has been replicated in 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` |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| RX interval | Minimum required receive interval. | `extensive` |
| Local min TX interval | Minimum amount of time, in seconds, between control packet transmissions on the local system. | `extensive` |
| Local min RX interval | Minimum amount of time, in seconds, between control packet detections on the local system. | `extensive` |
| Remote min TX interval | Minimum amount of time, in seconds, between control packet transmissions on the remote system. | `extensive` |
| Remote min TX interval | Minimum amount of time, in seconds, between control packet detections on the remote system. | `extensive` |
| 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 |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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` |
| Authentication | Summary status of BFD authentication:<br><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><br>• `keychain`—Name of the security authentication keychain associated with the specified BFD session.<br><br>• `algo`—BFD authentication algorithm being used: `keyed-md5`, `keyed-sha-1`, `meticulous-keyed-md5`, `meticulous-keyed-sha-1`, or `simple-password`.<br><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><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 |

**Table 8: show bfd session Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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><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 extensive

```
user@host> show bfd session extensive                                  Detect
Transmit
Address                State     Interface    Time    Interval  Multiplier
10.9.1.33              Up        et-1/0/1.0   6.000   2.000      3
 Client PRPD:6, TX interval 1.000, RX interval 1.000
 Session up time 00:00:01
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Session type: Single hop BFD (Echo Lite)
 Local discriminator 16
 Echo TX interval 2.000, echo detection interval 6.000
  Session ID: 0
```

```
1 sessions, 1 clients
Cumulative transmit rate 0.5 pps, cumulative receive rate 0.5 pps
```

## 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.

Option `client` introduced in Junos OS Release 12.3R3.

Option `subscriber` introduced in Junos OS Release 15.1 for the MX Series.

### RELATED DOCUMENTATION

*clear bfd session*

Understanding BFD for Static Routes for Faster Network Failure Detection

*Understanding BFD for OSPF*

Understanding BFD for BGP

Understanding Bidirectional Forwarding Detection Authentication for PIM

Configuring BFD for PIM

Understanding BFD for IS-IS

# show as-path

## Syntax

```
show as-path
<brief | detail>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show as-path
<brief | detail>
```

## 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

## Output Fields

Table 9 on page 435 lists the output fields for the `show as-path` command. Output fields are listed in the approximate order in which they appear.

**Table 9: 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 |

**Table 9: show as-path Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| *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><br>• `I`—IGP.<br><br>• `E`—EGP.<br><br>• `?`—Incomplete; typically, the AS path was aggregated.<br><br>• `Atomic`—Route is an aggregate of several route prefixes.<br><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
 ...
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show as-path domain

## Syntax

```
show as-path domain
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show as-path domain
```

# 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

# Output Fields

Table 10 on page 441 lists the output fields for the `show as-path domain` command. Output fields are listed in the approximate order in which they appear

**Table 10: 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. |

**Table 10: show as-path domain Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *Number* Paths | Number of known AS paths. |
| Flags | Information about the AS path:<br><br>• **ASLoop**—Path contains an AS loop.<br><br>• **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute.<br><br>• **Local**—Path was created by local aggregation.<br><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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show as-path summary

## Syntax

```
show as-path summary
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show as-path summary
```

## 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

## Output Fields

lists the output fields for the `show as-path summary` command. Output fields are listed in the approximate order in which they appear.

**Table 11: 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. |

## Syntax

```
show chassis forwarding-options
```

## 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 Primary, not for the other members. Values remain the same across all members. All configuration changes for the Unified Forwarding Table are made through the Primary.

## Required Privilege Level

view

# Output Fields

lists the output fields for the `show chassis forwarding-options` command. Output fields are listed in the approximate order in which they appear.

**Table 12: show chassis forwarding-options Output Fields**

| Field Name | Field Description |
|---|---|
| profile name | Name of profile configured:<br><br>• `custom-profile` (QFX5200 only)<br><br>• `l2-profile-one`<br><br>• `l2-profile-three` (default)<br><br>• `l2-profile-two`<br><br>• `l3-profile`<br><br>• `lpm-profile`<br><br>• `vxlan-gbp-profile` (EX4100, EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches)<br><br>• `vxlan-gbp-l2-profile` (EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches)<br><br>• `vxlan-gbp-l3-profile` (EX4400, EX4650, and QFX5120-32C and QFX5120-48Y switches) |
| 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. |

**Table 12: show chassis forwarding-options Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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 (vxlan-gbp-l2-profile on EX4400 Series switch)**

```
user@host> show chasiss forwarding-options
UFT Configuration:
vxlan-gbp-l2-profile
num-65-127-prefix = 1K
--------------Bank details for various types of entries-----------------
Entry type                    Dedicated Bank Size(K)    Shared Bank Size(K)
L2(mac)                       16                        16 * num shared banks
L3(unicast & multicast)       16                        16 * num shared banks
Exact Match                   0                         8 * num shared banks
```

**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
```

## Release Information

Command introduced in Junos OS Release 13.2

Support added to QFX5200 switches in Junos OS Release 15.1X53-D30

*Configuring the Unified Forwarding Table on Switches*

*Example: Configuring a Unified Forwarding Table Custom Profile*

# show interfaces routing summary

**IN THIS SECTION**

## Syntax

```
show interfaces routing summary
<interface-name>
<logical-system (all | logical-system-name)>
```

## 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 Routing Devices. For information about related operational mode commands for routing instances and protocols, see the CLI Explorer.

## Required Privilege Level

view

# Output Fields

lists the output fields for the `show interfaces routing summary` command. Output fields are listed in the approximate order in which they appear.

**Table 13: 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route

**IN THIS SECTION**

## 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>
<te-ipv6-prefix-ipv6-addr te-ipv6-prefix-ipv6-addr>
<te-ipv6-prefix-node-iso te-ipv6-prefix-node-iso>
<rib-sharding (main | rib-shard-name)>
```

## Syntax (EX Series Switches)

```
show route
<all>
<destination-prefix>
<private>
```

## 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. |
| **programmed detail** | (Optional) Display API-programmed routes. |
| **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. |
| **te-ipv6-prefix-ipv6-addr** *te-ipv6-prefix-ipv6-addr* | (Optional) Filter IPv6 node addresses from the traffic-engineering IPv6 prefix. |
| **te-ipv6-prefix-node-iso** *te-ipv6-prefix-node-iso* | (Optional) Filter IPv6 routes with the specified ISO circuit ID in the traffic-engineering IPv6 prefix. |
| **rib-sharding (main |** *rib-shard-name*) | (Optional) Display the rib shard name. |

## Required Privilege Level

view

## Output Fields

describes the output fields for the `show route` command. Output fields are listed in the approximate order in which they appear.

**Table 14: 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><br> • `active` (routes that are active). <br><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><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><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><br> • `hidden` (routes that are not used because of a routing policy). |

**Table 14: 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><br>• *MPLS-label* (for example, 80001).<br><br>• *interface-name* (for example, ge-1/0/2).<br><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><br>  • *neighbor-address*—Address of the neighbor.<br><br>  • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord.<br><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><br>  • *vc-id*—Virtual circuit identifier.<br><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><br>• +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br><br>• - —A hyphen indicates the last active route.<br><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><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. |

**Table 14: show route Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *weeks:days hours:minutes:seconds* | How long the route been known (for example, 2w4d 13:11:14, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds). |
| metric | Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value. |
| localpref | Local preference value included in the route. |
| from | Interface from which the route was received. |
| AS path | AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: <br><br> • `I`—IGP. <br><br> • `E`—EGP. <br><br> • `?`—Incomplete; typically, the AS path was aggregated. <br><br> When AS path numbers are included in the route, the format is as follows: <br><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><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><br> • `( )`—Parentheses enclose a confederation. <br><br> • `( [ ] )`—Parentheses and brackets enclose a confederation set. <br><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 14: show route Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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-Over 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><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><br>• Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br><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><br>• Valid—Indicates that the prefix and autonomous system pair are found in the database. |
| to | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.<br><br>If the destination is Discard, traffic is dropped. |

**Table 14: show route Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| via | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:<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><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><br>• *lsp-path-name*—Name of the LSP used to reach the next hop.<br><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

The following sample output shows route hierarchy for translation route.

```
user@host> show route 10.1.1.1

C1.inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.1/32       *[PRPD/10] 00:16:50, metric 2
                   >  to 192.0.2.2 via ge-0/0/1.0
```

## show route forwarding-table matching 10.1.1.1

```
user@host> show route forwarding-table matching 10.1.1.1
Routing table: C1.inet
Internet:
Destination        Type RtRef Next hop        Type Index    NhRef Netif

10.1.1.1/32        user    0                   indr  1048574    4


                                               comp     624    2
```

## show route 10.1.1.1 extensive expanded-nh

```
user@host> show route 10.1.1.1extensive expanded-nh
C1.inet
C1.inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
10.1.1.1/32 (1 entry, 1 announced)
Installed-nexthop:
Indr (0xc5c207c) ::44.0.0.1
  Krt_inh (0xc6fd004) Index:1048574 PNH: ::44.0.0.1
    Translate-comp (0xc5c2144) Index:624 v4tov6 src ::22.0.0.1 dest ::44.0.0.1
```

## show route te-ipv6-prefix-ipv6-addr

```
user@host> show route te-ipv6-prefix-ipv6-addr 10::10

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 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: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet6.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

lsdist.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both


IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/1216
                     *[IS-IS/15] 00:07:58
                         Fictitious
```

## show route te-ipv6-prefix-node-iso

```
user@host> show route te-ipv6-prefix-node-iso 0100.0100.0100.00

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 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: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

inet6.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

lsdist.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/
1216
                     *[IS-IS/15] 00:08:46
                         Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:21:0:1::1/128 } ISIS-L1:0 }/
1216
                     *[IS-IS/15] 00:08:46
                         Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:abcd::128:207:200:16/128 } ISIS-
L1:0 }/1216
                     *[IS-IS/15] 00:08:46
                         Fictitious
```

### 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 programmed detail**

```
user@host> show route programmed detail
inet.0: 36 destinations, 37 routes (36 active, 0 holddown, 0 hidden)
100.75.1.0/27 (2 entries, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 0
                Address: 0xcc38a10
                Next-hop reference count: 1
                Next hop: 100.30.1.2 via ge-0/0/2.0 weight 0x1, selected
                Session Id: 0x0
                Next hop: via fti0.1001 weight 0x8001
                Session Id: 0x0
              State: <Active Int NSR-incapable Programmed>
                Age: 37
                Validation State: unverified
                Announcement bits (1): 0-KRT
                AS path: I
```

## Release Information

Command introduced before Junos OS Release 7.4.

Option `private` introduced in Junos OS Release 9.5.

Option `private` introduced in Junos OS Release 9.5 for EX Series switches.

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.

Understanding IS-IS Configuration

Verifying and Managing Junos OS Enhanced Subscriber Management

# show route active-path

**IN THIS SECTION**

## 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>
```

## 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

## 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`.

**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
```

## Release Information

Command introduced in Junos OS Release 8.0.

# show route all

**IN THIS SECTION**

- Syntax | **476**
- Syntax (EX Series Switches) | **476**
- Description | **476**
- Options | **476**
- Required Privilege Level | **476**
- Output Fields | **477**
- Sample Output | **477**
- Release Information | **478**

## Syntax

```
show route all
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route all
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*show route brief*

*show route detail*

# show route aspath-regex

**IN THIS SECTION**

## Syntax

```
show route aspath-regex regular-expression
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route aspath-regex regular-expression
```

## 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

## 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)

...
```

# Release Information

Command introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

*Example: Using AS Path Regular Expressions*

# show route best

## Syntax

```
show route best destination-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route best destination-prefix
<brief | detail | extensive | terse>
```

## 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

## 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 484.

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*show route brief*

*show route detail*

*show route extensive*

*show route terse*

# show route brief

**IN THIS SECTION**

- Syntax | **487**
- Syntax (EX Series Switches) | **487**
- Description | **487**
- Options | **487**
- Required Privilege Level | **488**

## Syntax

```
show route brief
<destination-prefix>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route brief
<destination-prefix>
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*show route*

*show route all*

*show route best*

# show route cumulative

**IN THIS SECTION**

## Syntax

```
show route cumulative
<fabric>
<logical-system (all | logical-system-name)>
<vpn-family (inet.0 | inet6.0)>
```

## 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

## Output Fields

| Field Name | Field Description |
|---|---|
| destinations | Number of destinations for which there are VRF routes in the routing table. |

*(Continued)*

| Field Name | Field Description |
|------------|-------------------|
| `routes` | Number of VRF routes in the routing table:<br><br>• `active`—Number of routes that are active.<br><br>• `holddown`—Number of VRF routes that are in the hold-down state before being declared inactive.<br><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)
```

## Release Information

Command introduced in Junos OS Release 13.3.

# show route detail

## Syntax

```
show route detail
<destination-prefix>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route detail
<destination-prefix>
```

## 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

## Output Fields

describes the output fields for the `show route detail` command. Output fields are listed in the approximate order in which they appear.

describes all possible values for the Next-hop Types output field.

describes all possible values for the State output field. A route can be in more than one state (for example, `<Active NoReadvrt Int Ext>`).

describes the possible values for the Communities output field.

**Table 15: 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. |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *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><br>• `holddown` (routes that are in the pending state before being declared inactive)<br><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><br>• *MPLS-label* (for example, 80001).<br><br>• *interface-name* (for example, ge-1/0/2).<br><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><br>  • *neighbor-address*—Address of the neighbor.<br><br>  • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.<br><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><br>  • *vc-id*—Virtual circuit identifier.<br><br>  • *source*—Source of the advertisement: `Local` or `Remote`.<br><br>  • *source*—Source of the advertisement: `Local` or `Remote`. |

**Table 15: 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><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><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 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| [*protocol*, *preference*] | Protocol from which the route was learned and the preference value for the route.<br><br>• +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br><br>• - —A hyphen indicates the last active route.<br><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><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.<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:<br><br>• Both Signed `Preference2` values<br><br>    • Route A = -101<br><br>    • Route B = -156<br><br>  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.<br><br>• Unsigned `Preference2` values<br><br>  Now consider both unsigned `Preference2` values:<br><br>    • Route A = 4294967096<br><br>    • Route B = 200<br><br>  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.<br><br>• Combination of signed and unsigned `Preference2` values |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| | 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:<br><br>• Route A = -200<br><br>• Route B = 200<br><br>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. |
| 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 16 on page 504. |
| 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. |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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><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 17 on page 506. |
| Local AS | AS number of the local routing device. |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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 15: 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><br>• I—IGP.<br><br>• E—EGP.<br><br>• Recorded—The AS path is recorded by the sample process (sampled).<br><br>• ?—Incomplete; typically, the AS path was aggregated.<br><br>When AS path numbers are included in the route, the format is as follows:<br><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><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><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><br>• ( )—Parentheses enclose a confederation.<br><br>• ( [ ] )—Parentheses and brackets enclose a confederation set.<br><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 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| validation-state | (BGP-learned routes) Validation status of the route:<br><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><br>• `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br><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><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. |
| 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. |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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 18 on page 509 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 may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. |

**Table 15: show route detail Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| Statistics ID Group | Indicates the Kernel ID number and Statistics ID number. |
| Statistics | Indicates the number of packets and data transferred. |

**Table 16: Next-hop Types Output Field Values**

| Next-Hop Type | Description |
| --- | --- |
| Broadcast (bcast) | Broadcast next hop. |
| 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. |

**Table 16: 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><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><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 17: State Output Field Values**

| Value | Description |
| --- | --- |
| Accounting | Route needs accounting. |
| Active | Route is active. |
| Always Compare MED | Path with a lower multiple exit discriminator (MED) is available. |
| AS path | Shorter AS path is available. |
| Cisco Non-deterministic MED selection | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| Clone | Route is a clone. |
| Cluster list length | Length of cluster list sent by the route reflector. |
| Delete | Route has been deleted. |
| Ex | Exterior route. |
| Ext | BGP route received from an external BGP neighbor. |
| FlashAll | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| Hidden | Route not used because of routing policy. |
| IfCheck | Route needs forwarding RPF check. |
| IGP metric | Path through next hop with lower IGP metric is available. |

**Table 17: State Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| 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). |
| NotInstall | Route not to be installed in the forwarding table. |

**Table 17: State Output Field Values** *(Continued)*

| Value | Description |
|-------|-------------|
| 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. |
| ProtectionCand | Indicates paths requesting protection. |
| ProtectionPath | Indicates the route entry that can be used as a protection path. |
| 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. |

**Table 17: State Output Field Values** *(Continued)*

| Value | Description |
|-------|-------------|
| Unusable path | Path is not usable because of one of the following conditions:<br><br>• The route is damped.<br><br>• The route is rejected by an import policy.<br><br>• The route is unresolved. |
| Update source | Last tiebreaker is the lowest IP address value. |

**Table 18: 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. |

**Table 18: Communities Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| 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 be0); 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. |
| ipv6-origin | Defines the source of the IPv6 unicast address in a policy match condition. |
| ipv6-target | Defines the VPN IPv6 target unicast address used in a policy match condition. The target has the *128-bit IP address:16-bit number* format. For example, 2001:db8:9999::9. |
| ipv6-extended | Defines the extended format of the IPv6 unicast address in a policy match. For example, `ipv6-extended:203:<2001:db8::7>:67` |
| 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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                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:    65069
                    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:    65069
                    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:    65069
                    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:    65069
                    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:     65069
                 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:     65069 Peer AS:     65069
                 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:    65099
                    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)
```

```
10.1.1.3/32 (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)
10.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.8 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 programmed detail (with statistics information)

```
user@host> show route programmed detail

inet.0: 104 destinations, 104 routes (103 active, 0 holddown, 1 hidden)
```

```
inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
192.0.2.0/24 (1 entry, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 0
                Address: 0xa667694
                Next-hop reference count: 15
                Next hop: 198.51.100.0 via ae0.0
                Label operation: Push 50008, Push 3400(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3400: None;
                Label element ptr: 0x84e7e90
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 10, Stats IDs = { 536870923 }
                Statistics: Packets 1380, Bytes 681720
                Next hop: 198.51.100.1 via ae0.1
                Label operation: Push 50008, Push 3401(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3401: None;
                Label element ptr: 0x84e7ee0
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 11, Stats IDs = { 536870924 }
                Statistics: Packets 1444, Bytes 713336
                Next hop: 198.51.100.2 via ae0.2
                Label operation: Push 50008, Push 3410(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3410: None;
                Label element ptr: 0xa1926a0
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 12, Stats IDs = { 536870925 }
                Statistics: Packets 1420, Bytes 701480
                Next hop: 198.51.100.3 via ge-0/0/2.0
```

```
Label operation: Push 50008, Push 3411(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3411: None;
Label element ptr: 0xa1926c8
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 13, Stats IDs = { 536870926 }
Statistics: Packets 1486, Bytes 734084
Next hop: 198.51.100.4 via ge-0/0/2.1, selected
Label operation: Push 50008, Push 3420(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3420: None;
Label element ptr: 0xa1926f0
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 14, Stats IDs = { 536870927 }
Statistics: Packets 1399, Bytes 691106
Next hop: 198.51.100.5 via ge-0/0/2.2
Label operation: Push 50008, Push 3421(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3421: None;
Label element ptr: 0xa192718
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 15, Stats IDs = { 536870928 }
Statistics: Packets 1455, Bytes 718770
Next hop: 198.51.100.6 via ge-0/0/4.0
Label operation: Push 50008, Push 3450(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3450: None;
Label element ptr: 0xa192740
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
```

```
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 16, Stats IDs = { 536870929 }
                    Statistics: Packets 1407, Bytes 695058
                    Next hop: 198.51.100.7 via ge-0/0/4.1
                    Label operation: Push 50008, Push 3460(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3460: None;
                    Label element ptr: 0xa192768
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 17, Stats IDs = { 536870930 }
                    Statistics: Packets 1418, Bytes 700492
                    Next hop: 198.51.100.8 via ge-0/0/4.2
                    Label operation: Push 50008, Push 3470(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3470: None;
                    Label element ptr: 0xa192790
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 18, Stats IDs = { 536870931 }
                    Statistics: Packets 0, Bytes 0
                    State: <Active Int NSR-incapable Programmed>
                    Age: 12:46
                    Validation State: unverified
                    Announcement bits (3): 0-Resolve tree 6 1-Resolve tree 1 3-Resolve_IGP_FRR task
                    AS path: I
                    Session-IDs associated:
                    Session-id: 370 Version: 0
                    Thread: junos-main
```

**show route detail (with IPv6 unicast address community)**

```
user@host> show route 203.0.112.1 detail
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
```

```
203.0.112.1/32 (1 entry, 1 announced)
       *BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 559
                Address: 0x7b5de84
                Next-hop reference count: 12, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 13.1.1.1
                Next hop: 13.1.1.1 via ge-0/0/0.0, selected
                Session Id: 320
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 8:41      Metric: 1
                Validation State: unverified
                Task: BGP_100.13.1.1.1
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Communities: target:100:65 target:172.16.7.7:2 ipv6-target:<2001:db8::5>:65
                ipv6-target:<2001:db8::5>:101 ipv6-origin:<2001:db8::6>:66
                ipv6-origin:<2001:db8::6>:102 ipv6-extended:100:<2001:db8::7>:67
                ipv6-extended:256:<2001:db8::7>:103
                Accepted
                Localpref: 100
                Router ID: 2.2.2.2
                Thread: junos-main
```

**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: 65002 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:     65001 Peer AS:     65002
                   Age: 5:04:43
                      Validation State: unverified
                   Task: BGP_2.10.1.1.6+58198
                   AS path: 65002 I
                   Accepted MultipathContrib
                   Localpref: 100
                   Router ID: 172.16.1.3
```

### show route detail (with BGP, DeletePending)

```
user@host> show route detail
2001:db8:2:1:10.1.1.12/30 (1 entry, 0 announced)
       *BGP    Preference: 170/-101
               Route Distinguisher: 65002:1
               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:     65002 Peer AS:     65002
                    Age: 2w1d 17:42:45     Metric2: 1
                    Validation State: unverified
                    Task: BGP_10.2.1.1.4+55190
                    AS path: I
                    Communities: target:65002:1
                    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, 10.1.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 10.1.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                   RPF Nexthops :
                        ge-1/2/20.0, 10.1.96.1, Label: 301584, weight: 0xfffe
                        ge-1/2/19.0, 10.1.6.1, 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
                    Source MAC Address: 00:00:5e:00:52:01
                    Destination Address: 2001:db8:10:1:1:1:0:1
                    Destination UDP Port: 4790
                    VXLAN Flags: 0x08
...
```

## Release Information

Command introduced before Junos OS Release 7.4.

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.

`ipv6-origin`, `ipv6-target`, and `ipv6-extended Community` output fields supported in Junos OS Release and Junos OS Evolved Release 23.1.

# show route exact

## Syntax

```
show route exact destination-prefix
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route exact destination-prefix
<brief | detail | extensive | terse>
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*show route*

*show route detail*

*show route extensive*

*show route terse*

# show route export

**IN THIS SECTION**

- Syntax | **532**
- Syntax (EX Series Switches) | **532**
- Description | **532**
- Options | **532**
- Required Privilege Level | **533**
- Output Fields | **533**

# 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>
```

# 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. |

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

## Output Fields

lists the output fields for the `show route export` command. Output fields are listed in the approximate order in which they appear.

**Table 19: 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` |

**Table 19: show route export Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Flags | (`instance` keyword only) Flags for this feature on this instance:<br><br>• `config auto-policy`—The policy was deduced from the configured IGP export policies.<br><br>• `cleanup`—Configuration information for this instance is no longer valid.<br><br>• `config`—The instance was explicitly configured. | detail |
| Options | (`instance` keyword only) Configured option displays the type of routing tables the feature handles:<br><br>• `unicast`—Indicates *instance*.inet.0.<br><br>• `multicast`—Indicates *instance*.inet.2.<br><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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route export vrf-target

**IN THIS SECTION**

## Syntax

```
show route export vrf-target
<brief | detail>
<community community--regular-expression>
<logical-system (all | logical-system-name)>
```

## 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

## Output Fields

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 20: 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 |
| **Family** | Routing table entries for the specified family. | **brief** none |
| *type-of-routing-table(s)* | Type of routing tables the feature handles:<br><br>• **unicast**—Indicates *instance*.inet.0.<br><br>• **multicast**—Indicates *instance*.inet.2.<br><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** |

**Table 20: show route export vrf-target Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| **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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route extensive

## Syntax

```
show route extensive
<destination-prefix>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route extensive
<destination-prefix>
```

# 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

# Output Fields

Table 21 on page 540 describes the output fields for the `show route extensive` command. Output fields are listed in the approximate order in which they appear.

**Table 21: 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. |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *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><br>• `holddown` (routes that are in the pending state before being declared inactive).<br><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><br>• *MPLS-label* (for example, 80001 ).<br><br>• *interface-name* (for example, ge-1/0/2).<br><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><br>   • *neighbor-address*—Address of the neighbor.<br><br>   • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.<br><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><br>   • *vc-id*—Virtual circuit identifier.<br><br>   • *source*—Source of the advertisement: `Local` or `Remote`. |
| TSI | Protocol header information. |

**Table 21: 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><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><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><br>• `+`—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br><br>• `-` —A hyphen indicates the last active route.<br><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><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). |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| Next-hop type | Type of next hop. |
| 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><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. |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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). |
| 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><br>• 0x4000 indicates passive next hops. |
| State | State of the route (a route can be in more than one state). |
| 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. |

**Table 21: 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><br> • `Active preferred`—Currently active route was selected over this route. <br><br> • `Always compare MED`—Path with a lower multiple exit discriminator (MED) is available. <br><br> • `AS path`—Shorter AS path is available. <br><br> • `Cisco Non-deterministic MED selection`—Cisco nondeterministic MED is enabled and a path with a lower MED is available. <br><br> • `Cluster list length`—Path with a shorter cluster list length is available. <br><br> • `Forwarding use only`—Path is only available for forwarding purposes. <br><br> • `IGP metric`—Path through the next hop with a lower IGP metric is available. <br><br> • `IGP metric type`—Path with a lower OSPF link-state advertisement type is available. <br><br> • `Interior > Exterior > Exterior via Interior`—Direct, static, IGP, or EBGP path is available. <br><br> • `Local preference`—Path with a higher local preference value is available. <br><br> • `Next hop address`—Path with a lower metric next hop is available. <br><br> • `No difference`—Paths from all neighbors are identical for all necessary route attributes in some special scenarios, such as BGP L2VPN/VPLS, and there is no difference. <br><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><br> • `Number of gateways`—Path with a higher number of next hops is available. <br><br> • `Origin`—Path with a lower origin code is available. <br><br> • `OSPF version`—Path does not support the indicated OSPF version. <br><br> • `RIB preference`—Route from a higher-numbered routing table is available. <br><br> • `Route destinguisher`—64-bit prefix added to IP subnets to make them unique. |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| | • `Route metric or MED comparison`—Route with a lower metric or MED is available. <br><br> • `Route preference`—Route with a lower preference value is available. <br><br> • `Router ID`—Path through a neighbor with a lower ID is available. <br><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><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. |
| 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. |
| Task | Name of the protocol that has added the route. |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| 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><br>• `I`—IGP.<br><br>• `E`—EGP.<br><br>• `Recorded`—The AS path is recorded by the sample process (sampled).<br><br>• `?`—Incomplete; typically, the AS path was aggregated.<br><br>When AS path numbers are included in the route, the format is as follows:<br><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><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><br>• `( )`—Parentheses enclose a confederation.<br><br>• `( [ ] )`—Parentheses and brackets enclose a confederation set.<br><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 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
| --- | --- |
| validation-state | (BGP-learned routes) Validation status of the route:<br><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><br>• `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br><br>• `Unverified`—Indicates that origin validation is not enabled for the BGP peers.<br><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. |
| route status | Indicates the status of a BGP route:<br><br>• `Accepted`—The specified BGP route is imported by the default BGP policy.<br><br>• `Import`—The route is imported into a Layer 3 VPN routing instance.<br><br>• `Import-Protect`—A remote instance egress that is protected.<br><br>• `Multipath`—A BGP multipath active route.<br><br>• `MultipathContrib`—The route is not active but contributes to the BGP multipath.<br><br>• `Protect`—An egress route that is protected.<br><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. |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| Communities | Community path attribute for the route. |
| DeletePending | The DeletePending flag indicates that a BGP route needs to be processed due to a BGP peer down event. |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). |

**Table 21: show route extensive Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| Statistics ID Group | Indicates the Kernel ID number and Statistics ID number. |
| Statistics | Indicates the number of packets and data transferred. |

## 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 programmed extensive (with statistics information)**

```
user@host> show route programmed extensive

inet.0: 104 destinations, 104 routes (103 active, 0 holddown, 1 hidden)

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
192.0.2.0/24 (1 entry, 1 announced)
        *Static Preference: 5/100
                Next hop type: Router, Next hop index: 0
                Address: 0xa667694
                Next-hop reference count: 15
                Next hop: 198.51.100.0 via ae0.0
                Label operation: Push 50008, Push 3400(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Load balance label: Label 50008: None; Label 3400: None;
                Label element ptr: 0x84e7e90
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Statistics ID Group: Kernel ID = 10, Stats IDs = { 536870923 }
                Statistics: Packets 1380, Bytes 681720
```

```
                    Next hop: 198.51.100.1 via ae0.1
                    Label operation: Push 50008, Push 3401(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3401: None;
                    Label element ptr: 0x84e7ee0
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 11, Stats IDs = { 536870924 }
                    Statistics: Packets 1444, Bytes 713336
                    Next hop: 198.51.100.2 via ae0.2
                    Label operation: Push 50008, Push 3410(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3410: None;
                    Label element ptr: 0xa1926a0
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 12, Stats IDs = { 536870925 }
                    Statistics: Packets 1420, Bytes 701480
                    Next hop: 198.51.100.3 via ge-0/0/2.0
                    Label operation: Push 50008, Push 3411(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3411: None;
                    Label element ptr: 0xa1926c8
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 13, Stats IDs = { 536870926 }
                    Statistics: Packets 1486, Bytes 734084
                    Next hop: 198.51.100.4 via ge-0/0/2.1, selected
                    Label operation: Push 50008, Push 3420(top)
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3420: None;
                    Label element ptr: 0xa1926f0
                    Label parent element ptr: 0x0
                    Label element references: 3
```

```
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 14, Stats IDs = { 536870927 }
Statistics: Packets 1399, Bytes 691106
Next hop: 198.51.100.5 via ge-0/0/2.2
Label operation: Push 50008, Push 3421(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3421: None;
Label element ptr: 0xa192718
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 15, Stats IDs = { 536870928 }
Statistics: Packets 1455, Bytes 718770
Next hop: 198.51.100.6 via ge-0/0/4.0
Label operation: Push 50008, Push 3450(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3450: None;
Label element ptr: 0xa192740
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 16, Stats IDs = { 536870929 }
Statistics: Packets 1407, Bytes 695058
Next hop: 198.51.100.7 via ge-0/0/4.1
Label operation: Push 50008, Push 3460(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 50008: None; Label 3460: None;
Label element ptr: 0xa192768
Label parent element ptr: 0x0
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Statistics ID Group: Kernel ID = 17, Stats IDs = { 536870930 }
Statistics: Packets 1418, Bytes 700492
Next hop: 198.51.100.8 via ge-0/0/4.2
Label operation: Push 50008, Push 3470(top)
```

```
                    Label TTL action: prop-ttl, prop-ttl(top)
                    Load balance label: Label 50008: None; Label 3470: None;
                    Label element ptr: 0xa192790
                    Label parent element ptr: 0x0
                    Label element references: 3
                    Label element child references: 0
                    Label element lsp id: 0
                    Session Id: 0x0
                    Statistics ID Group: Kernel ID = 18, Stats IDs = { 536870931 }
                    Statistics: Packets 0, Bytes 0
                    State: <Active Int NSR-incapable Programmed>
                    Age: 12:46
                    Validation State: unverified
                    Announcement bits (3): 0-Resolve tree 6 1-Resolve tree 1 3-Resolve_IGP_FRR task
                    AS path: I
                    Session-IDs associated:
                    Session-id: 370 Version: 0
                    Thread: junos-main
```

## show route extensive (BGP-SRTE routes)

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
9.9.9.9-1 <c>/64 (1 entry, 0 announced):
        **SPRING-TE Preference: 8
                Next hop type: Indirect, Next hop index: 0
                Address: 0xdc33080
                Next-hop reference count: 1
                Next hop type: Router, Next hop index: 0
                Next hop: 1.2.2.2 via ge-0/0/2.0, selected
                Label element ptr: 0xdf671d0
                Label parent element ptr: 0x0
                Label element references: 11
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x0
                Protocol next hop: 299920
                Label operation: Push 800040
                Label TTL action: prop-ttl
                Load balance label: Label 800040: None;
                Composite next hop: 0xcd4f950 - INH Session ID: 0x0
```

```
                       Indirect next hop: 0xdc99a84 - INH Session ID: 0x0 Weight 0x1
                       State: <Active Int>
                       Local AS:   100
                       Age: 5d 17:37:19     Metric: 1     Metric2: 16777215
                       Validation State: unverified
                       Task: SPRING-TE
                       AS path:
                       SRTE Policy State:
                       SR Preference/Override: 200/100
                       Tunnel Source: Static configuration
                       Composite next hops: 1
                               Protocol next hop: 299920 Metric: 0
                               Label operation: Push 800040
                               Label TTL action: prop-ttl
                               Load balance label: Label 800040: None;
                               Composite next hop: 0xcd4f950 - INH Session ID: 0x0
                               Indirect next hop: 0xdc99a84 - INH Session ID: 0x0 Weight 0x1
                               Indirect path forwarding next hops: 1
                                       Next hop type: Router
                                       Next hop: 1.2.2.2 via ge-0/0/2.0
                                       Session Id: 0x0
                                       299920 /52 Originating RIB: mpls.0
                                         Metric: 0 Node path count: 1
                                         Forwarding nexthops: 1
                                               Next hop type: Router
                                               Next hop: 1.2.2.2 via ge-0/0/2.0
                                               Session Id: 0x141
```

## show route extensive (AIGP)

```
user@host> show route 3.3.3.3/32 extensive
inet.0: 18 destinations, 23 routes (18 active, 0 holddown, 1 hidden)
3.3.3.3/32 (2 entries, 1 announced)
TSI:
KRT in-kernel 3.3.3.3/32 -> {Push 299872}
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 570
                Address: 0x9760194
                Next-hop reference count: 7, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 12.12.12.2
                Next hop: 12.12.12.2 via ge-0/0/0.0, selected
```

```
            Load balance label: Label 299872: None;
            Session Id: 333
            State: <Active Ext>
            Peer AS:   200
            Age: 3d 4:16:17
            Validation State: unverified
            AIGP: 67
            Effective AIGP: 67 (AIGP plus IGP metric)
            Task: BGP_200_100.12.12.12.2
            Announcement bits (3): 0-KRT 3-Resolve tree 2 4-Resolve_IGP_FRR task
            AS path: 200 300 I
            Accepted
            Localpref: 100
            Router ID: 128.51.141.40
            Session-IDs associated:
            Session-id: 332 Version: 3
            Thread: junos-main
```

## Release Information

Command introduced before Junos OS Release 7.4.

`DeletePending` flag added to the command output in Junos OS Release 19.4R1.

# show route export vrf-target

**IN THIS SECTION**

- Release Information | **568**

## Syntax

```
show route export vrf-target
<brief | detail>
<community community--regular-expression>
<logical-system (all | logical-system-name)>
```

## 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

# Output Fields

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 22: 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 |
| **Family** | Routing table entries for the specified family. | **brief** none |
| *type-of-routing-table(s)* | Type of routing tables the feature handles: <br><br> • **unicast**—Indicates *instance*.inet.0. <br><br> • **multicast**—Indicates *instance*.inet.2. <br><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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route forwarding-table

## 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>
```

## 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. |
| 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

## Output Fields

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 23: 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 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Enabled protocols | The features and protocols that have been enabled for a given routing table. This field can contain the following values:<br><br>• BUM hashing—BUM hashing is enabled.<br><br>• MAC Stats—Mac Statistics is enabled.<br><br>• Bridging—Routing instance is a normal layer 2 bridge.<br><br>• No VLAN—No VLANs are associated with the bridge domain.<br><br>• All VLANs—The `vlan-id all` statement has been enabled for this bridge domain.<br><br>• Single VLAN—Single VLAN ID is associated with the bridge domain.<br><br>• MAC action drop—New MACs will be dropped when the MAC address limit is reached.<br><br>• Dual VLAN—Dual VLAN tags are associated with the bridge domain<br><br>• No local switching—No local switching is enabled for this routing instance..<br><br>• Learning disabled—Layer 2 learning is disabled for this routing instance.<br><br>• MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached.<br><br>• VPLS—The VPLS protocol is enabled.<br><br>• No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance.<br><br>• ACKed by all peers—All peers have acknowledged this routing instance.<br><br>• BUM Pruning—BUM pruning is enabled on the VPLS instance.<br><br>• Def BD VXLAN—VXLAN is enabled for the default bridge domain.<br><br>• EVPN—EVPN protocol is enabled for this routing instance. | All levels |

**Table 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| | • Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain.<br><br>• Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain.<br><br>• L2 backhaul—Layer 2 backhaul is enabled.<br><br>• FRR optimize—Fast reroute optimization<br><br>• MAC pinning—MAC pinning is enabled for this bridge domain.<br><br>• MAC Aging Timer—The MAC table aging time is set per routing instance.<br><br>• EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation.<br><br>• PBBN—This routing instance is configured as a provider backbone bridged network.<br><br>• PBN—This routing instance is configured as a provider bridge network.<br><br>• ETREE—The ETREE protocol is enabled on this EVPN routing instance.<br><br>• ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance.<br><br>• Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain.<br><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` |

**Table 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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><br> • `cloned` (`clon`)—(TCP or multicast only) Cloned route. <br><br> • `destination` (`dest`)—Remote addresses directly reachable through an interface. <br><br> • `destination down` (`iddn`)—Destination route for which the interface is unreachable. <br><br> • `interface cloned` (`ifcl`)—Cloned route for which the interface is unreachable. <br><br> • `route down` (`ifdn`)—Interface route for which the interface is unreachable. <br><br> • `ignore` (`ignr`)—Ignore this route. <br><br> • `interface` (`intf`)—Installed as a result of configuring an interface. <br><br> • `permanent` (`perm`)—Routes installed by the kernel when the routing table is initialized. <br><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 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| Flags | Route type flags:<br><br>• `none`—No flags are enabled.<br><br>• `accounting`—Route has accounting enabled.<br><br>• `cached`—Cache route.<br><br>• `incoming-iface` *interface-number*—Check against incoming interface.<br><br>• `prefix load balance`—Load balancing is enabled for this prefix.<br><br>• `rt nh decoupled`—Route has been decoupled from the next hop to the destination.<br><br>• `sent to PFE`—Route has been sent to the Packet Forwarding Engine.<br><br>• `static`—Static route. | `extensive` |
| Next hop | IP address of the next hop to the destination.<br><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` |

**Table 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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><br>• `broadcast (bcst)`—Broadcast.<br><br>• `deny`—Deny.<br><br>• `discard (dscd)` —Discard.<br><br>• `hold`—Next hop is waiting to be resolved into a unicast or multicast type.<br><br>• `indexed (idxd)`—Indexed next hop.<br><br>• `indirect (indr)`—Indirect next hop.<br><br>• `local (locl)`—Local address on an interface.<br><br>• `routed multicast (mcrt)`—Regular multicast next hop.<br><br>• `multicast (mcst)`—Wire multicast next hop (limited to the LAN).<br><br>• `multicast discard (mdsc)`—Multicast discard.<br><br>• `multicast group (mgrp)`—Multicast group member.<br><br>• `receive (recv)`—Receive.<br><br>• `reject (rjct)`—Discard. An ICMP unreachable message was sent.<br><br>• `resolve (rslv)`—Resolving the next hop.<br><br>• `unicast (ucst)`—Unicast.<br><br>• `unilist (ulst)`—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.<br><br>• `VxLAN Local`—EVPN Type 5 route in kernel. | `detail extensive` |
| Index | Software index of the next hop that is used to route the traffic for a given prefix. | `detail extensive` `none` |

**Table 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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` |
| Alternate forward nh index | Index number of the alternate next hop interface. Seen with **multicast** option only. | `extensive` |
| Next-hop L3 Interface | The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the **multicast** option. | `extensive` |
| Next-hop L2 Interface | The next hop layer 2 interface. Seen with the **multicast** option only. | `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` |

**Table 23: show route forwarding-table Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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` |
| sftw | On Junos OS Evolved devices, `sftw` is another type of next-hop. A software next-hop can contain a link to another single next-hop, and hence it often works as an indirect next-hop. The usage of `sftw` helps to reduce churn in Junos OS Evolved objects by providing a common object where topology updates can be reflected. Further MPLS states can also be attached to these software next-hops. When a route is created, a `sftw` next-hop seems to be created whether it is the management port or the loopback interface. | `detail extensive` none |

## 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 summary

```
user@switch> show route forwarding-table summary


Routing table: default.inet
Internet:
        user:         6 routes
        perm:         5 routes
        intf:         8 routes
        dest:        12 routes
        ifdn:         1 routes
        iddn:         3 routes
```

## show route forwarding-table extensive

```
user@switch> show route forwarding-table extensive


Routing table: default.inet [Index 0]
Internet:


Destination:  default
  Route type: user
  Route reference: 2                 Route interface-index: 0
  Flags: sent to PFE, rt nh decoupled
  Nexthop: 0:12:f2:21:cf:0
  Next-hop type: unicast             Index: 333     Reference: 5
  Next-hop interface: me0.0


Destination:  default
  Route type: permanent
  Route reference: 0                 Route interface-index: 0
```

```
    Flags: none
    Next-hop type: reject            Index: 36        Reference: 2


Destination:  0.0.0.0/32
  Route type: permanent
  Route reference: 0                 Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: discard             Index: 34        Reference: 1


Destination:  2.2.2.0/24
  Route type: interface
  Route reference: 0                 Route interface-index: 66
  Flags: sent to PFE
  Next-hop type: resolve             Index: 1309      Reference: 1
  Next-hop interface: ae0.0


Destination:  2.2.2.0/32
  Route type: destination
  Route reference: 0                 Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 2.2.2.0
  Next-hop type: receive             Index: 1307      Reference: 1
  Next-hop interface: ae0.0


Destination:  2.2.2.1/32
  Route type: destination
  Route reference: 0                 Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 0:21:59:cc:89:c0
  Next-hop type: unicast             Index: 1320      Reference: 1
  Next-hop interface: ae0.0


Destination:  2.2.2.2/32
  Route type: interface
  Route reference: 0                 Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 2.2.2.2
  Next-hop type: local               Index: 1308      Reference: 2


Destination:  2.2.2.2/32
  Route type: destination
  Route reference: 0                 Route interface-index: 66
  Flags: none
```

```
  Nexthop: 2.2.2.2
  Next-hop type: local             Index: 1308     Reference: 2


Destination:  2.2.2.255/32
  Route type: destination
  Route reference: 0               Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 2.2.2.255
  Next-hop type: broadcast         Index: 1306     Reference: 1
  Next-hop interface: ae0.0
```

## show route forwarding-table ccc

```
user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
ge-0/0/0.10  (CCC) user     0 3.3.3.2            Push 300112  1343     2 ae1.0
```

## show route forwarding-table family (MPLS)

```
user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
0                  user     0                    recv    49     3
1                  user     0                    recv    49     3
2                  user     0                    recv    49     3
299776             user     0                    Pop   1334     2 ge-0/0/0.10
299792             user     0                    Pop   1339     2 ge-0/0/0.14
299808             user     0                    Pop   1341     2 ge-0/0/0.2
299824             user     0                    Pop   1344     2 ge-0/0/0.11
299840             user     0                    Pop   1345     2 ge-0/0/0.13
299856             user     0                    Pop   1346     2 ge-0/0/0.18
299872             user     0                    Pop   1347     2 ge-0/0/0.16
299888             user     0                    Pop   1348     2 ge-0/0/0.7
299904             user     0                    Pop   1349     2 ge-0/0/0.20
299920             user     0                    Pop   1350     2 ge-0/0/0.19
```

```
299936             user     0                    Pop    1351     2 ge-0/0/0.17
299952             user     0                    Pop    1352     2 ge-0/0/0.9
299968             user     0                    Pop    1353     2 ge-0/0/0.1
299984             user     0                    Pop    1354     2 ge-0/0/0.12
300000             user     0                    Pop    1355     2 ge-0/0/0.8
300016             user     0                    Pop    1356     2 ge-0/0/0.4
300032             user     0                    Pop    1357     2 ge-0/0/0.5
300048             user     0                    Pop    1358     2 ge-0/0/0.3
300064             user     0                    Pop    1359     2 ge-0/0/0.15
ge-0/0/0.1  (CCC) user     0 3.3.3.2             Push 300064  1340     2 ae1.0
ge-0/0/0.2  (CCC) user     0 3.3.3.2             Push 299872  1328     2 ae1.0
ge-0/0/0.3  (CCC) user     0 3.3.3.2             Push 299792  1323     2 ae1.0
ge-0/0/0.4  (CCC) user     0 3.3.3.2             Push 300016  1337     2 ae1.0
ge-0/0/0.5  (CCC) user     0 3.3.3.2             Push 299824  1325     2 ae1.0
ge-0/0/0.7  (CCC) user     0 3.3.3.2             Push 299920  1331     2 ae1.0
ge-0/0/0.8  (CCC) user     0 3.3.3.2             Push 299840  1326     2 ae1.0
ge-0/0/0.9  (CCC) user     0 3.3.3.2             Push 299888  1329     2 ae1.0
ge-0/0/0.10 (CCC) user     0 3.3.3.2             Push 300112  1343     2 ae1.0
ge-0/0/0.11 (CCC) user     0 3.3.3.2             Push 299776  1322     2 ae1.0
ge-0/0/0.12 (CCC) user     0 3.3.3.2             Push 299952  1333     2 ae1.0
ge-0/0/0.13 (CCC) user     0 3.3.3.2             Push 300096  1342     2 ae1.0
ge-0/0/0.14 (CCC) user     0 3.3.3.2             Push 299984  1335     2 ae1.0
ge-0/0/0.15 (CCC) user     0 3.3.3.2             Push 299936  1332     2 ae1.0
ge-0/0/0.16 (CCC) user     0 3.3.3.2             Push 299808  1324     2 ae1.0
ge-0/0/0.17 (CCC) user     0 3.3.3.2             Push 300000  1336     2 ae1.0
ge-0/0/0.18 (CCC) user     0 3.3.3.2             Push 300032  1338     2 ae1.0
ge-0/0/0.19 (CCC) user     0 3.3.3.2             Push 299904  1330     2 ae1.0
ge-0/0/0.20 (CCC) user     0 3.3.3.2             Push 299856  1327     2 ae1.0
```

**show route forwarding-table family (IPv6)**

```
user@switch> show route forwarding-table family inet6
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-switch.inet6
```

```
Internet6:
Destination       Type RtRef Next hop       Type Index NhRef Netif
default           perm   0                  rjct   530    1
::/128            perm   0                  dscd   528    1
2:1::3a00/312     user   0                  indr 131070    2
                                            comp   572    1
2:1::3a82/320     user   0                  indr 131071    3
                                            comp   573    1
2:1::3af0/320     user   0                  indr 131071    3
                                            comp   573    1
2:1:0:ff00::/56   user   0                  mdsc   529    2
ff00::/8          perm   0                  mdsc   529    2
ff02::1/128       perm   0 ff02::1          mcst   526    1


Routing table: __master.anon__.inet6
Internet6:
Destination       Type RtRef Next hop       Type Index NhRef Netif
default           perm   0                  rjct   554    1
::/128            perm   0                  dscd   552    1
ff00::/8          perm   0                  mdsc   553    1
ff02::1/128       perm   0 ff02::1          mcst   550    1
```

## show route forwarding-table label

```
user@switch> show route forwarding-table label 29976

Routing table: default.mpls
MPLS:
Destination       Type RtRef Next hop       Type Index NhRef Netif
299776            user   0                  Pop   1334    2 ge-0/0/0.10
```

## show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3

Routing table: default.inet
Internet:
```

**show route forwarding-table multicast**

```
user@switch> show route forwarding-table multicast

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop        Type Index NhRef Netif
224.0.0.0/4        perm    1                   mdsc   35    1
224.0.0.1/32       perm    0 224.0.0.1         mcst   31    3
224.0.0.5/32       user    1 224.0.0.5         mcst   31    3

Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop        Type Index NhRef Netif
224.0.0.0/4        perm    0                   mdsc 1289    1
224.0.0.1/32       perm    0 224.0.0.1         mcst 1285    1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop        Type Index NhRef Netif
ff00::/8           perm    0                   mdsc   43    1
ff02::1/128        perm    0 ff02::1           mcst   39    1
```

show route forwarding-table destination extensive (EVPN Type 5 route with Type 2 and Type 5 route coexistence)

```
user@device> show route forwarding-table destination 10.1.1.20 table vrf1 extensive
Routing table: vrf1.inet [Index 9]
Internet:
Destination:  10.1.1.20/32
  Route type: user
  Route reference: 0                     Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, VxLAN Local
  Nexthop:
  Next-hop type: composite          Index: 2694     Reference: 7
  Next-hop type: indirect           Index: 524326   Reference: 2
  Next-hop type: unilist            Index: 524288   Reference: 5
  Nexthop: 10.1.1.1
  Next-hop type: unicast            Index: 1724     Reference: 15
```

```
  Next-hop interface: xe-0/0/1.0    Weight: 0x0

  Nexthop: 10.1.1.4  Next-hop type: unicast             Index: 1725     Reference: 15

  Next-hop interface: xe-0/0/4.0    Weight: 0x0
```

**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;
        }
    }
}
```

show route forwarding-table (Junos OS Evolved devices with next-hop type 'sftw')

```
user@host> show route forwarding-table | grep sftw
Internet:
Destination        Type RtRef Next hop        Type Index NhRef Netif
default            user   0                    sftw   7021    1 re0:mgmt-0.0
10.21.1.0/32       user   0                    sftw   7022    1 et-0/0/32.0

user@host> show route forwarding-table | grep sftw
Internet:
Destination        Type RtRef Next hop        Type Index NhRef Netif
default            user   0                    sftw   5021    1 re0:mgmt-0.0
10.22.1.0/32       user   0                    sftw   5022    1 et-0/0/32.0
```

## 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.

**RELATED DOCUMENTATION**

*show route instance*

# show route forwarding-table interface-name

**IN THIS SECTION**

## Syntax

```
show route forwarding-table interface-name interface-name
<detail | extensive>
<all>
```

## 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

## Output Fields

Table 24 on page 590 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 24: 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** |
| Network | Network information:<br><br>• **<Link>**—Physical interface, not a logical interface.<br><br>• **<PtoP>**—Point-to-point network.<br><br>• **ipaddress**—Network address. | All levels |

**Table 24: show route forwarding-table interface-name Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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


 ...
```

## Release Information

Command introduced in Junos OS Release 9.6.

# show route hidden

**IN THIS SECTION**

# Syntax

```
show route hidden
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

# 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

## 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 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)
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route inactive-path

## Syntax

```
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route inactive-path
<brief | detail | extensive | terse>
```

## 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

## 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.

**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)
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*show route*

*show route active-path*

*show route detail*

*show route extensive*

*show route terse*

# show route inactive-prefix

## 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>
```

## 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

## 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 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route instance

## Syntax

```
show route instance
<brief | detail | summary>
<instance-name>
<logical-system (all | logical-system-name)>
<operational>
```

## Syntax (EX Series Switches and QFX Series)

```
show route instance
<brief | detail | summary>
<instance-name>
<operational>
```

## 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

## Output Fields

Table 25 on page 609 lists the output fields for the `show route instance` command. Output fields are listed in the approximate order in which they appear.

**Table 25: 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` |
| 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` |

**Table 25: show route instance Output Fields** *(Continued)*

| Field Name | Field Description | Level of Output |
|---|---|---|
| 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><br>• Pending:*protocol-name*—List of protocols that have not yet completed graceful restart for this routing table.<br><br>• Complete—All protocols have restarted for this routing table. | `detail` |
| Primary rib | Primary table for this routing instance. | `brief none`<br>`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 (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
```

## Release Information

Command introduced before Junos OS Release 7.4.

**RELATED DOCUMENTATION**

Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling

*Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

# show route label-switched-path

**IN THIS SECTION**

## 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>
```

## 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

## 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

10.1.1.1/32        [MPLS/7] 00:00:06, metric 0
                > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

```
10.3.3.3/32          *[MPLS/7] 00:00:06, metric 0
                     > to 10.111.1.9 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


10.2.2.2/32          *[MPLS/7] 00:00:06, metric 0
                     > to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
10.4.4.4/32          *[MPLS/7] 00:00:06, metric 0
                        to 10.111.1.9 via s0-0/0/0, label-switched-path abc
                     > to 10.111.1.9 via s0-0/0/0, label-switched-path xyz
                        to 10.111.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32       [MPLS/7] 00:00:06, metric 0
                     > to 10.111.1.9 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route localization

**IN THIS SECTION**

## Syntax

```
show route localization
```

## Description

(T320, T640, and T1600 routers only) Display route localization details.

## Options

detail                    Display detailed output.

## Required Privilege Level

view

## Output Fields

Table 26 on page 617 lists the output fields for the `show route localization` command. Output fields are listed in the approximate order in which they appear.

**Table 26: 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

**command-name**

```
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
```

**command-name**

```
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
```

## Release Information

Command introduced in Junos OS Release 11.4.

### RELATED DOCUMENTATION

*Example: Configuring Packet Forwarding Engine FIB Localization*

# show route martians

**IN THIS SECTION**

## 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>
```

## 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

## Output Fields

Table 27 on page 620 lists the output fields for the `show route martians` command. Output fields are listed in the approximate order in which they appear

**Table 27: 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
...
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route next-hop

## Syntax

```
show route next-hop next-hop
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route next-hop next-hop
<brief | detail | extensive | terse>
```

## 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

## 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 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)
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route protocol

**IN THIS SECTION**

## 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>
```

## Syntax (PTX and ACX Series Switches)

```
show route protocol protocol
(index index | prefix prefix)
<brief | detail | extensive | terse>
```

## Operational mode

`cli` mode

`cli-pfe` mode.

## Description

Display the route entries in the routing table that were learned from a particular protocol. Displays details about routes installed in `cli-pfe` mode after FIB compression of FIB compression statistics (PTX and ACX series devices only).

## 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

## 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 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 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 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 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 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 rsvp extensive

```
user@host> show route protocol rsvp extensive
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)


inet.3: 3 destinations, 4 routes (3 active, 0 holddown, 0 hidden)
192.168.4.4/32 (2 entries, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 0
                Address: 0xaad4fbc
                Next-hop reference count: 2, key opaque handle: 0x0, non-key opaque handle: 0x0
                Next hop: 192.168.0.1 via ge-0/0/0.0 weight 0x1, selected
                Label-switched-path PE1_to_PE2_1
                Label operation: Push 299920
                Label TTL action: prop-ttl
                Load balance label: Label 299920: None;
                Label element ptr: 0x9938be8
```

```
                Label parent element ptr: 0x0
                Label element references: 3
                Label element child references: 1
                Label element lsp id: 2
                Session Id: 0
                Next hop: 192.168.1.1 via ge-0/0/1.0 weight 0x8001 uflags Mark connection
protection tlv
                Label-switched-path Bypass->192.168.0.1
                Label operation: Push 299920
                Label TTL action: prop-ttl
                Load balance label: Label 299920: None;
                Label element ptr: 0x9938eb8
                Label parent element ptr: 0x9938be8
                Label element references: 2
                Label element child references: 0
                Label element lsp id: 2
                Session Id: 0
                State: <Active Int>
                Local AS:    100
                Age: 1:20        Metric: 100
                Validation State: unverified
                Actual IGP metric: 2
                Task: RSVP
                Announcement bits (2): 0-Resolve tree 3 3-Resolve tree 1
                AS path: I
                Thread: junos-main
```

**show route protocol index prefix detail (ACX7100-32C, ACX7100-48L, and ACX7509 Devices)**

Displays if a route is installed or not when FIB compression is enabledIf the route is not installed, the command displays the details of the parent route installed, that compressed the route.

```
user@root:pfe> show route proto 2 index 0 prefix 200.1.0.16 detail
Protocol: 2
Protocol     : IPv4
Table        : default
Prefix       : 200.1.0.16 (primary)
NH           : 25035 (software)
Flags        : 0x00008000
Details      :
```

```
guid          : 889058232562
type          :  user
nhid          :  25035
Forwarding state:
installed?    :  no
(Installed parent: 200.1.0/23)
```

## Release Information

Command introduced before Junos OS Release 7.4.

`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.

### RELATED DOCUMENTATION

*show route*

*show route detail*

*show route extensive*

*show route terse*

# show route range

**IN THIS SECTION**

- Syntax | **637**
- Syntax (EX Series Switches) | **637**
- Description | **637**
- Options | **637**

## 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>
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route resolution

**IN THIS SECTION**

## 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>
```

## 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

## Output Fields

Table 28 on page 645 describes the output fields for the `show route resolution` command. Output fields are listed in the approximate order in which they appear.

**Table 28: 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. |

**Table 28: show route resolution Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| Originating RIB | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix. |
| Metric | Metric associated with the forwarding next hop. |
| Node path count | Number of nodes in the path. |
| Forwarding next hops | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

*Example: Configuring Route Resolution on PE Routers*

# 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>
```

## 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

## 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>
```

## Release Information

Command introduced in Junos OS Release 8.5.

# show route source-gateway

## 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>
```

## 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

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

# show route summary

## 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
```

## Operational mode

`cli` mode

`cli-pfe` mode.

# Description

Display summary statistics about the entries in the routing table. Additionally, displays summary of FIB compression statistics in `cli-pfe` mode. (PTX and ACX series devices).

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

## Output Fields

Table 29 on page 664 lists the output fields for the `show route summary` command. Output fields are listed in the approximate order in which they appear.

**Table 29: show route summary Output Fields**

| Field Name | Field Description |
|---|---|
| **Router ID** | Address of the local routing device. |
| **Highwater Mark** | High-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system. Shows the following details:<br><br>• RIB Unique destination routes<br><br>• RIB routes<br><br>• FIB routes<br><br>• VRF type routing instances<br><br>The high-water mark data can also be viewed in the syslog at the **LOG_NOTICE** level. |

**Table 29: show route summary Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *routing-table-name* | Name of the routing table example:<br><br>• **inet.0**— List all routes for IPv4.<br><br>• **junos-main::inet.0**—Lists all active routes for Ipv4.<br><br>• **junos-bgpshard0::inet.0**—Lists IPv4 routes present in bgpshard0 thread.<br><br>• **__raass__inet.inet.0**—Lists IPv4 routes stored in RaaS Server that are received from the clients.<br><br>• **__raasc__inet.inet.0**—Lists IPv4 routes stored in RaaS Clients that are used for local resolution.<br><br>• **inet.6**— List all routes for IPv6.<br><br>• **junos-main::inet6.0**— Lists all active routes for IPv6.<br><br>• **junos-bgpshard0::inet6.0**—Lists IPv6 routes present in bgpshard0 thread.<br><br>• **__raass__inet.inet6.0**—Lists IPv6 routes stored in RaaS Server that are received from the clients.<br><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. |
| **routes** | Number of routes in the routing table:<br><br>• **active**—Number of routes that are active.<br><br>• **holddown**—Number of routes that are in the hold-down state before being declared inactive.<br><br>• **hidden**—Number of routes that are not used because of routing policy. |

**Table 29: show route summary Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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><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>    `Restart Pending: OSPF LDP VPN`<br><br>This indicates that `OSPF`, `LDP`, and `VPN` protocols did not restart for `LDP.inet.0` routing table.<br><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 `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><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><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. |
| **Local** | Local routes. |

**Table 29: show route summary Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *protocol-name* | Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**. |

Output fields for CLI-PFE mode (ACX7100-32C, ACX7100-48L, and ACX7509 Devices) that display route table statistics at the PFE (for example, IPv4, IPv6, MPLS Route Tables)

| | |
|---|---|
| Index | Routing table index number. |
| Routes | Number of active routes |
| Size(b) | Size of the packets #chk w SME. |
| Prefixes | Actual number of routes published by the peer devices. |
| Aggr | Number of aggregated routes generated by the PFE after compression that includes valid next hops |
| Installed | Number of routes installed on the ACX device. |
| Comp(%) | Route table compression percentage. |

## Sample Output

### show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52

Highwater Mark (All time / Time averaged watermark)
    RIB Unique destination routes: 51 at 2020-06-15 13:54:01 / 51
    RIB routes                   : 52 at 2020-06-15 13:54:01 / 52
    FIB routes                   : 33 at 2020-06-15 13:54:01 / 33
    VRF type routing instances   : 0 at 2020-06-15 13:54:00
```

```
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 (cli-pfe mode in ACX7100-32C, ACX7100-48L, and ACX7509 Devices)**

The following command shows the route summary information in `cli-pfe` mode only:

```
user@host:pfe> show route summary

IPv4 Route Tables:
Index      Routes      Size(b) Prefixes     Aggr     Installed  Comp(%)
--------   ----------  ----------  ---------  ---------  ----------  ------
Default    10095472  1569927632   5000016    4999986       11525   100
1                 0           0         0          0           0     0
51               10        1040         5          0           5     0
36738            10        1040         5          0           5     0


IPv6 Route Tables:
Index      Routes      Size(b) Prefixes     Aggr     Installed  Comp(%)
--------   ----------  ----------  ---------  ---------  ----------  ------
0       IPv6      1300025              1299963          62       100

1       IPv6            0                    0           0         0

51      IPv6            6                    0           6         0

36738   IPv6            6                    0           6         0
MPLS Route Tables:
Index      Routes      Size(b) Prefixes     Aggr     Installed  Comp(%)
--------   ----------  ----------  ---------  ---------  ----------  ------
Default          14        1456         7          0           7     -
52                2         208         1          0           1     -


IPv6 Route Tables:
Index      Routes      Size(b) Prefixes     Aggr     Installed  Comp(%)
--------   ----------  ----------  ---------  ---------  ----------  ------
Default          27        2808        13          0          16  4294967273
1                 0           0         0          0           0     0
51               12        1248         6          0           6     0
53               12        1248         6          0           6     0
36738            12        1248         6          0           6     0


CLNP Route Tables:
Index      Routes      Size(b) Prefixes     Aggr     Installed  Comp(%)
--------   ----------  ----------  ---------  ---------  ----------  ------
```

| Default | 4 | 416 | 2 | 0 | 2 | - |
| 51 | 2 | 208 | 1 | 0 | 1 | - |

## 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

`rib-sharding` option introduced for cRPD Release 20.1R1.

# show route table

**IN THIS SECTION**

-

-

-

-

-

-

## 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>
```

## Description

Display the route entries in a particular routing table.

## Options

**brief | detail | extensive |**     (Optional) Display the specified level of output.
**terse**

| logical-system (all \| *logical-system-name*) | (Optional) Perform this operation on all logical systems or on a particular logical system. This option is only supported on Junos OS. |
|---|---|
| *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

## Output Fields

describes the output fields for the `show route table` command. Output fields are listed in the approximate order in which they appear.

**Table 30: show route table Output Fields**

| Field Name | Field Description |
|---|---|
| *routing-table-name* | Name of the routing table (for example, inet.0). |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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><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>    `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><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><br>• `holddown` (routes that are in the pending state before being declared inactive)<br><br>• `hidden` (routes that are not used because of a routing policy) |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| *route-destination*<br>(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><br>• *MPLS-label* (for example, 80001).<br><br>• *interface-name* (for example, ge-1/0/2).<br><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><br>    • *neighbor-address*—Address of the neighbor.<br><br>    • *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: `NoCtrlWord` or `CtrlWord`.<br><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><br>    • `vc-id`—Virtual circuit identifier.<br><br>    • *source*—Source of the advertisement: `Local` or `Remote`.<br><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><br>    • *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.<br><br>    • *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.<br><br>    • *IP address length*—(1 octet) Length of IP address in bits.<br><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. |

**Table 30: show route table 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><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><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><br>• `+`—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br><br>• `-`—A hyphen indicates the last active route.<br><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><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). |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| Next-hop type | Type of next hop. For a description of possible values for this field, see Table 31 on page 683. |
| 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><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). |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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 32 on page 685. |
| 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. |
| 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. |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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><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><br>• I—IGP.<br><br>• E—EGP.<br><br>• Recorded—The AS path is recorded by the sample process (sampled).<br><br>• ?—Incomplete; typically, the AS path was aggregated.<br><br>When AS path numbers are included in the route, the format is as follows:<br><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><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><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><br>• ( )—Parentheses enclose a confederation.<br><br>• ( [ ] )—Parentheses and brackets enclose a confederation set.<br><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 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| validation-state | (BGP-learned routes) Validation status of the route:<br><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><br>• `Unknown`—Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br><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><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. |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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 33 on page 689 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. |

**Table 30: show route table Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| 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. |
| 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. |

describes all possible values for the Next-hop Types output field.

**Table 31: Next-hop Types Output Field Values**

| Next-Hop Type | Description |
|---|---|
| Broadcast (bcast) | Broadcast next hop. |

**Table 31: Next-hop Types Output Field Values** *(Continued)*

| Next-Hop Type | Description |
|---|---|
| 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. |
| Multicast group (mgrp) | Multicast group member. |
| Receive (recv) | Receive. |

**Table 31: Next-hop Types Output Field Values** *(Continued)*

| Next-Hop Type | Description |
|---|---|
| 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><br> To qualify as a 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><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. |

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 32: State Output Field Values**

| Value | Description |
|---|---|
| Accounting | Route needs accounting. |
| Active | Route is active. |

**Table 32: State Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| Always Compare MED | Path with a lower multiple exit discriminator (MED) is available. |
| AS path | Shorter AS path is available. |
| Cisco Non-deterministic MED selection | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| Clone | Route is a clone. |
| Cluster list length | Length of cluster list sent by the route reflector. |
| Delete | Route has been deleted. |
| Ex | Exterior route. |
| Ext | BGP route received from an external BGP neighbor. |
| FlashAll | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| Hidden | Route not used because of routing policy. |
| IfCheck | Route needs forwarding RPF check. |
| IGP metric | Path through next hop with lower IGP metric is available. |
| Inactive reason | Flags for this route, which was not selected as best for a particular destination. |
| Initial | Route being added. |

**Table 32: State Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| 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). |
| 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. |

**Table 32: State Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| 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><br>• The route is damped.<br><br>• The route is rejected by an import policy.<br><br>• The route is unresolved. |
| Update source | Last tiebreaker is the lowest IP address value. |
| VxlanLocalRT | Route is an EVPN Type 5 route (IP prefix route). |

describes the possible values for the Communities output field.

**Table 33: 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. |
| GBP Tag | Assigned Group Based Policy (GBP) tag from 1 through 65535. |
| 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 be0); 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*. |

**Table 33: Communities Output Field Values** *(Continued)*

| Value | Description |
|---|---|
| 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. |
| evpn-mcast-flags | Identifies the value in the multicast flags extended community, which includes bits that indicate whether snooping or optimized intersubnet multicast (OISM) is enabled. For example:<br><br>• 0x01—IGMP snooping bit<br><br>• 0x02—MLD snooping bit<br><br>• 0x08—OISM bit |
| 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 evpn-vs2.evpn.0 extensive

```
user@host> show route table evpn-vs2.evpn.0 extensive
evpn-vs2.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
Restart Complete
2:10.3.3.3:2::601: :2c:6b:f5:f2:87:f0/304 MAC/IP (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.3.3.3:2
                Next hop type: Indirect, Next hop index: 0
                Address: 0x76b66dc
                Next-hop reference count: 32
                Source: 10.2.2.2
                Protocol next hop: 10.3.3.3
                Indirect next hop: 0x2 no-forward INH Session ID: 0
                State: <Secondary Active Ext>
                Local AS:   102 Peer AS:   201
                Age: 22:56:12      Metric2: 1
                Validation State: unverified
                Task: BGP_201.10.2.2.2
                Announcement bits (1): 0-evpn-vs2-evpn
                AS path: 201 203 I
                Communities: target:5:5 encapsulation:vxlan(0x8) evpn-default-gateway gbp-tag:
200
```

**show route table bgpmcast.0 extensive**

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



4:3:0:0:32:10.2.20.20:32:225.1.1.1:10.4.4.4:10.20.20.20/240 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Tunnel type: AnyEncap, RPF tunnel:, Remote end point: 10.1.4.4
                Tunnel type: AnyEncap, Remote end point: 10.2.4.4
                Tunnel type: AnyEncap, Remote end point: 10.3.4.4
                Tunnel type: AnyEncap, Remote end point: 10.4.6.4
                Next hop type: Indirect, Next hop index: 0
                Address: 0xc54639c
                Next-hop reference count: 1
                Source: 10.1.1.1
                Protocol next hop: 10.20.20.20
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                State: <Active Int Ext>
                Local AS:   65100 Peer AS:   65100
                Age: 30        Metric2: 0
                Validation State: unverified
                Task: BGP_100.1.1.1.1
                Announcement bits (1): 0-bgpmcast global task
                AS path: I
                Communities: target:10.4.4.4:0
                Accepted
                Localpref: 100
                Router ID: 10.1.1.1
                Indirect next hops: 1
                        Protocol next hop: 10.20.20.20
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.4.20.20 via ge-0/0/5.0
                                Session Id: 0x140
                                10.20.20.20/32 Originating RIB: inet.0
                                  Node path count: 1
                                  Forwarding nexthops: 1
                                        Next hop type: Router
```

```
                                     Next hop: 10.4.20.20 via ge-0/0/5.0
                                     Session Id: 0x140
```

## 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 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: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both


NODE { AS:100 ISO:0003.0003.0303.00 ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:16
                      Fictitious
NODE { AS:100 ISO:1282.0404.9202.00 ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
NODE { AS:100 ISO:1282.0404.9202.02 ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
LINK { Local { AS:100 ISO:0003.0003.0303.00 }.{ IPv4:23.0.0.2 IPv6:23::2 } Remote { AS:100
ISO:1282.0404.9202.02 }.{ } ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.00 }.{ IPv4:23.0.0.1 IPv6:23::1 } Remote { AS:100
ISO:1282.0404.9202.02 }.{ } ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.02 }.{ } Remote { AS:100 ISO:0003.0003.0303.00 }.{ }
ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
LINK { Local { AS:100 ISO:1282.0404.9202.02 }.{ } Remote { AS:100 ISO:1282.0404.9202.00 }.{ }
ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:07
                      Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0 }/1216
                  *[IS-IS/18] 10:19:16
                      Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:32.32.32.32/32 } ISIS-L2:0 }/
1216
                  *[IS-IS/18] 10:19:16
                      Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:33.33.33.33/32 } ISIS-L2:0 }/
1216
                  *[IS-IS/18] 10:19:16
                      Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:100.100.100.0/32 } ISIS-L2:0 }/
1216
                  *[IS-IS/18] 10:19:16
                      Fictitious
PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv4:128.204.49.193/32 } ISIS-L2:0 }/
```

```
1216
                      *[IS-IS/18] 10:19:16
                          Fictitious
PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0 }/1216
                      *[IS-IS/18] 10:19:07
                          Fictitious
PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv4:128.204.49.202/32 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:07
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:3::3/128 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:16
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:23::/64 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:16
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:100:100:100::100/128 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:16
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:0003.0003.0303.00 } { IPv6:abcd::128:204:49:193/128 } ISIS-
L2:0 }/1216
                      *[IS-IS/18] 10:19:16
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:2::2/128 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:07
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:23::/64 } ISIS-L2:0 }/
1216
                      *[IS-IS/18] 10:19:07
                          Fictitious
IPV6 PREFIX { Node { AS:100 ISO:1282.0404.9202.00 } { IPv6:abcd::128:204:49:202/128 } ISIS-
L2:0 }/1216
                      *[IS-IS/18] 10:19:07
                          Fictitious
```

## show route table lsdist.0 (BGP Confederation Enabled)

```
user@host>show route table lsdist.0 detail
lsdist.0: 72 destinations, 72 routes (72 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:17 Member-ASN:65002 ISO:0004.0404.0400.00 ISIS-L2:0 }/1216
                  *[BGP/170] 00:00:23, localpref 100
                    AS path: (65002) I, validation-state: unverified
                  >  to 11.14.1.2 via ge-0/0/0.0
NODE { AS:17 Member-ASN:65002 ISO:0005.0505.0500.00 ISIS-L2:0 }/1216
                  *[BGP/170] 00:00:23, localpref 100
                    AS path: (65002) I, validation-state: unverified
                  >  to 11.14.1.2 via ge-0/0/0.0


LINK { Local { AS:17 Member-ASN:65001 ISO:0003.0303.0300.00 }.{ IPv4:13.14.1.1
IPv6:abcd::13:14:1:1 } Remote { AS:17 Member-ASN:65001 ISO:0004.0404.0400.00 }.
{ IPv4:13.14.1.2 } ISIS-L2:0 }/1216
                  *[IS-IS/18] 00:00:48
                      Fictitious
LINK { Local { AS:17 Member-ASN:65001 ISO:0003.0303.0300.00 }.{ IPv4:13.14.2.1
IPv6:abcd::13:14:2:1 } Remote { AS:17 Member-ASN:65001 ISO:0004.0404.0400.00 }.
{ IPv4:13.14.2.2 } ISIS-L2:0 }/1216
                  *[IS-IS/18] 00:00:48
                      Fictitious
```

## show route table lsdist.0 detail

```
user@host> show route table lsdist.0 detail
lsdist.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
NODE { AS:200 ISO:1282.2113.1154.00 ISIS-L1:0 }/1216 (1 entry, 1 announced)
*IS-IS  Preference: 15
Level: 1
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5b3054
Next-hop reference count: 14
………..
………..
………..
```

```
Area membership:
47 00 05 80 ff f8 00 00 00 01 08 00 01
SPRING-Capabilities:
- SRGB block [Start: 800000, Range: 4096, Flags: 0xc0]
SPRING-Algorithms:
- Algo: 0
SPRING Flex-Algorithms Definition:
- Flex-Algo: 129
Metric: 0, Calc: 0, priority: 129
- Flags: 0x02, - Inc Any: 0x00040000, - Exclude: 0x00008000, - Inc All: 0x00004000
………….
………….
………….
PREFIX { Node { AS:200 ISO:1282.2113.3158.00 } { IPv4:128.220.13.196/32 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
*IS-IS  Preference: 15
Level: 1
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5b3054
Next-hop reference count: 14
Next hop:
State: <Active NotInstall>
Local AS:   200
Age: 16:16:25
………….
………….
Prefix SID: 10, Flags: 0xe0, Algo: 0
Prefix SID: 780, Flags: 0xe0, Algo: 129
Flex Algo: 129, Flex Algo Metric: 10
………….
………….
IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:4000::/64 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 588
                Address: 0x7660f64
                Next-hop reference count: 12, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 320
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 1d 10:20:38
```

```
                        Validation State: unverified
                        Task: BGP_100.21.0.2.1
                        Announcement bits (1): 0-TED Export
                        AS path: 100 I
                        Accepted
                        SRv6 Locator Flags: 0, Algo: 0, Metric: 0
                        Localpref: 100
                        Router ID: 100.100.100.0
                        Thread: junos-main
```

## show route table lsdist.0 detail

```
user@host> show route table lsdist.1 detail
SPRING Flex-Algorithms Definition:
                - Flex-Algo: 128
                  Metric: 2, Calc: 1, Priority: 10
                   - Flags: 0X80000000, - Inc Any: 0x00000002, - Exclude: 0x00000004, - Inc
All: 0x00000002
```

## show route table lsdist.0 extensive

The output of the `show route table lsdist.0 extensive` command to include IPv6 Prefix attributes.

```
user@host> show route table lsdist.0 extensive
lsdist.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
NODE { AS:100 ISO:0100.0100.0100.00 ISIS-L1:0 }/1216 (1 entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x75ace70
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0x7661124
                Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 0
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 17
                Validation State: unverified
```

```
                  Task: BGP_100.21.0.2.1
                  Announcement bits (1): 0-TED Export
                  AS path: 100 I
                  Accepted
                  IPv4 Router-ids:
                    100.100.100.0
                  Area border router: No
                  External router: No
                  Attached: No
                  Overload: No
                  Hostname: R0
                  Area membership:
                    49 00 05
                  SPRING-Algorithms:
                   - Algo: 0
                   - Algo: 1
                  SRV6 Capable: - Flags: 0
                  SRV6 Node MSD:
                   - Type: 41, Value: 6
                   - Type: 42, Value: 7
                   - Type: 43, Value: 5
                   - Type: 44, Value: 6
                   - Type: 45, Value: 6
                  Localpref: 100
                  Router ID: 100.100.100.0
                  Thread: junos-main

PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv4:100.100.100.0/32 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x0
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0x7661124
                Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                Source: 21.0.2.1
                Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                Session Id: 0
                State: <Active Ext>
                Local AS:   200 Peer AS:   100
                Age: 28
                Validation State: unverified
                Task: BGP_100.21.0.2.1
```

```
                 Announcement bits (1): 0-TED Export
                 AS path: 100 I
                 Accepted
                 Localpref: 100
                 Router ID: 100.100.100.0
                 Thread: junos-main


IPV6 PREFIX { Node { AS:100 ISO:0100.0100.0100.00 } { IPv6:10::10/128 } ISIS-L1:0 }/1216 (1
entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x0
        *BGP    Preference: 170/-101
                 Next hop type: Router, Next hop index: 0
                 Address: 0x7661124
                 Next-hop reference count: 5, key opaque handle: 0x0, non-key opaque handle: 0x0
                 Source: 21.0.2.1
                 Next hop: 21.0.2.1 via ge-0/0/0.0, selected
                 Session Id: 0
                 State: <Active Ext>
                 Local AS:   200 Peer AS:   100
                 Age: 28
                 Validation State: unverified
                 Task: BGP_100.21.0.2.1
                 Announcement bits (1): 0-TED Export
                 AS path: 100 I
                 Accepted
                 Localpref: 100
                 Router ID: 100.100.100.0
                 Thread: junos-main
```

## 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.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 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
```

## Release Information

Command introduced before Junos OS Release 7.4.

Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

### RELATED DOCUMENTATION

show route summary

# show route terse

**IN THIS SECTION**

- Syntax | **705**
- Syntax (EX Series Switches) | **705**
- Description | **705**
- Options | **705**
- Required Privilege Level | **705**
- Output Fields | **706**
- Sample Output | **709**
- Release Information | **710**

## Syntax

```
show route terse
<logical-system (all | logical-system-name)>
```

## Syntax (EX Series Switches)

```
show route terse
```

## 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

# Output Fields

describes the output fields for the `show route terse` command. Output fields are listed in the approximate order in which they appear.

**Table 34: 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><br>• `active` (routes that are active)<br><br>• `holddown` (routes that are in the pending state before being declared inactive)<br><br>• `hidden` (routes that are not used because of a routing policy) |
| *route key* | Key for the state of the route:<br><br>• `+`—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.<br><br>• `-` —A hyphen indicates the last active route.<br><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. |

**Table 34: show route terse Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| V | Validation status of the route:<br><br>• ?—Not evaluated. Indicates that the route was not learned through BGP.<br><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><br>• N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database.<br><br>• V—Valid. Indicates that the prefix and autonomous system pair are found in the database. |
| Destination | Destination of the route. |

**Table 34: show route terse Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| P | Protocol through which the route was learned:<br><br>• `A`—Aggregate<br><br>• `B`—BGP<br><br>• `C`—CCC<br><br>• `D`—Direct<br><br>• `G`—GMPLS<br><br>• `I`—IS-IS<br><br>• `L`—L2CKT, L2VPN, LDP, Local<br><br>• `K`—Kernel<br><br>• `M`—MPLS, MSDP<br><br>• `O`—OSPF<br><br>• `P`—PIM<br><br>• `R`—RIP, RIPng<br><br>• `S`—Static<br><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. |

**Table 34: show route terse Output Fields** *(Continued)*

| Field Name | Field Description |
|---|---|
| Next hop | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. |
| AS path | AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:<br><br>• I—IGP.<br><br>• E—EGP.<br><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
```

## Release Information

Command introduced before Junos OS Release 7.4.