MPLS Feature Guide for QFX Series Switches

Release

14.1X53

Modified: 2016-12-05

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

MPLS Feature Guide for QFX Series Switches
14.1X53

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

Part 2 MPLS

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- QFX Series standalone switches

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**

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

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

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

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/

- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

LDP

CHAPTER 1

Using LDP

LDP Introduction

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

Junos OS LDP Protocol Implementation

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

LDP Operation

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Logical Interfaces*.

Related Documentation

- *Logical Interfaces*

Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- Tunneling LDP LSPs in RSVP LSPs Overview on page 5
- Label Operations on page 5

Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.

Label Operations

Figure 1 on page 6 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *Label Description*.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 1: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs

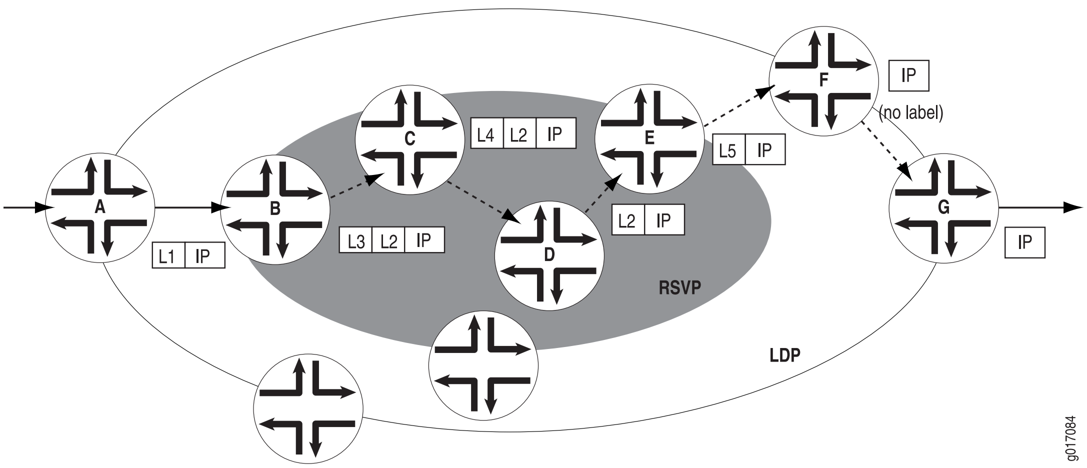

Figure 2 on page 6 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 2: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs

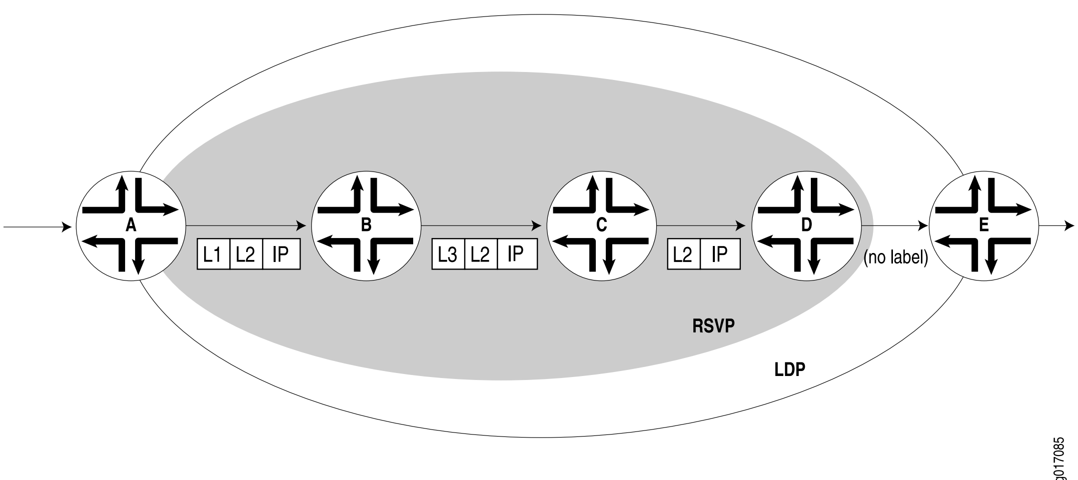

LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- Discovery Messages on page 7
- Session Messages on page 7
- Advertisement Messages on page 7
- Notification Messages on page 8

Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- Basic discovery—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Extended discovery—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

Session Messages

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

Advertisement Messages

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

Notification Messages

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

LDP Session Protection

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router, the LDP sessions are maintained as follows:

1. An LDP session is established between a router and a remote neighboring router.
2. If all of the direct links between the routers go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers is reestablished, the LDP session is not restarted. The routers simply exchange LDP hellos with each other over the direct link. They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.

LDP Graceful Restart

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

Minimum LDP Configuration

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
  interface interface-name;
}
```

To enable LDP on all interfaces, specify **all** for ***interface-name***.

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

Enabling and Disabling LDP

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
  interface interface-name;
}
```

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

To enable LDP on all interfaces, specify **all** for ***interface-name***.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {
  disable;
}
```

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

Enabling Strict Targeted Hello Messages for LDP

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

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

Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

import [*policy-names*];

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

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. Table 3 on page 11 lists the only **from** operators that apply to LDP received-label filtering.

Table 3: from Operators That Apply to LDP Received-Label Filtering

from Operator	Description
interface	Matches on bindings received from a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings received from the specified LDP router ID
next-hop	Matches on bindings received from a neighbor advertising the specified interface address
route-filter	Matches on bindings with the specified prefix

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

`user@host>` **show ldp database**


```
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
  ldp {
    import only-32;
    ...
  }
}
policy-options {
  policy-statement only-32 {
    term first {
      from {
        route-filter 0.0.0.0/0 upto /31;
      }
      then reject;
    }
    then accept;
  }
}
```

Accept **131.108/16** or longer from router ID **10.10.255.2** and accept all prefixes from all other neighbors:

```
[edit]
protocols {
  ldp {
    import nosy-neighbor;
    ...
  }
}
policy-options {
  policy-statement nosy-neighbor {
    term first {
      from {
        neighbor 10.10.255.2;
        route-filter 131.108.0.0/16 orlonger accept;
        route-filter 0.0.0.0/0 orlonger reject;
      }
    }
    then accept;
  }
}
```


Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

export [*policy-name*];

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

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in Table 4 on page 13.

Table 4: to Operators for LDP Outbound-Label Filtering

to Operator	Description
interface	Matches on bindings sent to a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings sent to the specified LDP router ID
next-hop	Matches on bindings sent to a neighbor advertising the specified interface address

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for **10.10.255.6/32** to any neighbors:

```
[edit protocols]
ldp {
  export block-one;
}
policy-options {
  policy-statement block-one {
    term first {
      from {
        route-filter 10.10.255.6/32 exact;
      }
      then reject;
    }
    then accept;
  }
}
```

Send only **131.108/16** or longer to router ID **10.10.255.2**, and send all prefixes to all other routers:

```
[edit protocols]
ldp {
  export limit-lsps;
}
policy-options {
  policy-statement limit-lsps {
    term allow-one {
      from {
        route-filter 131.108.0.0/16 orlonger;
      }
      to {
        neighbor 10.10.255.2;
      }
      then accept;
    }
    term block-the-rest {
      to {
        neighbor 10.10.255.2;
      }
      then reject;
    }
    then accept;
  }
}
```


Specifying the Transport Address Used by LDP

Routers must first establish a TCP session between each other before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

To configure the LDP transport address, include the transport-address statement:

```
transport-address (router-id | interface);
```

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

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

Related Documentation

- transport-address on page 85

Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval interval;
  no-penultimate-hop;
}
```

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

This section includes the following topics:

LDP Statistics Output

The following sample output is from an LDP statistics file:

```
FEC                 Type       Packets     Bytes    Shared
10.255.350.448/32   Transit          0         0        No
                    Ingress          0         0        No
10.255.350.450/32   Transit          0         0       Yes
                    Ingress          0         0        No
10.255.350.451/32   Transit          0         0        No
                    Ingress          0         0        No
220.220.220.1/32    Transit          0         0       Yes
                    Ingress          0         0        No
220.220.220.2/32    Transit          0         0       Yes
                    Ingress          0         0        No
220.220.220.3/32    Transit          0         0       Yes
                    Ingress          0         0        No
May 28 15:02:05, read 12 statistics in  00:00:00 seconds
```

The LDP statistics file includes the following columns of data:

- **read**—Number of bytes of data passed by the FEC since its LSP came up.
- **read**—FEC for which LDP traffic statistics are collected.
- **read**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {
  no-penultimate-hop;
}
```


For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.

NOTE: When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

```
FEC                 Type          Packets          Bytes        Shared
10.255.245.218/32   Transit             0              0            No
                    Ingress             4            246            No
10.255.245.221/32   Transit     statistics disabled
                    Ingress     statistics disabled
13.1.1.0/24         Transit     statistics disabled
                    Ingress     statistics disabled
13.1.3.0/24         Transit     statistics disabled
                    Ingress     statistics disabled
```

LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.
- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 17
- Tracing LDP Protocol Traffic Within FECs on page 18
- Examples: Tracing LDP Protocol Traffic on page 19

Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:


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

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

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory /var/log. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

  You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.
- **path**—Trace label-switched path operations.
- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

Tracing LDP Protocol Traffic Within FECs

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
  ldp {
    traceoptions {
```



```
            file ldp size 10m files 5;
            flag error;
        }
    }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
    ldp {
        traceoptions {
            file ldp size 10m files 5 world-readable;
            flag packets receive;
            flag binding;
        }
        interface all {
        }
    }
}
```

Trace LDP protocol traffic for an FEC associated with the LSP:

```
[edit]
protocols {
    ldp {
        traceoptions {
            flag route filter match-on fec policy filter-policy-for-ldp-fec;
        }
    }
}
```

Example: Configuring LDP Downstream on Demand

This example shows how to configure LDP downstream on demand. LDP is commonly configured using downstream unsolicited advertisement mode, meaning label advertisements for all routes are received from all LDP peers. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between the access and aggregation networks and to reduce the processing requirements for the control plane.

Downstream nodes could potentially receive tens of thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the entire MPLS network, the downstream aggregation node can be configured using LDP downstream on demand to only request the label bindings for the FECs corresponding to the loopback addresses of those egress nodes on which it has services configured.

Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the downstream-on-demand statement at the **[edit protocols ldp session]** hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

Configuration

Configuring LDP Downstream on Demand

Step-by-Step Procedure

To configure a LDP downstream on demand policy and then configure that policy and enable LDP downstream on demand on the LDP session:

1. Configure the downstream on demand policy (DOD-Request-Loopbacks in this example).

   This policy causes the router to forward label request messages only to the FECs that are matched by the DOD-Request-Loopbacks policy.

   ```
   [edit policy-options]
   user@host# set prefix-list Request-Loopbacks 10.1.1.1/32
   user@host# set prefix-list Request-Loopbacks 10.1.1.2/32
   user@host# set prefix-list Request-Loopbacks 10.1.1.3/32
   user@host# set prefix-list Request-Loopbacks 10.1.1.4/32
   user@host# set policy-statement DOD-Request-Loopbacks term 1 from prefix-list
     Request-Loopbacks
   user@host# set policy-statement DOD-Request-Loopbacks term 1 then accept
   ```

2. Specify the DOD-Request-Loopbacks policy using the **dod-request-policy** statement at the **[edit protocols ldp]** hierarchy level.

   The policy specified with the **dod-request-policy** statement is used to identify the prefixes to send label request messages. This policy is similar to an egress policy or an import policy. When processing routes from the inet.0 routing table, the Junos OS software checks for routes matching the **DOD-Request-Loopbacks** policy (in this example). If the route matches the policy and the LDP session is negotiated with DOD advertisement mode, label request messages are sent to the corresponding downstream LDP session.

   ```
   [edit protocols ldp]
   user@host# set dod-request-policy DOD-Request-Loopbacks
   ```

3. Include the **downstream-on-demand** statement in the configuration for the LDP session to enable downstream on demand distribution mode.

   [edit protocols ldp]
   user@host# **set session 1.1.1.1** downstream-on-demand

Distributing LDP Downstream on Demand Routes into Labeled BGP

Step-by-Step Procedure

To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute_ldp** in this example).

   [edit policy-options]
   user@host# **set policy-statement redistribute_ldp term 1 from protocol ldp**
   user@host# **set policy-statement redistribute_ldp term 1 from tag 1000**
   user@host# **set policy-statement redistribute_ldp term 1 then accept**

2. Include the LDP route policy, **redistribute_ldp** in the BGP configuration (as a part of the BGP group configuration **ebgp-to-abr** in this example).

   BGP forwards the LDP routes based on the **redistribute_ldp** policy to the remote PE router

   [edit protocols bgp]
   user@host# **set group ebgp-to-abr type external**
   user@host# **set group ebgp-to-abr local-address 192.168.0.1**
   user@host# **set group ebgp-to-abr peer-as 65319**
   user@host# **set group ebgp-to-abr local-as 65320**
   user@host# **set group ebgp-to-abr neighbor 192.168.6.1 family inet unicast**
   user@host# **set group ebgp-to-abr neighbor 192.168.6.1 family inet labeled-unicast rib inet.3**
   user@host# **set group ebgp-to-abr neighbor 192.168.6.1 export redistribute_ldp**

Step-by-Step Procedure

To restrict label propagation to other routers configured in downstream unsolicited mode (instead of downstream on demand), configure the following policies:

1. Configure the **dod-routes** policy to accept routes from LDP.

   user@host# **set policy-options policy-statement dod-routes term 1 from protocol ldp**
   user@host# **set policy-options policy-statement dod-routes term 1 from tag 1145307136**
   user@host# **set policy-options policy-statement dod-routes term 1 then accept**

2. Configure the **do-not-propagate-du-sessions** policy to not forward routes to neighbors **1.1.1.1**, **2.2.2.2**, and **3.3.3.3**.

   user@host# **set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 1.1.1.1**
   user@host# **set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 2.2.2.2**
   user@host# **set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 3.3.3.3**
   user@host# **set policy-options policy-statement do-not-propagate-du-sessions term 1 then reject**

3. Configure the **filter-dod-on-du-sessions** policy to prevent the routes examined by the **dod-routes** policy from being forwarded to the neighboring routers defined in the **do-not-propagate-du-sessions** policy.

   user@host# **set policy-options policy-statement filter-dod-routes-on-du-sessions term 1 from policy dod-routes**
   user@host# **set policy-options policy-statement filter-dod-routes-on-du-sessions term 1 to policy do-not-propagate-du-sessions**

4. Specify the **filter-dod-routes-on-du-sesssion** policy as the export policy for BGP broup **ebgp-to-abr**.

   [edit protocols bgp]
   user@host# **set group ebgp-to-abr neighbor 192.168.6.2 export filter-dod-routes-on-du-sessions**

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

```
user@host#

show policy-options
prefix-list Request-Loopbacks {
    10.1.1.1/32;
    10.1.1.2/32;
    10.1.1.3/32;
    10.1.1.4/32;
}
policy-statement DOD-Request-Loopbacks {
    term 1 {
        from {
            prefix-list Request-Loopbacks;
        }
        then accept;
    }
}
policy-statement redistribute_ldp {
    term 1 {
        from {
            protocol ldp;
            tag 1000;
        }
        then accept;
    }
}

user@host#

show protocols ldp
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
    downstream-on-demand;
}

user@host#

show protocols bgp
group ebgp-to-abr {
    type external;
```

```
        local-address 192.168.0.1;
        peer-as 65319;
        local-as 65320;
        neighbor 192.168.6.1 {
            family inet {
                unicast;
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export redistribute_ldp;
        }
    }
```

Verification

Verifying Label Advertisement Mode

Purpose Confirm that the configuration is working properly.

Use the **show ldp session** command to verify the status of the label advertisement mode for the LDP session.

Action Issue the **show ldp session** and **show ldp session detail** commands:

- The following command output for the **show ldp session** command indicates that the **Adv. Mode** (label advertisement mode) is **DOD** (meaning the LDP downstream on demand session is operational):

```
user@host> show ldp session
  Address           State        Connection     Hold time  Adv. Mode
1.1.1.2             Operational  Open             22          DOD
```

- The following command output for the **show ldp session detail** command indicates that the **Local Label Advertisement mode** is **Downstream unsolicited**, the default value (meaning downstream on demand is not configured on the local session). Conversely, the **Remote Label Advertisement mode** and the **Negotiated Label Advertisement mode** both indicate that **Downstream on demand** is configured on the remote session

```
user@host> show ldp session detail
Address: 1.1.1.2, State: Operational, Connection: Open, Hold time: 24
  Session ID: 1.1.1.1:0--1.1.1.2:0
  Next keepalive in 4 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: configured-tunneled
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 1.1.1.1, Remote address: 1.1.1.2
  Up for 17:54:52
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: disabled, Helper mode: enabled,
  Remote - Restart: disabled, Helper mode: enabled
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream on demand
  Negotiated Label Advertisement mode: Downstream on demand
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    1.1.1.2
```

Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to

respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see "Configuring the Delay Before LDP Neighbors Are Considered Down" on page 26.

Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

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

Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
  hello-interval seconds;
}
```

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

Configuring the Delay Before LDP Neighbors Are Considered Down

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.

NOTE: By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see "Configuring the LDP Timer for Hello Messages" on page 25.

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an

LDP peer router advertises a shorter hold time than the value you have configured, the peer router's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router. This enables the routers to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

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

Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
  hold-time seconds;
}
```

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

Configuring the Interval for LDP Keepalive Messages

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see "Configuring the Delay Before LDP Neighbors Are Considered Down" on page 26.

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

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

Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

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

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

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

Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- Enabling Graceful Restart on page 29
- Disabling LDP Graceful Restart or Helper Mode on page 29
- Configuring Reconnect Time on page 30
- Configuring Recovery Time and Maximum Recovery Time on page 30

Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router. To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {
  graceful-restart {
    disable;
  }
}
```

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

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {
  graceful-restart {
    helper-disable;
  }
}
```

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

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {
  reconnect-time seconds;
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

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

Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```
graceful-restart {
  maximum-neighbor-recovery-time seconds;
  recovery-time seconds;
}
```

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

Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

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

NOTE: If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems *logical-system-name* policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

```
[edit protocols]
ldp {
  egress-policy connected-only;
}
policy-options {
  policy-statement connected-only {
    from {
      protocol direct;
    }
    then accept;
  }
}
```

Configuring LDP LSP Traceroute

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the **[edit protocols ldp]** hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {
  disable;
  exp exp-value;
  fanout fanout-value;
  frequency minutes;
  paths number-of-paths;
  retries retry-attempts;
  source address;
  ttl ttl-value;
  wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit protocols ldp oam fec *address*]**

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.
- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.

Configuring Miscellaneous LDP Properties

The following sections describe how to configure a number of miscellaneous LDP properties:

Configuring LDP to Use the IGP Route Metric

Use the **track-igp-metric** statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the **track-igp-metric** statement:

```
track-igp-metric;
```

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

Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the **no-forwarding** statement, you can prevent ingress routes from being added to the inet.0 routing table instead of the inet.3 routing table even if you enabled the **traffic-engineering bgp-igp** statement at the **[edit protocols mpls]** or the **[edit logical-systems *logical-system-name* protocols mpls]** hierarchy level. By default, the **no-forwarding** statement is disabled.

To omit ingress routes from the inet.0 routing table, include the **no-forwarding** statement:

```
no-forwarding;
```

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

Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Multiple Instances for Label Distribution Protocol Feature Guide*.

Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the **explicit-null** statement:

```
explicit-null;
```

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

NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see *Label Description* and *Label Allocation*.

Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see “Enabling and Disabling LDP” on page 10). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from source;
      to destination;
      ldp-tunneling;
    }
  }
```



```
}
```

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

Related Documentation

- Tunneling LDP LSPs in RSVP LSPs Overview on page 5

Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the **ignore-lsp-metrics** statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section "Enabling LDP over RSVP-Established LSPs" on page 34.

Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.

To configure an MD5 signature for an LDP TCP connection, include the **session** and **authentication-key** statement:

```
session address {
  authentication-key md5-authentication-key;
}
```

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

Use the **session** statement to configure the address for the remote end of the LDP session.

The ***md5-authentication-key*** (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

You can also configure an authentication key update mechanism for the LDP routing protocol. This mechanism allows you to update authentication keys without interrupting associated routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).

To configure the authentication key update mechanism, include the **key-chain** statement at the **[edit security authentication-key-chains]** hierarchy level, and specify the **key** option to create a keychain consisting of several authentication keys.

```
[edit security authentication-key-chains]
key-chain key-chain-name {
  key key {
    secret secret-data;
    start-time yyyy-mm-dd.hh:mm:ss;
  }
}
```

To configure the authentication key update mechanism for the LDP routing protocol, include the **authentication-key-chain** statement at the **[edit protocols ldp]** hierarchy level to associate the protocol with the **[edit security authentication-key-chains]** authentication keys.

```
[edit protocols ldp]
group group-name {
  neighbor address {
    authentication-key-chain key-chain-name;
  }
}
```

For more information about the authentication key update feature, see *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*.

Configuring LDP Session Protection

An LDP session is normally created between a pair of routers that are connected by one or more links. The routers form one hello adjacency for every link that connects them and associate all the adjacencies with the corresponding LDP session. When the last hello adjacency for an LDP session goes away, the LDP session is terminated. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished.

You can configure the Junos OS to leave the LDP session between two routers up even if there are no hello adjacencies on the links connecting the two routers by configuring the **session-protection** statement. You can optionally specify a time in seconds using the **timeout** option. The session remains up for the duration specified as long as the routers maintain IP network connectivity.

```
session-protection {
```

```
  timeout seconds;
}
```

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

Disabling SNMP Traps for LDP

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *SNMP MIBs and Traps Reference* and *Interpreting the Enterprise-Specific LDP MIB*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```
log-updown {
  trap disable;
}
```

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

Configuring LDP Synchronization with the IGP on LDP Links

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

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

Configuring LDP Synchronization with the IGP on the Router

You can configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

To configure the time the LDP waits before informing the IGP that the LDP neighbor and session are operational, include the **igp-synchronization** statement and specify a time in seconds for the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

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

Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid resignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

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

Ignoring the LDP Subnet Check

In Junos OS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```


This statement can be included at the following hierarchy levels:

- **[edit protocols ldp interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols ldp interface *interface-name*]**

CHAPTER 2

Configuration Statements for LDP

allow-subnet-mismatch

Syntax allow-subnet-mismatch;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp interface *interface-name*],
[edit protocols ldp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.

Default The source address in the LDP link hello packet is matched against the interface address.

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

Related Documentation

- Ignoring the LDP Subnet Check on page 38

authentication-algorithm

Syntax authentication-algorithm *algorithm*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* protocols ldp session *session-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],
[edit logical-systems *logical-system-name* routing-options bmp],
[edit logical-systems *logical-system-name* routing-options bmp station *station-name*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],
[edit protocols ldp session *session-address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit routing-instances *routing-instance-name* protocols ldp session *session-address*],
[edit routing-options bmp],
[edit routing-options bmp station *station-name*]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced for BGP in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric
Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced for BMP in Junos OS Release 13.3.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure an authentication algorithm type.

Options *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

Default: hmac-sha-1-96

NOTE: The default is not displayed in the output of the show bgp bmp command unless a key or key-chain is also configured.

Required Privilege Level

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

Related Documentation

- *Example: Configuring Route Authentication for BGP*
- *Configuring BGP Monitoring Protocol Version 3*

authentication-key (Protocols LDP)

Syntax

authentication-key *md5-authentication-key*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp session *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *address*],
[edit protocols ldp session *address*],
[edit routing-instances *routing-instance-name* protocols ldp session *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

Required Privilege Level

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

Related Documentation

- Configuring the TCP MD5 Signature for LDP Sessions on page 35

authentication-key-chain (Protocols LDP)

Syntax authentication-key-chain *key-chain*;

Hierarchy Level [edit logical-systems *name* protocols ldp session *address*],
[edit logical-systems *name* routing-instances *instance-name* protocols ldp session *address*],
[edit protocols ldp session *address*],
[edit routing-instances *instance-name* protocols ldp session *address*]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the **0.0.0.0/allow** statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

Options ***key-chain***—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

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

Related Documentation

- *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*
- Configuring Miscellaneous LDP Properties on page 33

deaggregate

Syntax deaggregate | no-deaggregate;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Control forwarding equivalence class (FEC) deaggregation on the router. The use of the **deaggregate** statement in LDP is a standard practice that we recommend for LDP deployments.

Default Deaggregation is disabled on the router.

Options **deaggregate**—Deaggregate FECs.

no-deaggregate—Aggregate FECs.

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

Related Documentation

- *Configuring FEC Deaggregation*

disable (Protocols LDP)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp graceful-restart],
[edit logical-systems *logical-system-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options graceful-restart],
[edit protocols ldp graceful-restart],
[edit protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.

Default LDP is enabled on interfaces configured with the LDP **interface** statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the **[edit routing-options]** hierarchy level.

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

Related Documentation

- Enabling and Disabling LDP on page 10
- Configuring LDP Graceful Restart on page 28

dod-request-policy

Syntax dod-request-policy *dod-request-policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit protocols ldp]

Release Information Statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy.

Options ***dod-request-polcy-name***—Specify the name of the downstream on demand request policy.

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

Related Documentation

- Example: Configuring LDP Downstream on Demand on page 20

downstream-on-demand

Syntax downstream-on-demand;

Hierarchy Level [edit logical systems *logical-system-name* protocols ldp session *session-address*],
[edit protocols ldp session *session-address*]

Release Information Statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.

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

Related Documentation

- Example: Configuring LDP Downstream on Demand on page 20

egress-policy

Syntax egress-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Control the prefixes advertised into LDP.

Default Only the loopback address is advertised.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation

- Configuring the Prefixes Advertised into LDP from the Routing Table on page 31

explicit-null (Protocols LDP)

Syntax explicit-null;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Advertise label 0 to the egress router of a label-switched path (LSP).

Default If you do not include the **explicit-null** statement in the MPLS configuration, label 3 (implicit null) is advertised.

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

Related Documentation

- Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 34

export (Protocols LDP)

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation
- Filtering Outbound LDP Label Bindings on page 13

fec

Syntax

```
fec fec-address {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
      remove-nexthop;
      remove-route;
    }
    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  no-bfd-liveness-detection;
  periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-systems-name* protocols ldp oam],
[edit protocols ldp oam]

Release Information

Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 12.2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).

Options

fec-address—Specify the FEC address.

The other statements are explained separately.

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

Related Documentation
- *Configuring BFD for LDP LSPs*

graceful-restart (Protocols LDP)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-neighbor-recovery-time value;
  reconnect-time seconds;
  recovery-time value;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Configure LDP graceful restart on the LDP master protocol instance or for a specific routing instance.

NOTE: When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

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

Related Documentation
- Configuring LDP Graceful Restart on page 28

hello-interval (Protocols LDP)

Syntax

hello-interval *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ldp targeted-hello],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp targeted-hello],
[edit protocols ldp interface *interface-name*],
[edit protocols ldp targeted-hello],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp targeted-hello]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for LDP targeted hellos added in Junos OS Release 9.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the **hello-interval** statement.

Options

seconds—Length of time between transmission of hello packets.

Range: 1 through 65,535 seconds

Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages

Required Privilege Level

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

Related Documentation

- Configuring the LDP Timer for Hello Messages on page 25

helper-disable (LDP)

Syntax helper-disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances *routing-instance-name* protocols ldp graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.

Default Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.

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

Related Documentation

- Configuring LDP Graceful Restart on page 28

hold-time (Protocols LDP)

Syntax

hold-time *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ldp targeted-hello],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp targeted-hello],
[edit protocols ldp interface *interface-name*],
[edit protocols ldp targeted-hello],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp targeted-hello]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for LDP targeted hellos added in Junos OS Release 9.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the **hold-time** statement.

Options

seconds—Hold-time value.

Range: 1 through 65,535 seconds

Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages

Required Privilege Level

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

Related Documentation

- Configuring the Delay Before LDP Neighbors Are Considered Down on page 26

ignore-lsp-metrics

Syntax ignore-lsp-metrics;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts],
[edit protocols ospf traffic-engineering shortcuts]

Release Information Statement introduced in Junos OS Release 7.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Cause OSPF to ignore the RSVP LSP metric.

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.

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

Related Documentation

- Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 35

igp-synchronization

Syntax igp-synchronization holddown-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

Options **holddown-interval** ***seconds***—Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.

Default: 10 seconds

Range: 10 through 60 seconds

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

Related Documentation

- Configuring LDP Synchronization with the IGP on the Router on page 38

import (Protocols LDP)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation

- Filtering Inbound LDP Label Bindings on page 11

interface (Protocols LDP)

Syntax

interface *interface-name* {
  disable;
  hello-interval *seconds*;
  hold-time *seconds*;
  transport-address (interface | loopback);
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Enable LDP on one or more router interfaces.

Default

LDP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify **all**.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Enabling and Disabling LDP on page 10

keepalive-interval

Syntax keepalive-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Set the keepalive interval value.

Options *seconds*—Keepalive value.

Range: 1 through 65,535

Default: 10 seconds

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

Related Documentation

- Configuring the Interval for LDP Keepalive Messages on page 27

keepalive-timeout

Syntax

keepalive-timeout *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.

Options

seconds—Keepalive timeout value.

Range: 1 through 65,535

Default: 30 seconds

Required Privilege Level

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

Related Documentation

- Configuring the LDP Keepalive Timeout on page 28

l2-smart-policy

Syntax l2-smart-policy;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.

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

Related Documentation

- *Configuring LDP IPv4 FEC Filtering*

label-withdrawal-delay

Syntax label-withdrawal-delay *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Virtual Chassis and Virtual Chassis Fabric.

Description Delay the withdrawal of labels to reduce router workload during IGP convergence.

Options ***seconds***—Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.

Default: 60 seconds

Range: 0 through 300 seconds

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

Related Documentation

- Configuring the Label Withdrawal Timer on page 38

ldp

Syntax

```
ldp {
  (deaggregate | no-deaggregate);
  egress-policy [ policy-names ];
  explicit-null;
  export [ policy-names ];
  graceful-restart {
    disable;
    helper-disable;
    maximum-neighbor-recovery-time seconds;
    reconnect-time seconds;
    recovery-time seconds;
  }
  import [ policy-names];
  interface (interface-name | all) {
    disable;
    hello-interval seconds;
    hold-time seconds;
    transport-address (interface | router-id);
  }
  keepalive-interval seconds;
  keepalive-timeout seconds;
  log-updown {
    trap disable;
  }
  no-forwarding;
  oam {
    bfd-liveness-detection {
      detection-time threshold milliseconds;
      ecmp;
      failure-action {
        remove-nexthop;
        remove-route;
      }
      holddown-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      minimum-transmit-interval milliseconds;
      multiplier detection-time-multiplier;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
    }
    fec fec-address {
      bfd-liveness-detection {
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
          remove-nexthop;
          remove-route;
        }
```

```
            holddown-interval milliseconds;
            ingress-policy ingress-policy-name;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            version (0 | 1 | automatic);
          }
          no-bfd-liveness-detection;
          periodic-traceroute {
            disable;
            exp exp-value;
            fanout fanout-value;
            frequency minutes;
            paths number-of-paths;
            retries retry-attempts;
            source address;
            ttl ttl-value;
            wait seconds;
          }
        }
        ingress-policy ingress-policy-name;
        periodic-traceroute {
          disable;
          exp exp-value;
          fanout fanout-value;
          frequency minutes;
          paths number-of-paths;
          retries retry-attempts;
          source address;
          ttl ttl-value;
          wait seconds;
        }
      }
      p2mp;
      policing {
        fec fec-address {
          ingress-traffic filter-name;
          transit-traffic filter-name;
        }
      }
      preference preference;
      session address {
        authentication-algorithm algorithm;
        authentication-key authentication-key;
        authentication-key-chain key-chain-name;
      }
      strict-targeted-hellos;
      traceoptions {
        file filename <files number <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
```

```
  }
  track-igp-metric;
  traffic-statistics {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval interval;
    no-penultimate-hop;
  }
  transport-address (address | interface | router-id);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Enable LDP routing on the router or switch.

You must include the **ldp** statement in the configuration to enable LDP on the router or switch.

Default

LDP is disabled on the router.

Options

The other statements are explained separately.

Required Privilege Level

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

Related Documentation

- Minimum LDP Configuration on page 9
- Enabling and Disabling LDP on page 10

ldp-synchronization

Syntax

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ospf interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf interface *interface-name*],
[edit protocols ospf interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 7.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.

Options

The other statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring LDP Synchronization with the IGP on LDP Links on page 37

ldp-tunneling

Syntax

```
ldp-tunneling;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Enable the LSP to be used for LDP tunneling.

Required Privilege Level

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

Related Documentation

- Enabling LDP over RSVP-Established LSPs on page 34

log-updown (Protocols LDP)

Syntax

```
log-updown {
  trap disable;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Disable LDP traps on the router, logical system, or routing instance.

Options

trap disable—Disable LDP traps.

Default: LDP traps are enabled on the router.

Required Privilege Level

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

Related Documentation

- Disabling SNMP Traps for LDP on page 37

maximum-neighbor-recovery-time

Syntax maximum-neighbor-recovery-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances *routing-instance-name* protocols ldp graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4. Statement changed from **maximum-recovery-time** to **maximum-neighbor-recovery-time** in Junos OS Release 9.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.

Options *seconds*—Configure the maximum recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

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

Related Documentation

- Configuring Recovery Time and Maximum Recovery Time on page 30
- *Configuring Graceful Restart Options for LDP*
- *no-strict-lsa-checking*
- *recovery-time*

no-forwarding

Syntax

no-forwarding;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Do not add ingress routes to the inet.0 routing table even if **traffic-engineering bgp-igp** (configured at the **[edit protocols mpls]** hierarchy level) is enabled.

Default

The **no-forwarding** statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when **traffic-engineering bgp-igp** is enabled.

Required Privilege Level

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

Related Documentation

- Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 33
- *Configuring Virtual-Router Routing Instances in VPNs*

policing (Protocols LDP)

Syntax

```
policing {
  fec fec-address {
    ingress-traffic filter-name;
    transit-traffic filter-name;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Enable policing of forwarding equivalence classes (FECs) for LDP.

Options

fec *fec-address*—Specify the address for the FEC.

ingress-traffic *filter-name*—Specify the name of the filter for policing ingress FEC traffic.

transit-traffic *filter-name*—Specify the name of the filter for policing transit FEC traffic.

Required Privilege Level

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

Related Documentation

- *Configuring Policers for LDP FECs*

preference (Protocols LDP)

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Set the route preference level for LDP routes.

Options *preference*—Preferred value.

Range: 0 through 255

Default: 9

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

Related Documentation

- Configuring LDP Route Preferences on page 28

reconnect-time

Syntax reconnect-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances *routing-instance-name* protocols ldp graceful-restart]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.

Options ***seconds***—Time required for reconnection.

Range: 30 through 300

Default: 60 seconds

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

Related Documentation

- Configuring LDP Graceful Restart on page 28 on *LDP Feature Guide for Routing Devices*
- *Configuring Graceful Restart Options for LDP*

recovery-time

Syntax recovery-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances *routing-instance-name* protocols ldp graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the amount of time a router waits for LDP to restart gracefully.

Options ***seconds***—Configure the recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

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

Related Documentation

- Configuring Recovery Time and Maximum Recovery Time on page 30

session (ldp)

Syntax

session *address* {
  authentication-algorithm *algorithm*;
  authentication-key *authentication-key*;
  authentication-key-chain *key-chain-name*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
authentication-algorithm statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the address for the remote end of the LDP session.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring the TCP MD5 Signature for LDP Sessions on page 35

session-protection

Syntax

session-protection {
  timeout *seconds*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Description

Configure when an LDP session is torn down and resignaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.

Options

timeout ***seconds***—Time in seconds before the LDP session is torn down and resignaled.

Range: 1 through 65,535 seconds

Required Privilege Level

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

Related Documentation

- Configuring LDP Session Protection on page 36

strict-targeted-hellos

Syntax strict-targeted-hellos;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.

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

Related Documentation

- Enabling Strict Targeted Hello Messages for LDP on page 10

targeted-hello

Syntax

```
targeted-hello {
  hello-interval seconds;
  hold-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the LDP timer and LDP hold time for targeted hellos.

Options

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring the LDP Timer for Hello Messages on page 25
- Configuring the Delay Before LDP Neighbors Are Considered Down on page 26

traceoptions (Protocols LDP)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.

match-on address option for the **filter** flag modifier added in Junos OS Release 10.4.

nsr-synchronization and **p2mp-nsr-synchronization** operations for **flag** statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify LDP protocol-level trace options.

Default

The default LDP protocol-level trace options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

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 **ldp-log**. We recommend that you place LDP tracing output in the file **ldp-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.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **address**—Operation of address and address withdrawal messages
- **binding**—Label-binding operations

- **error**—Error conditions
- **event**—Protocol events
- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**— Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
  - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
    - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
    - **fec**—Filter based on the FEC associated with the traced object.
  - **policy *policy-name***—Specify the filter policy.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent all users 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 this 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.

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: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

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

Required Privilege Level

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

Related Documentation

- Tracing LDP Protocol Traffic on page 17
- *Network Management Administration Guide for Routing Devices*

track-igp-metric

Syntax

track-igp-metric;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

Required Privilege Level

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

Related Documentation

- Configuring LDP to Use the IGP Route Metric on page 33

traffic-statistics (Protocols LDP)

Syntax

```
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-penultimate-hop;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.

Options

file *filename*—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

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

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you also must include the **size** statement to specify the maximum file size.

interval *seconds*—(Optional) Specify the interval at which the statistics are polled and written to the file.

Default: 300 seconds (5 minutes)

no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.

no-world-readable—(Optional) Prevent all users from reading the log file.

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

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: 1 MB

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable log file access for all users.

Required Privilege Level

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

Related Documentation

- Collecting LDP Statistics on page 15

transport-address

Syntax transport-address (interface | router-id);

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* protocols ldp interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit protocols ldp interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

Default router-id

Options **interface**—The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

router-id—The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address.

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

Related Documentation

- Specifying the Transport Address Used by LDP on page 15

CHAPTER 3

Monitoring Commands for LDP

clear ldp neighbor

Syntax

clear ldp neighbor
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>

Description

Tear down Label Distribution Protocol (LDP) neighbor connections.

Options

none—Tear down connections with all LDP neighbors for all routing instances.

instance *instance name*—(Optional) Clear the LDP session for the specified routing instance only.

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

neighbor—(Optional) Clear an LDP session for the specified neighbor (IP address) only.

Required Privilege Level

clear

Related Documentation

• show ldp neighbor on page 105

List of Sample Output

clear ldp neighbor on page 88

Output Fields

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

Sample Output

clear ldp neighbor

```
user@host> clear ldp neighbor
```


clear ldp session

Syntax

clear ldp session
<*destination*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Clear Label Distribution Protocol (LDP) sessions.

Options

none—Clear LDP sessions for all destinations for all routing instances.

destination—(Optional) Clear an LDP session for the specified destination (IP address).

instance *instance-name*—(Optional) Clear the LDP session for the specified routing instance only.

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

Required Privilege Level

clear

Related Documentation

- show ldp session on page 113

List of Sample Output

clear ldp session on page 89

Output Fields

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

Sample Output

clear ldp session

```
user@host> clear ldp session
```

clear ldp statistics

Syntax

clear ldp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Set all Label Distribution Protocol (LDP) statistics to zero.

Options

none—Set all LDP statistics to zero for all routing instances.

instance ***instance-name***—(Optional) Clear the LDP session for the specified routing instance only.

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

Required Privilege Level

clear

Related Documentation

- show ldp statistics on page 119
- show ldp traffic-statistics on page 123

List of Sample Output

clear ldp statistics on page 90

Output Fields

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

Sample Output

clear ldp statistics

`user@host>` **clear ldp statistics**

ping mpls ldp

Syntax

ping mpls ldp *fec*
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<p2mp root-addr *ip-address* lsp-id *identifier*>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
size and **sweep** options introduced in Junos OS Release 9.6.
instance option introduced in Junos OS Release 10.0.
p2mp, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Check the operability of MPLS LDP-signaled label-switched path (LSP) connections.
Type Ctrl+c to interrupt a **ping mpls** command.

Options

count ***count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination ***address***—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp ***forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance ***routing-instance-name***—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

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

p2mp root-addr ***ip-address*** **lsp-id** ***identifier***—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

size ***bytes***—(Optional) Size of the LSP ping request packet (**88** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source ***source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *Junos OS MPLS Applications Library for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls ldp fec count on page 92
ping mpls ldp p2mp root-addr lsp-id on page 92

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls ldp p2mp root-addr lsp-id

```
user@host>ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
```

```
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
        Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
        Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```


show ldp database

Syntax

show ldp database
<brief | detail | extensive>
<inet | l2circuit>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<session *session*>
p2mp

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display entries in the LDP database.

Options

none—Display standard information about all entries in the LDP database for all routing instances.

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

inet | l2circuit—(Optional) Display only IPv4 or Layer 2 circuit bindings.

instance *instance-name*—(Optional) Display routing instance information for the specified instance only.

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

session *session*—(Optional) Display database for the specified session only. ***session*** is the destination address of the LDP session.

p2mp—(Optional) Display point-to-multipoint binding information.

Required Privilege Level

view

List of Sample Output

show ldp database (master) on page 96
show ldp database (standby) on page 97
show ldp database l2circuit detail on page 98
show ldp database l2circuit extensive on page 98
show ldp database p2mp (master) on page 98
show ldp database p2mp (standby) on page 99
show ldp database p2mp (master) on page 99
show ldp database p2mp (standby) on page 99
show ldp database session on page 100
show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 100
show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 101

Output Fields Table 5 on page 95 describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

Table 5: show ldp database Output Fields

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>Input label database</td><td>Label received from the other router.</td><td>All levels</td></tr>
<tr><td>Output label database</td><td>Label advertised to the other router.</td><td>All levels</td></tr>
<tr><td>session-identifier</td><td>Session identifier, which includes the local and remote label space identifiers.</td><td>All levels</td></tr>
<tr><td>Label</td><td>Label binding to a route prefix.</td><td>All levels</td></tr>
<tr><td>Prefix</td><td>Route prefix.

It can be one of the following values:

• IP prefix.
• Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
• Layer 2 encapsulation type.

Layer 2 encapsulation types are displayed in the format L2CKT control word status encapsulation-type vc-number, for example, L2CKT CtflWord FRAME RELAY VC 2

• control-word-status—Displays whether the use of the control word has been negotiated for this virtual circuit:
  ◦ NoCtrlWord
  ◦ CtrlWord
• encapsulation-type—Encapsulation type:
  ◦ FRAME RELAY
  ◦ ATM AAL5
  ◦ ATM CELL
  ◦ VLAN
  ◦ ETHERNET
  ◦ CISCO_HDLC
  ◦ PPP
• VC number—Virtual circuit number. It can have any numeric value.
• (Stale)—When you display the LDP database for the neighbor of a restarting router, the bindings leaned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.</td><td>All levels</td></tr>
<tr><td>MTU</td><td>MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.</td><td>detail</td></tr>
</table>

Table 5: show ldp database Output Fields *(continued)*

Field Name	Field Description	Level of Output
VCCV Control Channel types	Virtual Circuit Connection Verification (VCCV) control channel types. • **MPLS router alert label** • **MPLS PW label with TTL=1**	**extensive**
VCCV Control Verification types	The only valid VCCV control verification type is **LSP ping**.	**extensive**
TDM payload size	Size of the Time Division Multiplex (TDM) payload.	All levels
TDM bitrate	Bit rate for the TDM traffic.	All levels
Requested VLAN ID	(VLANs) VLAN identifier of the Layer 2 circuit.	**detail**
Cell bundle size	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	**detail**

State	State of the label binding:
	• **Active**—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.
	• **New**—New label that has not yet been distributed.
	  ◦ **MapRcv**—Waiting to receive a label mapping message.
	  ◦ **MapSend**—Waiting to send a label mapping message.
	  ◦ **RelRcv**—Waiting to receive a label release message.
	  ◦ **RelRsnd**—Waiting to receive a label release message before resending label mapping message.
	  ◦ **RelSend**—Waiting to send a label release message.
	  ◦ **ReqSend**—Waiting to send a label request message.
	  ◦ **W/dSend**—Waiting to send a label withdrawal message.
		detail

Age	Time elapsed since the binding was created.	**detail**

Sample Output

show ldp database (master)

```
user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
 299840      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
      3      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
 299776      L2CKT CtrlWord VLAN VC 100
            MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R
```

```
Bit: 1
            State: Active
            Age: 9:35
            Entropy Label Capability: No
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
      3      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
 299776      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
```

show ldp database (standby)

```
user@host> show ldp database extensive

Input label database, 10.255.107.236:0--10.255.107.234:0
  Label     Prefix
 299808      10.255.107.230/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
 301136      10.255.107.232/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
      3      10.255.107.234/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0
  Label     Prefix
 302480      10.255.107.236/32
            State: Active
            Age: 1d 2:46:36
            Standby binding state:
                Map messages: 1
                Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.234:0
  Label     Prefix
 299904      10.255.107.230/32
```



```
                 State: Active
                 Age: 1d 2:46:36
 299936           10.255.107.232/32
                 State: Active
                 Age: 1d 2:46:36
 299872           10.255.107.234/32
                 State: Active
                 Age: 1d 2:46:36
      3           10.255.107.236/32
                 State: Active
                 Age: 1d 2:46:36
 299952           P2MP root-addr 10.255.107.230, lsp-id 16777217
                 State: Active
                 Age: 1d 2:46:36
```

show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
 Input label database, 10.255.245.44:0--10.255.245.45:0
    Label      Prefix
   100176      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
               Cell bundle size: 80
               State: Active
               Age: 9:48
   100256      L2CKT CtrlWord FRAME RELAY VC 101
               MTU: 4470
               State: Active
               Age: 9:48

 Output label database, 10.255.245.44:0--10.255.245.45:0
    Label      Prefix
   100048      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
               Cell bundle size: 80
               State: Active
               Age: 9:48
   100112      L2CKT CtrlWord FRAME RELAY VC 101
               MTU: 4470
               State: Active
               Age: 9:48
```

show ldp database l2circuit extensive

```
user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
  Label      Prefix
 299872      L2CKT CtrlWord PPP VC 100
             MTU: 4470
             VCCV Control Channel types:
                 MPLS router alert label
                 MPLS PW label with TTL=1
             VCCV Control Verification types:
                 LSP ping
  Label      Prefix
             State: Active
             Age: 19:23:08
```

show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive


Input label database, 10.255.107.232:0--10.255.107.236:0
```

```
  Label     Prefix
 569649      P2MP root-addr 10.255.107.232, lsp-id 16777217
            State: Active
            Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
  Label     Prefix
 299888      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 2d 6:41:35
```

show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
 299968      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 4d 22:21:57
          Standby binding state:
              Map messages: 1
              Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
      3      P2MP root-addr 10.255.107.232, lsp-id 1
            State: Active
            Age: 4d 22:21:57
```

show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive


Input label database, 10.255.107.232:0--10.255.107.236:0
  Label     Prefix
 569649      P2MP root-addr 10.255.107.232, lsp-id 16777217
            State: Active
            Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
  Label     Prefix
 299888      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 2d 6:41:35
```

show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
 299968      P2MP root-addr 10.255.107.230, lsp-id 16777217
```

```
            State: Active
            Age: 4d 22:21:57
          Standby binding state:
               Map messages: 1
               Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
  Label     Prefix
      3      P2MP root-addr 10.255.107.232, lsp-id 1
            State: Active
            Age: 4d 22:21:57
```

show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
  Label     Prefix
 100002     10.255.245.197/32
 100003     10.255.245.196/32
 100004     10.0.0.194/32
      3     10.1.1.195/32
 100000     L2CKT NoCtrlWord FRAME RELAY VC 1
 100001     L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
  Label     Prefix
 100003     10.255.245.197/32
 100004     10.1.1.195/32
 100002     10.255.245.196/32
      3     10.0.0.194/32
 100000     L2CKT CtrlWord FRAME RELAY VC 2
 100001     L2CKT NoCtrlWord FRAME RELAY VC 1
```

show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp database
Input label database, 1.1.1.2:0--1.1.1.3:0
  Label     Prefix
 299808      1.1.1.2/32
      3      1.1.1.3/32
 299792      1.1.1.6/32
 299776      10.255.2.227/32
 299840      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
  Label     Prefix
      3      1.1.1.2/32
 299776      1.1.1.3/32
 299808      1.1.1.6/32
 299792      10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299792      1.1.1.3/32
      3      1.1.1.6/32
 299776      10.255.2.227/32
 299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
 299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
```



```
 299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
  Label     Prefix
      3      1.1.1.2/32
 299776      1.1.1.3/32
 299808      1.1.1.6/32
 299792      10.255.2.227/32
```

show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp database
Input label database, 10.255.2.227:0--1.1.1.3:0
  Label     Prefix
 299808      1.1.1.2/32
      3      1.1.1.3/32
 299792      1.1.1.6/32
 299776      10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
  Label     Prefix
 299856      1.1.1.2/32
 299776      1.1.1.3/32
 299792      1.1.1.6/32
      3      10.255.2.227/32

Input label database, 10.255.2.227:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299792      1.1.1.3/32
      3      1.1.1.6/32
 299776      10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.6:0
  Label     Prefix
 299856      1.1.1.2/32
 299776      1.1.1.3/32
 299792      1.1.1.6/32
      3      10.255.2.227/32
 299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
 299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
 299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7
```


show ldp fec-filters

Syntax

show ldp fec-filters
<*fec*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.

Options

fec—(Optional) Display FEC filter information for the specified FEC.

instance *instance-name*—(Optional) Display FEC filter information for the specified instance.

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

Required Privilege Level

view

List of Sample Output

show ldp fec-filters on page 102

Output Fields

Table 6 on page 102 lists the output fields for the **show ldp fec-filters** command. Output fields are listed in the approximate order in which they appear.

Table 6: show ldp fec-filters Output Fields

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

Sample Output

show ldp fec-filters

```
user@host> show ldp fec-filters 10/8
10.22.1.2/32
     Ingress: f1-10.22.1.2/32 (index: 3)
     Transit: (null) (index: 0)
```

show ldp interface

Syntax

show ldp interface
<brief | detail | extensive>
<*interface-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the status of Label Distribution Protocol (LDP)-enabled interfaces.

Options

none—Display standard status information about all LDP-enabled interface for all routing instances.

interface-name—(Optional) Display information for the specified interface.

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

instance *instance-name*—(Optional) Display information for the specified routing instance.

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

Required Privilege Level

view

List of Sample Output

show ldp interface extensive on page 104

Output Fields

Table 7 on page 103 describes the output fields for the **show ldp interface** command. Output fields are listed in the approximate order in which they appear.

Table 7: show ldp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface name.	All levels
Label space ID	Label space identifier that the router is advertising on the interface.	All levels
Nbr count	Number of neighbors on the interface.	All levels
Next hello	How long until the next hello packet is sent on this interface, in seconds.	All levels
Hello interval	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	**detail** **extensive**

Table 7: show ldp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Hold time	Configured hold time, in seconds.	**detail** **extensive**
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	**extensive**
Local hello interval	Locally configured hello interval.	**extensive**

Sample Output

show ldp interface extensive

```
user@host> show ldp interface extensive
Interface            Label space ID        Nbr count   Next hello
 fe-0/0/3.0           10.255.245.6:0           2           0
   Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
   Local hello interval: 2, Index: 69
```


show ldp neighbor

Syntax

show ldp neighbor
<brief | detail | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<neighbor-address>

Release Information

Command introduced before Junos OS Release 7.4.
neighbor-address option added in Junos OS Release 8.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description

Display Label Distribution Protocol (LDP) neighbor information.

Options

none—Display standard information about LDP neighbors for all routing instances.

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

instance *instance-name*—(Optional) Display information for the specified routing instance.

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

neighbor-address—(Optional) Display information about the specified LDP neighbor.

Required Privilege Level

view

Related Documentation

- clear ldp neighbor on page 88

List of Sample Output

show ldp neighbor extensive on page 106

Output Fields

Table 8 on page 105 describes the output fields for the **show ldp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 8: show ldp neighbor Output Fields

Field Name	Field Description	Level of Output
Address	IP address of the neighbor.	All levels
Interface	Interface over which the neighbor was discovered.	All levels
Label space ID	Label space identifier advertised by the neighbor.	All levels
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	**detail**

Table 8: show ldp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	**detail**
Up for	Length of time the LDP neighbor has been in operation.	**detail extensive**
Reference count	Reference count for the LDP neighbor.	**extensive**
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	**extensive**
Proposed local/peer	Hold time value proposed by the local router and the peer router.	**extensive**

Sample Output

show ldp neighbor extensive

```
user@host> show ldp neighbor extensive
Address            Interface          Label space ID         Hold Time
 192.168.37.23      so-1/0/0.0         10.255.245.5:0           44
   Transport address: 10.255.245.5, Configuration sequence: 6
   Up for 00:03:37
   Reference count: 1
   Hold time: 45, Proposed local/peer: 15/45
```

show ldp path

Syntax

show ldp path
<brief | detail | extensive>
<*destination*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Label Distribution Protocol (LDP) label-switched paths (LSPs).

Options

none—Display standard information about all LDP LSPs for all routing instances.

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

destination—(Optional) Restrict the output to entries that match the specified destination prefix.

instance *instance-name*—(Optional) Display information for the specified routing instance only.

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

Required Privilege Level

view

List of Sample Output

show ldp path extensive on page 108

Output Fields

Table 9 on page 107 describes the output fields for the **show ldp path** command. Output fields are listed in the approximate order in which they appear.

Table 9: show ldp path Output Fields

Field Name	Field Description
Output Session (label)	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.
Input Session (label)	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
route	MPLS route.
Attached route	Route corresponding to the LSP.
Ingress route	The router acts as the ingress for the LSP.

Table 9: show ldp path Output Fields *(continued)*

Field Name	Field Description
Reference count	Reference count for the LDP neighbor.
Transit route	Names of the forwarding equivalence class (FEC) filters on the transit routers.
Global label	MPLS label that is used globally.

Sample Output

show ldp path extensive

```
user@host> show ldp path extensive
Output Session (label)          Input Session (label)
10.255.14.220:0(3)              ( )
   Attached route:  10.255.14.221/32
   Reference count: 3, Global label: 3
10.255.14.220:0(100000)         10.255.14.220:0(3)
   Attached route:  10.255.14.220/32, Ingress route
   Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)         10.255.14.220:0(100001)
   Attached route:  10.255.14.214/32, Ingress route
   Reference count: 2, Transit route, Global label: 100001
```


show ldp route

Syntax

show ldp route
<brief | detail | extensive>
<*destination*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).

Options

none—Display standard information about all entries in the LDP internal topology table for all routing instances.

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

destination—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.

instance *instance-name*—(Optional) Display entries for the specified routing instance only.

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

Required Privilege Level

view

List of Sample Output

show ldp route detail on page 111
show ldp route extensive on page 111

Output Fields

Table 10 on page 109 describes the output fields for the **show ldp route** command. Output fields are listed in the approximate order in which they appear.

Table 10: show ldp route Output Fields

Field Name	Field Description
Destination	Destination prefix.
Next-hop intf/lsp/table	Interface that is the next hop to the destination prefix.
Next-hop address	IP address of the next hop.
Session ID	LDP session ID.

Table 10: show ldp route Output Fields *(continued)*

Field Name	Field Description
Route flags	Information about the route. For example, the **Ingress TTL propagate** flag indicates that the time-to-live (TTL) value is being propagated with the route.
Bound to outgoing label	The route has been bound to LSPs with the label being distributed for that LSP.
Topology entry	The topology that the route is bound to.
Ingress route status	Status of the ingress route. For example, it could be **Active** or **Inactive**.
Last modified	The length of time since the ingress route status last changed.

Sample Output

show ldp route detail

```
user@host> show ldp route 10.255.8.5 detail
Destination          Next-hop intf/lsp                Next-hop address
 10.255.8.5/32       f1
   Session ID 10.255.170.84:0--10.255.170.92:0
                     fe-0/0/0.0                       192.168.100.2
   Session ID 10.255.170.84:0--10.255.8.5:0
                     so-0/2/1.0
   Session ID 10.255.170.84:0--10.255.8.5:0
                     so-0/2/2.0
   Session ID 10.255.170.84:0--10.255.8.3:0
   Bound to outgoing label 299776, Topology entry: 0x8c38a80
   BFD dest addr   BFD state LSP-ping Next-hop addr   Next-hop intf/lsp
   127.0.0.64      up        up       192.168.100.2   fe-0/0/0.0
   127.0.1.64      up        up                       so-0/2/1.0
   127.0.2.64      up        up                       so-0/2/2.0
   127.0.3.64      up        up                       f1
.....
```

show ldp route extensive

```
user@host> show ldp route extensive

Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.0/30         ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.4/30         ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.8/30         ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.12/30        ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.16/30        ge-1/2/0.18
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.18/32
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.20/30        ge-1/2/1.21
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 10.0.0.21/32
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 192.168.0.1/32      ge-1/2/0.18                       10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination          Next-hop intf/lsp/table           Next-hop address
 192.168.0.2/32      ge-1/2/1.21                       10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
```

```
                    ge-1/2/0.18                      10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Route flags: None
Destination         Next-hop intf/lsp/table          Next-hop address
 192.168.0.3/32     ge-1/2/1.21                      10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Route flags: None
Destination         Next-hop intf/lsp/table          Next-hop address
 192.168.0.4/32     ge-1/2/1.21                      10.0.0.22
   Session ID 192.168.0.6:0--192.168.0.4:0
   Bound to outgoing label 299808, Topology entry: 0x92a483c
   Ingress route status: Active, Last modified: 00:01:19 ago
   Route flags: Ingress TTL propagate, Transit TTL propagate
Destination         Next-hop intf/lsp/table          Next-hop address
 192.168.0.5/32     ge-1/2/0.18                      10.0.0.17
   Session ID 192.168.0.6:0--192.168.0.5:0
   Bound to outgoing label 299792, Topology entry: 0x92a47f8
   Ingress route status: Active, Last modified: 00:01:19 ago
   Route flags: Ingress TTL propagate, Transit TTL propagate
Destination         Next-hop intf/lsp/table          Next-hop address
 192.168.0.6/32     lo0.6
   Bound to outgoing label 3, Topology entry: 0x92a4a5c
   Ingress route status: Inactive
   Route type: Egress route
   Route flags: None
```


show ldp session

Syntax

show ldp session
<brief | detail | extensive>
<*destination*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about Label Distribution Protocol (LDP) sessions.

Options

none—Display standard information about all LDP sessions for all routing instances.

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

destination—(Optional) Restrict LDP session display to the specified address.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If ***instance-name*** is omitted, information is displayed for the master instance.

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

Required Privilege Level

view

Related Documentation

- clear ldp session on page 89

List of Sample Output

show ldp session brief on page 117
show ldp session detail on page 117
show ldp session extensive on page 117

Output Fields

Table 11 on page 113 describes the output fields for the **show ldp session** command. Output fields are listed in the approximate order in which they appear.

Table 11: show ldp session Output Fields

Field Name	Field Description	Level of Output
Address	Transport address of the session.	any
State	State of the session: **Nonexistent**, **Connecting**, **Initialized**, **OpenRec**, **OpenSent**, **Operational**, or **Closing**. The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any
Connection	TCP connection state: **Closed**, **Opening**, or **Open**.	any

Table 11: show ldp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Hold time	Time remaining until the session will be closed, in seconds.	any
Session ID	LDP identifiers of the peers of this session.	detail extensive
Next keepalive	Time until next keepalive is sent, in seconds.	**detail extensive**
Active	Whether the local router is playing the active role in the session and during session establishment.	**detail extensive**
Passive	Whether the local router is playing the passive role in the session and during session establishment.	**detail extensive**
Maximum PDU	Maximum protocol data unit (PDU) size (packet size) for the session.	**detail extensive**
Hold time	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the **keepalive-timeout** statement configured at the **[edit protocols ldp]** hierarchy level.	**detail extensive**
Neighbor count	Number of neighbors that are contributing to the session.	**detail extensive**
Keepalive interval	Keepalive interval, in seconds.	**detail extensive**
Connect retry interval	TCP connection retry interval, in seconds.	**detail extensive**
Local address	Local transport address.	**detail extensive**
Remote address	Remote transport address.	**detail extensive**
Up for	Time that this session has been up.	**detail extensive**
Last down	Time since the session last went down.	**detail extensive**

Table 11: show ldp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Reason	Reason the session went down: • **Aborted graceful restart** • **Authentication key was changed** • **Bad type length value (TLV)** • **Bad protocol data unit (PDU) packets** • **Command-line interface (CLI) command** • **Connect time expired** • **Connection error** • **Connection reset** • **Error during initialization** • **Hold time expired** • **No adjacency or all adjacencies down** • **Notification received** • **Received notification from peer** • **Unexpected End of File (EOF)** • **Unknown reason**	**detail extensive**
Number of session flaps	Number of times the session changes from up to down.	**detail extensive**
Restarting	LDP is in the process of gracefully restarting.	**detail extensive**
Capabilities advertised	LDP capabilities advertised to a peer.	**detail extensive**
Capabilities received	LDP capabilities received from a peer.	**detail extensive**
Protection	Information about the status of MPLS LDP session protection.	**detail extensive**
restart complete in *nnn* msec	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	**detail extensive**
Local	Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. • **Restart**—Status of the grateful restart feature at the local end of the LDP session: **enabled** or **disabled**. • **Helper mode**—Status of the helper mode feature at the local end of the LDP session: **enabled** or **disabled**. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures. • **Reconnect time**—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is **60000 msec** and is not configurable. (**Reconnect timeout** refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, *Internet Draft Graceful Restart Mechanism for LDP.)*	**detail extensive**

Table 11: show ldp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Remote	Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent. • **Restart**—Status of the grateful restart feature at the remote end of the LDP session: **enabled** or **disabled**. • **Helper mode**—Status of the helper mode feature at the remote end of the LDP session: **enabled** or **disabled**. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures. • **Reconnect time**—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.	**detail extensive**
Local maximum recovery time	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	**detail extensive**
Next-hop addresses received	Next-hop addresses received on the session.	**detail extensive**
Queue depth	Number of messages that are queued for sending to the peers in the group.	**extensive**
Message type	Type of message being sent: • ◦ **Initialization**—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established. ◦ **Keeaplive**—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them. ◦ **Notification**—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer. ◦ **Address**—Message sent by an LSR to an LDP peer to advertise interface addresses. ◦ **Address withdraw**—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address. ◦ **Label mapping**—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC). ◦ **Label request**—Message sent by an LSR to an LDP peer to request a label mapping for an FEC. ◦ **Label withdraw**— Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping. ◦ **Label release**—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released. ◦ **Label abort**—Message sent by an LSR to an LDP peer to abort a label request message. • **Total**—Messages sent and received during the lifetime of the session. • **Last 5 seconds**—Messages sent and received during the current session.	**extensive**

Sample Output

show ldp session brief

```
user@host> show ldp session brief
  Address           State        Connection     Hold time
10.255.72.160       Operational  Open             21
10.255.72.164       Operational  Open             20
10.255.72.172       Operational  Open             21
```

show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 7 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3
  Up for 00:00:02
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    10.0.0.5
    10.0.0.33
```

show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 2 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3
  Up for 00:05:37
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    10.0.0.5
    10.0.0.33
```



```
  Queue depth: 0
Message type              Total                       Last 5 seconds
                       Sent      Received            Sent      Received
Initialization            1             1               0             0
Keepalive                33            33               1             1
Notification              0             0               0             0
Address                   1             1               0             0
Address withdraw          0             0               0             0
Label mapping             7             5               0             0
Label request             0             0               0             0
Label withdraw            3             1               0             0
Label release             1             3               0             0
Label abort               0             0               0             0
```

show ldp statistics

Syntax

show ldp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Label Distribution Protocol (LDP) statistics.

Options

none—Display LDP statistics for all routing instances.

instance *instance-name*—(Optional) Display information for the specified routing instance only.

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

Required Privilege Level

view

Related Documentation

- clear ldp statistics on page 90

List of Sample Output

show ldp statistics on page 122

Output Fields

Table 12 on page 119 lists the output fields for the **show ldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 12: show ldp statistics Output Fields

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.

Table 12: show ldp statistics Output Fields *(continued)*

Field Name	Field Description
Message type	LDP message types: • **Hello**—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor. • **Initialization**—Messages that indicate an LDP session has started. • **Keepalive**—Messages that ensure that the keepalive timeout is not exceeded. • **Notification**—Advisory information and signal error information. • **Address**—Messages with address information. • **Address withdrawal**—Messages regarding address withdrawal. • **Label mapping**—Messages with label mapping information. • **Label request**—Request for a label mapping from a neighboring router. • **Label withdrawal**—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • **Label release**—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • **Label abort**—Messages about label interruptions. • **All UDP**—All hello messages sent by LSRs to the well-known UDP port, 646. • **All TCP**—All LDP session messages.

Table 12: show ldp statistics Output Fields *(continued)*

Field Name	Field Description
Event type	LDP events and errors: • **Sessions opened**—Number of LDP sessions that have been opened. • **Sessions closed**—Number of LDP sessions that have been closed. • **Topology changes**—Number of changes to the known LDP topology. • **No interface**—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages. • **No session**—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers. • **No adjacency**—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected. • **Unknown version**—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. • **Malformed PDU**—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session. An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length). • **Malformed message**—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them. An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code. An LDP message is considered to be malformed If the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU. The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component. The LDP message is considered to be malformed if the message is missing one or more mandatory parameters. • **Unknown message type**—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message. • **Inappropriate message**—The message is not of the type that the receiver expects to receive. • **Malformed TLV**—The TLV lLength is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR. • **Bad TLV value**—The TLV Length is too large. • **Missing TLV**—The TLV is missing one or more mandatory parameters. • **PDU too large**—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.
Total	Total number of each event or error.
Last 5 seconds	Number of each event or error in the last 5 seconds.

Sample Output

show ldp statistics

```
user@host> show ldp statistics
Message type                  Total                       Last 5 seconds
                          Sent        Received          Sent       Received
Hello                      265           263               2              2
Initialization               2             2               0              0
Keepalive                  112           111               1              0
Notification                 0             0               0              0
Address                      2             2               0              0
Address withdraw             0             0               0              0
Label mapping                7             6               0              0
Label request                0             0               0              0
Label withdraw               2             0               0              0
Label release                0             2               0              0
Label abort                  0             0               0              0
All UDP                    265           263               2              2
All TCP                    123           121               1              0

Event type                              Total        Last 5 seconds

Sessions opened                             2                   0
Sessions closed                             0                   0
Topology changes                           11                   0

No interface                                0                   0
No session                                  0                   0
No adjacency                                0                   0
Unknown version                             0                   0
Malformed PDU                               0                   0
Malformed message                           0                   0
Unknown message type                        0                   0
Inappropriate message                       0                   0
Malformed TLV                               0                   0
Bad TLV value                               0                   0
Missing TLV                                 0                   0
PDU too large                               0                   0
```


show ldp traffic-statistics

Syntax

show ldp traffic-statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<p2mp>

Release Information

Statement introduced before Junos OS Release 7.4.
p2mp option added in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Label Distribution Protocol (LDP) traffic statistics.

NOTE: **If nonstop active routing features is configured, show ldp traffic-statistics command is not supported on backup Routing Engines.**

Options

none—Display LDP traffic statistics for all routing instances.

instance *instance-name*—(Optional) Display LDP traffic statistics for the specified routing instance only.

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

p2mp—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.

Additional Information

To collect output from this command on a periodic basis, configure the traffic-statistics statement for the LDP protocol. For more information, see the *Junos MPLS Applications Configuration Guide*.

Required Privilege Level

view

Related Documentation

- clear ldp statistics on page 90
- *Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain*
- *Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs*

List of Sample Output

show ldp traffic-statistics on page 124
show ldp traffic-statistics p2mp on page 125
show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 125
show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 125

Output Fields Table 13 on page 124 lists the output fields for the **show ldp traffic-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 13: show ldp traffic-statistics Output Fields

Field Name	Field Description
Message type	LDP message types.
FEC	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (**root_addr:lsp_id**). For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address (**root_addr:lsp_id/grp,src**).
Type	Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).
Packets	Number of packets passed by the FEC since its LSP came up.
Bytes	Number of bytes of data passed by the FEC since its LSP came up.
Shared	Whether a label is shared by prefixes: **Yes** or **No**. A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
Nexthop	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)
Label	For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop. 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.
Backup route	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.

Sample Output

show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
 FEC                Type          Packets           Bytes    Shared
  10.35.3.0/30      Transit             0               0     Yes
```

```
                     Ingress              0                  0    No
10.35.10.1/32       Transit               0                  0    Yes
                     Ingress              0                  0    No
 10.255.245.214/32   Transit              0                  0    No
                     Ingress             11                752    No
 192.168.37.36/30    Transit               0                  0    Yes
                     Ingress               0                  0    No

FEC(root_addr:lsp_id)  Nexthop            Packets           Bytes     Shared
10.255.72.160:16777217 192.168.8.81         152056           14597376    No
                       192.168.8.1          152056           14597376    No
                       192.168.8.65         152056           14597376    No
```

show ldp traffic-statistics p2mp

```
user@host> show ldp traffic-statistics p2mp
FEC(root_addr:lsp_id) Nexthop             Packets             Bytes Shared
10.255.72.160:16777217192.168.8.81         152056          14597376    No
                      192.168.8.1          152056          14597376    No
                      192.168.8.65         152056          14597376    No
```

show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp traffic-statistics p2mp
P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)     Nexthop                Packets              Bytes
 Shared
11.99.0.73:239.10.0.1,11.98.0.10  11.99.0.117             243408          121217184
 No
                                  11.99.0.13              236286          117670428
 No
11.99.0.73:239.10.0.2,11.98.0.10  11.99.0.117             248800          123902400
 No
                                  11.99.0.13              240759          119897982
 No
11.99.0.73:239.10.0.1,11.98.0.20  11.99.0.117             250286          124642428
 No
                                  11.99.0.13              243741          121383018
 No
11.99.0.73:239.10.0.2,11.98.0.20  11.99.0.117             252970          125979060
 No
                                  11.99.0.13              245218          122118564
 No
```

show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show ldp traffic-statistics p2mp

P2MP FEC Statistics:
```

```
FEC(root_addr:lsp_id/grp,src)     Nexthop             Packets             Bytes
 Shared
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568
                                  1.3.8.2                   0                 0
 No
                                  1.3.4.2                   0                 0
 No
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route
                                  1.3.4.2                   0                 0
 No
                                  1.3.8.2                   0                 0
 No
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600
                                  1.3.8.2                   0                 0
 No
                                  1.3.4.2                   0                 0
 No
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route
                                  1.3.4.2                   0                 0
 No
                                  1.3.8.2                   0                 0
 No
```


traceroute mpls ldp

Syntax

traceroute mpls <ldp> *fec*
<destination>
<detail>
<exp>
<fanout>
<logical-system>
<no-resolve>
<paths>
<retries>
<routing-instance>
<source>
<ttl>
<update>
<wait>

Release Information

Command introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use **traceroute mpls ldp** as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

fec—Specify the IP address and optional prefix of the forwarding equivalence class (FEC).

destination—(Optional) Specify the destination address to use when sending probes.

detail—(Optional) Display detailed output.

exp—(Optional) Specify the class-of-service to use when sending probes. The range of values is **0** through **7**. The default value is **7**.

fanout—(Optional) Specify the maximum number of nexthops to search per node. The range of values is **1** through **16**. The default value is **16**.

logical-system—(Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths—(Optional) Specify the number of paths to search. The range of values is **1** through **255**. The default value is **16**.

retries—(Optional) Specify the number of times to resend probe. values. The range of values is **1** through **9**. The default value is **3**.

routing-instance ***routing-instance-name***—(Optional) Specify the name of the routing instance for the traceroute attempt.

source ***source-address***—(Optional) Specify the source address of the outgoing traceroute packets.

ttl ***value***—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. The range of values is **1** through **125** and the default value is **64**.

wait ***seconds***—(Optional) Specify the number of seconds to wait before resending a probe. The range of values is **5** through **15** and the default value is **10** seconds.

Required Privilege Level network

List of Sample Output
traceroute mpls ldp on page 129
traceroute mpls ldp detail on page 129

Output Fields Table 14 on page 128 describes the output fields for the **traceroute mpls ldp** ***fec*** command and the **traceroute mpls ldp** ***fec*** **detail** commands. Output fields are listed in the approximate order in which they appear.

Table 14: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the **traceroute mpls ldp** ***fec*** command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is **null**.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is **null**.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**
Response time	Time for the echo request to reach the receiver.	**detail**

Table 14: traceroute mpls ldp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is **none**.	**detail**
Label Stack	Label stack used to forward the packet.	**detail**

Sample Output

traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4

  Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
  ttl     Label  Protocol     Address          Previous Hop     Probe Status
    1    100016  LDP          24.24.24.1       (null)           Success
    2    100000  LDP          20.20.20.2       24.24.24.1       Success
    3         3  LDP          22.22.22.4       20.20.20.2       Egress

Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail

  Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
  Hop 24.24.24.1 Depth 1
    Parent (null)
    Return code: Label switched at stack-depth 1
    Response time 165.93 msec
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.0 ~ 127.0.3.255
    Label Stack:
      Label 1 Value 100032 Protocol LDP

  Hop 20.20.20.2 Depth 2
    Parent 24.24.24.1
    Return code: Upstream interface index unknown label-switched at stack-depth
1
    Response time 19.05 msec
    Multipath type: IP bitmask
      Address Range 1: 127.0.0.0 ~ 127.0.3.255
    Label Stack:
      Label 1 Value 100000 Protocol LDP

  Hop 22.22.22.4 Depth 3
    Parent 20.20.20.2
    Return code: Egress-ok at stack-depth 1
    Response time 0.79 msec
    Multipath type: None
    Label Stack:
      Label 1 Value 3 Protocol LDP
```

PART 2

MPLS

CHAPTER 4

Using MPLS

MPLS Overview For QFX Switches and EX4600

MPLS is a protocol that uses labels to route packets instead of using IP addresses. In a traditional network, each switch performs an IP routing lookup, determines a next-hop based on its routing table, and then forwards a packet to that next-hop. With MPLS, only the first device does a routing lookup, and, instead of finding the next-hop, finds the ultimate destination along with a path to that destination. The path of an MPLS packet is called a label-switched path (LSP).

MPLS applies one or more labels to a packet so it can follow the LSP to the destination. Each switch pops off its label and sends the packet to the next switch label in the sequence.

The Junos OS includes everything you need to configure MPLS. You do not need to install any additional programs or protocols. MPLS is supported on switches with a subset of the commands supported on routers. The Junos MPLS-configured switches can interact with each other and with Junos MPLS-configured routers.

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports can be assigned different labels.
- A packet arriving at a particular provider edge (PE) switch can be assigned a label that is different from that of the same packet entering the network at a different PE switch. As a result, forwarding decisions that depend on the ingress PE switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

This topic describes:

Why Use MPLS?

MPLS reduces the use of the forwarding table by using labels instead of the forwarding table. The size of forwarding tables on a switch are limited by silicon and using exact matching for forwarding to destination devices is cheaper than buying more sophisticated hardware. In addition, MPLS allows you to control where and how traffic is routed on your network – this is called traffic engineering.

Some reasons to use MPLS instead of another switching solution are:

- MPLS can connect different technologies that would not otherwise be compatible---service providers have this compatibility issue when connecting clients with different autonomous systems in their networks. In addition, MPLS has a feature called Fast Reroute that provides alternate backups for paths – this prevents network degradation in case of a switch failure.
- • Other IP-based encapsulations such as Generic Route Encapsulation (GRE) or Virtual Extensible Local Area Networks (VXLAN) support only two levels of hierarchy, one for the transport tunnel and one piece of metadata. Using virtual servers means that you need multiple hierarchy levels. For example, one label is needed for top-of-rack (ToR), one label for the egress port that identifies the server, and one for the virtual server.

Why Not Use MPLS?

There are no protocols to auto-discover MPLS enabled nodes. MPLS protocol just exchanges label values for an LSP. They do not create the LSPs.

You must build the MPLS mesh, switch by switch. We recommend using scripts for this repetitive process.

MPLS hides suboptimal topologies from BGP where multiple exits may exist for the same route.

Large LSPs are limited by the circuits they traverse. You can work around this by creating multiple, parallel LSPs.

How Do I Configure MPLS?

There are three types of switches you must set up for MPLS:

- Label Edge Router/Switch (LER) or ingress node to the MPLS network. This switch encapsulates the packets.
- Label Switching Routers/Switches (LSR). One or more switches that transfer MPLS packets in the MPLS network.
- Egress router/switch is the final MPLS device that removes the last label before packets leave the MPLS network.

Service providers (SP) use the term provider router (P) for a backbone router/switch doing label switching only. The customer-facing router at the SP is called a provider edge router (PE). Each customer needs a customer edge router (CE) to communicate with the PE. Customer facing routers typically can terminate IP addresses, L3VPNs, L2VPNs/ pseudowires, and VPLS before packets are transferred to the CE.

Configure the MPLS LER (Ingress) Switch and the Egress Switch

To configure MPLS, you must first create one or more named paths on the ingress and egress routers. For each path, you can specify some or all transit routers in the path, or you can leave it empty.

Configure LSRs for MPLS

Configure one or more MPLS LSRs by following these steps:

1. Configure interfaces on each switch to transmit and receive MPLS packets using the usual interface command with MPLS appended. For example:

   ```
   [edit interfaces ge-0/0/0 unit 0] family mpls;
   ```

2. Add those same interfaces under [edit protocols mpls]. For example:

   ```
   [edit protocols mpls]
     interface ge-0/0/0;
   ```

3. Configure the interfaces on each switch to handle MPLS labels with a protocol. For example, for LDP:

   ```
   [edit protocols ldp]
     Interface ge-0/0/0.0;
   ```

   To watch a demo of these configurations, see https://www.youtube.com/watch?v=xegWBCUJ4tE.

What Does the MPLS Protocol Do?

Multiprotocol Label Switching (MPLS) is an Internet Engineering Task Force (IETF)-specified framework that provides for the designation, routing, forwarding and switching of traffic flows through the network. In addition, MPLS:

- Specifies mechanisms to manage traffic flows of various granularities, such as flows between different hardware, machines, or even flows between different applications.
- Remains independent of the layer-2 and layer-3 protocols.
- Provides a means to map IP addresses to simple, fixed-length labels used by different packet-forwarding and packet-switching technologies.
- Interfaces to existing routing protocols, such as Resource ReSerVation Protocol (RSVP) and Open Shortest PathFirst (OSPF).
- Supports IP, ATM, and Frame Relay layer-2 protocols.
- Uses these additional technologies:
  - FRR: MPLS Fast Reroute improves convergence during a failure by mapping out alternate LSPs in advance.
  - Link Protection/ Next-hop backup: A bypass LSP is created for every possible link failure.

- Node Protection/ Next-hop backup: A bypass LSP is created for every possible switch (node) failure.
- VPLS: Creates Ethernet multipoint switching service over MPLS and emulates functions of an L2 switch.
- L3VPN: IP-based VPN customers get individual virtual routing domains.

How Does MPLS Interface to Other Protocols?

Some of the protocols that work with MPLS are:

- RSVP-TE: Resource Reservation Protocol - Traffic Engineering reserves bandwidth for LSPs.
- LDP: Label Distribution Protocol is the defacto protocol used for distribution of MPLS packets and is usually configured to tunnel inside RSVP-TE.
- IGP: Interior Gateway Protocol is a routing protocol. Edge routers (PE-routers) run BGP between themselves to exchange external (customer) prefixes. Edge and core (P) routers run IGP (usually OSPF or IS-IS) to find optimum path toward BGP next hops. P- and PE-routers use LDP to exchange labels for known IP prefixes (including BGP next hops). LDP indirectly builds end-to-end LSPs across the network core.
- BGP: Border Gateway Protocol (BGP) is allows policy-based routing to take place, using TCP as its transport protocol on port 179 to establish connections. The Junos OS routing protocol software includes BGP version 4. You do not configure BGP---configuring interfaces with MPLS and LDP/RSVP establishes the labels and the ability to transmit packets. BGP automatically determines the routes packets take.
- OSPF and ISIS: These protocols are used for routing between the MPLS PE and CE. Open Shortest Path First (OSPF) is perhaps the most widely used interior gateway protocol (IGP) in large enterprise networks. IS-IS, another link-state dynamic routing protocol, is more common in large service provider networks. Assuming you're running L3VPN to your customers, on the SP edge between the PE and the CE you can run any protocol that your platform supports as a VRF aware instance.

If I Have Used Cisco MPLS, What Do I Need to Know?

Cisco Networks and Juniper Networks use different terminology.

What Cisco Calls:	Juniper Calls:
affinities	admin-groups
autoroute announce	TE shortcuts
forwarding adjacency	LSP-advertise
tunnel	LSP
make-before-break	adaptive

What Cisco Calls:	Juniper Calls:
application-window	adjust-interval
shared risk link groups	fate-sharing

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding MPLS Components on page 144
- Understanding MPLS Label Operations on page 148
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- *Junos OS MPLS Applications Library for Routing Devices*

MPLS Feature Support on QFX Series and EX4600 Switches

Multiprotocol Label Switching (MPLS) is a set of procedures for augmenting network layer packets with label stacks, thereby turning them into labeled packets. Service providers frequently use MPLS. Simply put, where traditional networks' routers each perform an IP lookup to determine the next hop, an MPLS network's first device does a routing lookup for the final destination instead of the next hop. A label is then applied to the packet—this is called packet switching. The final destination device removes the label.

A number of Juniper Networks switches are capable of running a subset of MPLS and can therefore communicate, not only with each other, but with Juniper Networks routers running MPLS. This topic describes the major MPLS features that are supported on QFX Series switches and on one EX Series switch, the EX4600. Be sure to check for any exceptions to this support in "MPLS Limitations on QFX Series and EX4600 Switches" on page 141.

NOTE: EX4600 uses the same chipset as QFX5100—this is why the EX4600 switch is discussed here along with QFX Series switches. Other EX Series switches also support MPLS but with a different feature set.

This topic describes:

- MPLS Commands Supported by QFX Series and EX4600 Switches on page 138
- MPLS Features Supported by QFX Series and EX4600 Switches on page 139

MPLS Commands Supported by QFX Series and EX4600 Switches

QFX Series and EX4600 switches support a subset of MPLS features. The command-line interface (CLI) for switches displays all MPLS related configuration statements, even those that are not supported. However, configuring those unsupported statements on a switch has no effect on the operation of the switch.

MPLS Features Supported by QFX Series and EX4600 Switches

Table 15 on page 139 lists the major MPLS features supported on QFX Series and EX4600 switches. It also lists the Juniper Networks Junos operating system (Junos OS) release in which they were introduced.

Table 15: MPLS Features with Junos OS Release Support

Feature	QFX3500	QFX5100	EX4600
QFX standalone switch as an MPLS provider edge (PE) switch or provider switch	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
MPLS provider edge (PE) switch or provider switch	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Route reflector for BGP labeled routes	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Label edge router (LER)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Label switch router (LSR)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Border Gateway Protocol (BGP) labeled unicast	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Carrier-over-carrier and inter-provider BGP L3 VPN	14.1X53-D15	14.1X53-D15 VC/VCF (14.1X53-D30)	not supported
Class of Service (CoS or QoS) for MPLS traffic	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Ethernet-over-MPLS (L2 circuit)	14.1X53-D10	14.1X53-D10 VC/VCF (14.1X53-D30)	15.1X53-D15
Entropy labels	not supported	not supported	not supported
Fast Reroute (FRR),one-to-one local protection and many-to-one local protection	14.1X53-D10	14.1X53-D10 VC/VCF (not supported)	not supported
FRR using detours and secondary LSP	not supported	not supported	not supported

Table 15: MPLS Features with Junos OS Release Support *(continued)*

Feature	QFX3500	QFX5100	EX4600
Firewall filters	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Flow-Aware Transport of Pseudowires (FAT) Flow Labels	not supported	not supported VC/VCF (not supported)	not supported
Graceful restart for Open Shortest Path First (OSPF) and Resource Reservation Protocol (RSVP)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	13.2X51-D25
Intermediate System to Intermediate System (IS-IS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
IP-over-MPLS label-switched paths (LSPs) both static and dynamic links	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
IPv6 over an MPLS IPv4 network (6PE) Layer 3 VPN 6PE	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
IPv6 over an MPLS core network	not supported	not supported VC/VCF (not supported)	not supported
Label Distribution Protocol (LDP) signaling over RSVP	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Layer 3 VPNs IPv4	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Integrated bridging and routing (IRB) over an MPLS core network	not supported	14.1X53-D40	not supported
MTU signaling in RSVP	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Object access method (OAM) including MPLS ping, traceroute and BFD	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Open Shortest Path First traffic engineering (OSPF TE)	14.1X53-D10	14.1X53-D15 VC/VCF (not supported)	14.1X53-D15

Table 15: MPLS Features with Junos OS Release Support *(continued)*

Feature	QFX3500	QFX5100	EX4600
OSPFv2 as an interior gateway protocol	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	13.2X51-D25
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	14.1X53-D10	14.1X53-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Resource Reservation Protocol (RSVP), bandwidth and auto-bandwidth	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP	14.1X53-D15	14.1X53-D15 VC/VCF (not supported)	not supported
RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
SNMP MIB support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15
Static and dynamic LSPs	12.2X50-D10	13.2X51-D10 VC/VCF (14.1X53-D30)	14.1X53-D15
Traffic Engineering (TE)	13.1X51-D10	13.1X51-D10	14.1X53-D15
TE auto-bandwidth and RSVP bandwidth	13.1X51-D10	13.1X51-D10	14.1X53-D15
Virtual routing and forwarding (VRF) label support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	14.1X53-D15

Related Documentation

- MPLS Limitations on QFX Series and EX4600 Switches on page 141
- MPLS Configuration Guidelines on page 217

MPLS Limitations on QFX Series and EX4600 Switches

Multiprotocol Label Switching (MPLS) is fully implemented on routers, while switches support a subset of the MPLS features. The limitations of each switch are listed in a separate section here, even though many of the limitations are duplicates that apply to more than one switch.

MPLS Limitations on QFX3500 Switches

- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute is supported, however the **include-all** and **include-any** options for fast reroute are not supported. For more information, see "Fast Reroute Overview" on page 166.
- MPLS-based circuit cross-connects (CCC) is not supported—only circuit-based pseudowires are supported.
- MTU signaling in RSVP and discovery is supported in Control Plane. However, this cannot be enforced in data plane.
- With Layer 2 (L2) circuit-based pseudowires, if multiple equal-cost RSVP label-switched paths (LSPs) are available to reach a Layer 2 Circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2 circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.

MPLS Limitations on QFX5100 and EX4600 Switches

- On a QFX5100 switch, you can observe traffic drop after changing your configuration to enable VLAN tagged for MPLS packets. As a result of packet capture, a QFX5100 switch can swap the wrong VLAN ID for MPLS packets.
- Even though both QFX5100 and EX4600 switches use the same chipset, MPLS support differs. EX4600 switches support only basic MPLS functionality while QFX5100 switches support some of the more advanced features. See "MPLS Feature Support on QFX Series and EX4600 Switches" on page 138 for details.
- On a QFX5100 switch, configuring Integrated Bridging and Routing (IRB) interfaces on the MPLS core is implemented on the switch by using TCAM rules. This is due to a chip limitation on the switch that only allows for a limited amount of TCAM space. 1K TCAM space is allocated for IRB. If multiple IRBs exist, make sure that you have enough available TCAM space on the switch. To check the TCAM space, see TCAM Filter Space Allocation and Verification in QFX Devices.

- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute (FRR) on regular interfaces is supported, however the **include-all** and **include-any** options for FRR are not supported. For more information, see "Fast Reroute Overview" on page 166.
- FRR is not supported on MPLS over IRB interfaces.
- MPLS-based circuit cross-connects (CCC) is not supported—only circuit-based pseudowires are supported.
- Configuring link aggregation groups (LAGs) on UNI ports for Layer 2 circuits is not supported.
- MTU signaling in RSVP and discovery is supported in the control plane. However, this cannot be enforced in data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSP's are available to reach a Layer 2 circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2 circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- These are hardware limitations for both EX4600 MPLS and QFX5100 MPLS switches:
  - Push of a maximum of 3 labels is supported in MPLS edge switch if label swap is not done.
  - Push of a maximum of 2 labels is supported in MPLS edge switch if label swap is done.
  - Pop at line rate is supported is supported for a maximum of 2 labels.
  - Global label space is supported but interface-specific label space is not supported.
  - MPLS ECMP on PHY node with BOS=1 is not supported for single labels.
  - QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, QFX5100, and QFX5200 switches.
  - On the QFX5100 switch, the MPLS MTU command can cause unexpected behavior—this is due to SDK chipset limitations on this platform.
- These LDP features are not supported on QFX5100 switches:
  - LDP multipoint
  - LDP link protection

- LDP bidirectional forwarding detection (BFD)
- LDP operation administration and management (OAM)
- LDP multicast-only fast reroute (MoFRR)

MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric

The following MPLS features are not supported by QFX5100 VC and QFX5100 VCF:

- Next-hop LSP Details in section 2.2.6.1
- BFD including BFD triggered FRR
- L2VPN based on BGP (VPWS / draft Kompella)
- VPLS
- Extended-vlan-ccc
- Pseudo-wire protection using Ethernet OAM
- Local switching of pseudo-wire
- Pseudowire fault detection based on VCCV
- QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, and QFX5100 switches.

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138

Understanding MPLS Components

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

IP Over MPLS for Customer Edge Interfaces

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See "Configuring MPLS on Provider Edge Switches" on page 179.

BGP Layer 3 VPN Configuration

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

Routing Instances for Layer 3 VPN

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- Interior Gateway Protocol on page 146
- MPLS Protocol on page 146
- RSVP on page 146
- Family mpls on page 147

Interior Gateway Protocol

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.

NOTE: The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to **family mpls**, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

RSVP

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

Family mpls

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.

NOTE: **You can enable family mpls on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.**

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 160
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187
- Configuring a Global MPLS EXP Classifier on page 175
- Configuring Ethernet over MPLS (L2 Circuit) on page 170

- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

Understanding MPLS Label Operations

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network header is analyzed and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. The packet is then forwarded to the next provider switch in the LSP. This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is forwarded based on its IP routing information.

This topic describes:

MPLS Label-Switched Paths and MPLS Labels

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the MPLS label (32 bits). Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

Figure 3 on page 149 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 3: Label Encoding

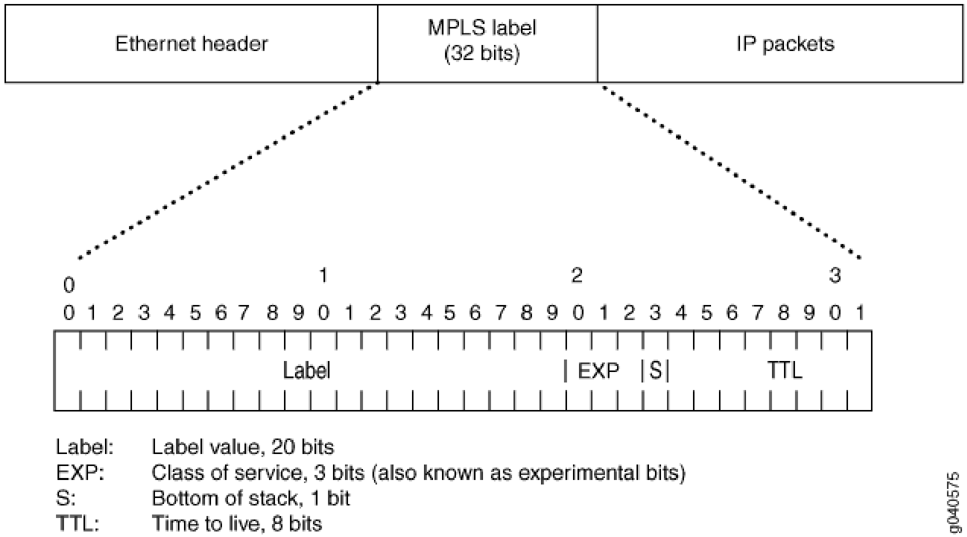

Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the range 0 through 15) have well-defined meanings. The following reserved labels are used by QFX Series and EX4600 devices:

- 0, IPv4 Explicit Null label—This value is valid only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 3, Implicit Null label—This label is used in the signaling protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

MPLS Label Operations

QFX Series and EX4600 devices support the following MPLS label operations:

- Push
- Pop
- Swap

NOTE: There is a limit with regard to the number of labels that QFX and EX4600 devices can affix (push operations) to the label stack or remove (pop operations) from the label stack.

- For Push operations—As many as three labels are supported.
- For Pop operations—As many as two labels are supported.

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 4 on page 150 shows an IP packet without a label arriving on the customer edge interface (ge-0/0/1) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination as the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (ge-0/0/5). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface ge-0/0/5 with label 100. The provider switch swaps label 100 with label 200 and forwards the MPLS packet through its core interface (ge-0/0/7) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (ge-0/0/7), removes the MPLS label, and sends the IP packet out of its customer edge interface (ge-0/0/1) to a destination that is beyond the tunnel.

Figure 4: MPLS Label Swapping

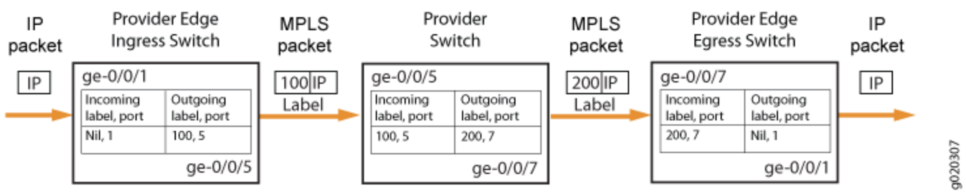

Figure 4 on page 150 shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each PE switch operates as both an ingress switch and an egress switch.

Penultimate-Hop Popping and Ultimate-Hop Popping

The switches enable penultimate-hop popping (PHP) by default with IP over MPLS configurations. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping the MPLS label.

- The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop switch removes the label and sends the packet to the egress PE switch.
- If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised and the egress PE switch of the LSP removes the label.

Related Documentation

- Understanding MPLS Components on page 144
- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183

- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

Understanding BGP

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes MP_REACH_NLRI and MP_UNREACH_NLRI, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- Autonomous Systems on page 152
- AS Paths and Attributes on page 152
- External and Internal BGP on page 153
- Multiple Instances of BGP on page 153

Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

AS Paths and Attributes

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*. Figure 5 on page 153 illustrates ASs, IBGP, and EBGP.

Figure 5: ASs, EBGP, and IBGP

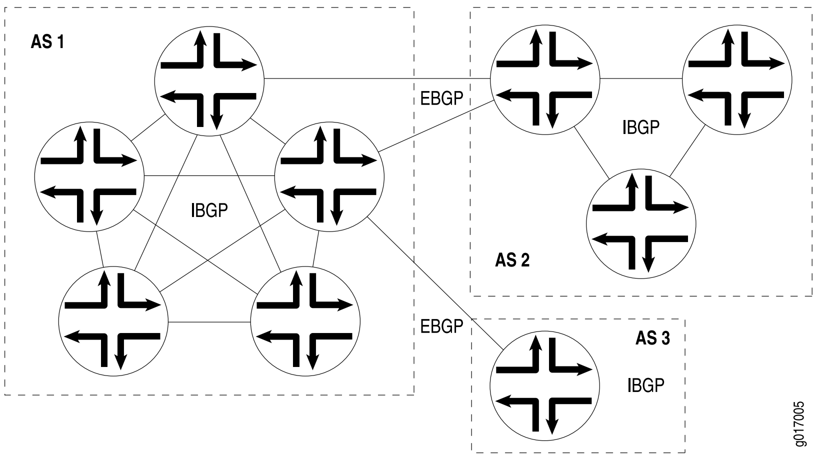

A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.

NOTE: When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

Related Documentation

- *BGP Routes Overview*
- *BGP Messages Overview*

IPv6 Layer 3 VPNs

The interfaces between the PE and CE routers of a Layer 3 VPN can be configured to carry IP version 6 (IPv6) traffic. IP allows numerous nodes on different networks to interoperate seamlessly. IPv4 is currently used in intranets and private networks, as well as the Internet. IPv6 is the successor to IPv4, and is based for the most part on IPv4.

In the Juniper Networks implementation of IPv6, the service provider implements an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers. The PE routers have both IPv4 and IPv6 capabilities. They maintain IPv6 VPN routing and forwarding (VRF) tables for their IPv6 sites and encapsulate IPv6 traffic in MPLS frames that are then sent into the MPLS core network.

IPv6 for Layer 3 VPNs is supported for BGP and for static routes.

IPv6 over Layer 3 VPNs is described in RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

For more information about IPv6, see the *Junos OS Routing Protocols Library for Routing Devices*.

Ethernet Pseudowire Overview

An Ethernet pseudowire is used to carry Ethernet or 802.3 Protocol Data Units (PDUs) over an MPLS network enabling service providers to offer emulated Ethernet services over existing MPLS networks. Ethernet or 802.3 PDUs are encapsulated within the pseudowire to provide a point-to-point Ethernet service. For the point-to-point Ethernet service, the following fault management features are supported:

- The IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM link-fault management on Ethernet point-to-point direct links or links across Ethernet repeaters.

  Ethernet OAM link-fault management can be used for physical link-level fault detection and management. It uses a new, optional sublayer in the data link layer of the OSI model. Ethernet OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. A system-wide implementation is not required; OAM can be deployed on particular interfaces of a router. Transmitted Ethernet OAM messages or OAM PDUs are of standard length, untagged Ethernet frames within the normal frame length limits in the range 64–1518 bytes.

- Ethernet connectivity fault management (CFM) to monitor the physical link between two routers.

  - Connection protection using the continuity check protocol for fault monitoring . The continuity check protocol is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.

  - Path protection using the linktrace protocol for path discovery and fault verification . Similar to IP traceroute, the linktrace protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.

Related Documentation

- *Configuring 802.3ah OAM Link-Fault Management*
- *Pseudowire Overview for ACX Series Universal Access Routers*
- *TDM Pseudowires Overview*
- *ATM Pseudowire Overview*

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. (For information about DSCP and IEEE 802.1p classifiers and general information about classifiers, see *Understanding CoS Classifiers*. For information about DSCP and IEEE 802.1p rewrite rules, see *Understanding CoS Rewrite Rules*.)

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

NOTE: There is no default EXP classifier. There is no default EXP rewrite rule. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, **set interfaces xe-0/0/35 unit 0 family mpls**.)

This topic includes:

EXP Classifiers

Unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to bits 000 through 111).

You can configure as many EXP classifiers as you want. However, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All switch interfaces use the global EXP classifier to classify MPLS traffic.

Only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.

NOTE: The switch uses only the outermost label of incoming EXP packets for classification.

NOTE: MPLS packets with 802.1Q tags are not supported.

EXP Rewrite Rules

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure as many EXP rewrite rules as you want, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

NOTE: **On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.**

Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class. See *Understanding CoS Output Queue Schedulers* and *Example: Configuring Queue Schedulers* for more information.

Related Documentation

- *Understanding CoS Classifiers*
- *Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces*
- Configuring a Global MPLS EXP Classifier on page 175
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187
- Configuring CoS Bits for an MPLS Network on page 170

Understanding Ethernet-over-MPLS (L2 Circuit)

Ethernet-over-MPLS allows sending Layer 2 (L2) Ethernet frames transparently over MPLS Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network This technology has applications in service provider, enterprise and data center environments. For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network.

NOTE: A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a dedicated LSP

• Ethernet-over-MPLS in Data Centers on page 159

Ethernet-over-MPLS in Data Centers

For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network. These data centers require L2 connectivity between them for the following reasons:

- To replicate the storage over Fiber Channel IP (FCIP). FCIP works only on the same broadcast domain.
- To run a dynamic routing protocol between the sites.
- To support High Availability clusters that interconnect the nodes hosted in the various data centers.

Related Documentation

• Configuring Ethernet over MPLS (L2 Circuit) on page 170

Understanding Using MPLS-Based Layer 3 VPNs on Switches

On the QFX Series and on EX4600, you can use MPLS-based Layer 3 virtual private networks (VPNs) to securely connect geographically diverse sites across an MPLS network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions.

A VPN uses a public telecommunications infrastructure, such as the Internet, to provide remote offices or individual users with secure access to their organization's network. VPNs are designed to provide the same level of performance and security as privately owned or leased networks but without the attendant costs.

This topic describes:

MPLS-Based Layer 3 VPNs

In Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks*. RFC 4364 defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

Customer networks, because they are private, can use either public or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. BGP/MPLS VPNs solve this problem by adding a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and on the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. Two different VPNs can use overlapping addresses. Each route within a VPN is assigned an MPLS label (for example, MPLS-ARCH, MPLS-BGP, or MPLS-ENCAPS). When BGP distributes a VPN route, it also distributes an MPLS label for that route. Before a customer data packet travels across the service provider's backbone, it is encapsulated along with the MPLS label that corresponds to the route within the customer's VPN that is the best match based on the packet's destination address. This MPLS packet is further encapsulated with another MPLS label or with an IP, so that it gets tunneled across the backbone to the egress provider edge (PE) switch. Thus, the backbone core switches do not need to know the VPN routes.

QFX5100 switches also support interprovider VPNs, and carrier-of-carriers VPNs. For more information, see "Interprovider and Carrier-of-Carriers VPNs" on page 166

Related Documentation

- Understanding MPLS Label Operations on page 148
- Understanding MPLS Components on page 144

- *Example: Configuring MPLS-Based Layer 2 VPNs*
- *Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches*

MPLS Stitching For Virtual Machine Connection

By using MPLS, the stitching feature of Junos OS provides connectivity between virtual machines that reside either on opposite sides of data center routers or in different data centers. An external controller, programmed in the data-plane, assigns MPLS labels to both virtual machines and servers. Then, the signaled MPLS labels are used between the data center routers, generating static link switched paths (LSPs), resolved over either BGP labeled unicast, RSVP or LDP, to provide the routes dictated by the labels.

When Would I Use Stitching?

There are several ways to connect virtual machines. One option when you have virtual machines on opposite sides of a router (or different data centers) is to use MPLS stitching. A typical topology for using MPLS stitching is shown in Figure 6 on page 161.

Figure 6: Virtual Machines on Either Side of Routers

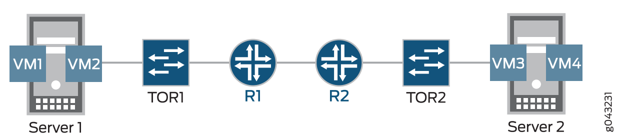

The above topology consists of the following MPLS layers: VMs | Servers | ToRs | Router ...... Router | ToRs | Servers | VMs

NOTE: The label on the left is the top of the label stack.

How Does MPLS Stitching Work?

With stitching, the MPLS static allocation of labels demultiplexes incoming traffic onto any device/entity in the next layer in the direction of traffic flow. Essentially, there is a label hierarchy that picks up labels for the correct top-of-rack switch, server, and virtual machine that receives traffic. Static label assignments are done between the top-of-rack switches and the virtual machines.

For example, imagine that traffic is sent from VM1 to VM3 in Figure 6 on page 161. When traffic exits Server1, its label stack is L1 | L2 | L3 where:

- L1 represents the egress top-of-rack switch ToR1.
- L2 represents the physical server, Server2, towards which the egress-side ToR will forward the traffic.
- L3 represents the virtual machine on Server2 to which Server2 should deliver the traffic.

Traffic arriving at ToR1 needs to be sent to ToR2. Since ToR1 and ToR2 are not directly connected, traffic must flow from ToR1 to ToR2 using label-switching starting on the outermost (top) label. Stitching has been added to static-LSP functionality to SWAP L1 to a l-BGP label that ToR2 advertises to ToR1. The label stack now must contain another label at the top to enable forwarding of the labeled packets between ToR1 and ToR2. An L-Top label is added if L-BGP is resolved over RSVP/LDP. If static LSP is resolved over L-BGP, then the top label is swapped with the L-BGP label and there is no L-Top label. When the traffic exits ToR1, the stack is: L-top | L-BGP | L2 | L3.

Traffic from ToR1 to ToR2 is then label switched over any signaled LSP.

When traffic arrives at ToR2, the top label is removed with PHP (popped) and the label stack becomes L-BGP | L2 | L3. Since L-BGP is a implicit null label, ToR2 pops the static LSP label L2 that corresponds to the egress server and then forwards the packet to the egress server using the static-LSP configuration on ToR2, which corresponds to a single-hop implicit-NULL LSP.

The outgoing stack becomes L3 and the next-hop is the egress server Server2.

When traffic arrives at the egress server Server2, Server2 pops L3 and delivers the packet to VM3.

How Do I Configure Stitching?

The new keyword **stitch** has been added under **transit** to resolve the remote next-hop. For example, instead of **set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 next-hop 10.9.82.47**, a top-of-rack switch redirects packets to another top-of-rack switch with **set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 stitch**. The **show mpls static-lsp** command has been extended to show the LSP state as 'InProgress' whenever the LSP is waiting for protocol next-hop resolution by resolver.

See the complete example for stitching at “Using MPLS Stitching with BGP to Connect Virtual Machines” on page 219 for more information.

Which Switches Support Stitching?

QFX5100, QFX3500 and EX4600 support the static LSP stitching feature.

Q&A

Q: Is link and node protection for the next-hop provided by MPLS stitching?
A: Link and node protection for the next-hop of transit LSP stitched to L-BGP LSP are not needed. That is provided by L-BGP LSP.

Q: Does stitching work with real L-BGP labels?
A: No, stitching does not work when the L-BGP label is a real label.

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Using MPLS Stitching with BGP to Connect Virtual Machines on page 219

Carrier-of-Carriers VPNs

The customer of a VPN service provider might be a service provider for the end customer. The following are the two main types of carrier-of-carriers VPNs (as described in RFC 4364:

- "Internet Service Provider as the Customer" on page 165—The VPN customer is an ISP that uses the VPN service provider's network to connect its geographically disparate regional networks. The customer does not have to configure MPLS within its regional networks.
- "VPN Service Provider as the Customer" on page 165—The VPN customer is itself a VPN service provider offering VPN service to its customers. The carrier-of-carriers VPN service customer relies on the backbone VPN service provider for inter-site connectivity. The customer VPN service provider is required to run MPLS within its regional networks.

Figure 7 on page 164 illustrates the network architecture used for a carrier-of-carriers VPN service.

Figure 7: Carrier-of-Carriers VPN Architecture

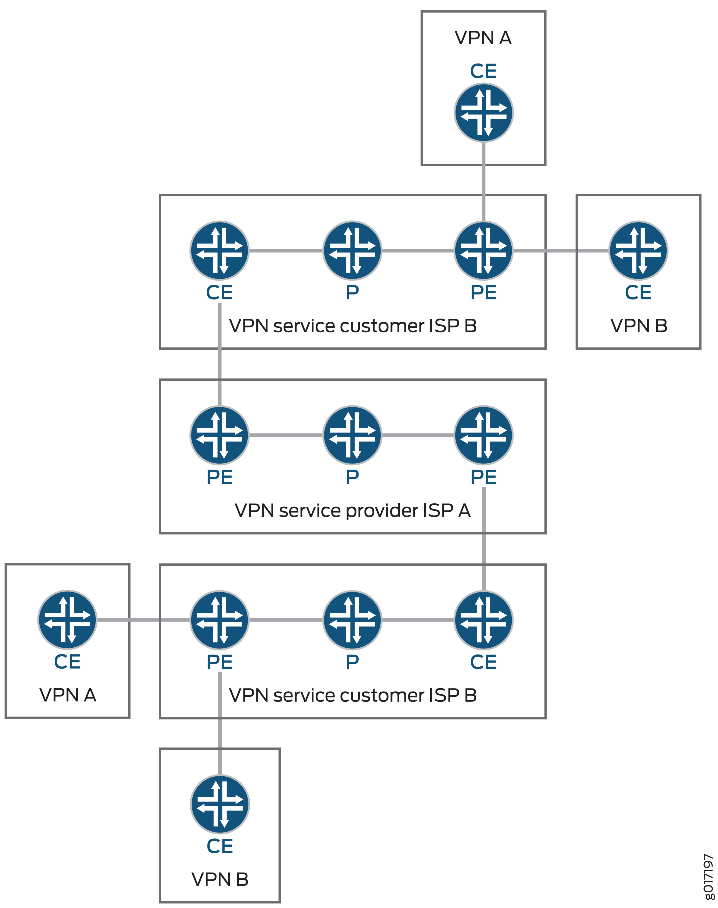

This topic covers the following:

- Internet Service Provider as the Customer on page 165
- VPN Service Provider as the Customer on page 165

Internet Service Provider as the Customer

In this type of carrier-of-carriers VPN configuration, ISP A configures its network to provide Internet service to ISP B. ISP B provides the connection to the customer wanting Internet service, but the actual Internet service is provided by ISP A.

This type of carrier-of-carriers VPN configuration has the following characteristics:

- The carrier-of-carriers VPN service customer (ISP B) does not need to configure MPLS on its network.
- The carrier-of-carriers VPN service provider (ISP A) must configure MPLS on its network.
- MPLS must also be configured on the CE routers and PE routers connected together in the carrier-of-carriers VPN service customer's and carrier-of-carriers VPN service provider's networks.

VPN Service Provider as the Customer

A VPN service provider can have customers that are themselves VPN service providers. In this type of configuration, also called a hierarchical or recursive VPN, the customer VPN service provider's VPN-IPv4 routes are considered external routes, and the backbone VPN service provider does not import them into its VRF table. The backbone VPN service provider imports only the customer VPN service provider's internal routes into its VRF table.

The similarities and differences between interprovider and carrier-of-carriers VPNs are shown in Table 16 on page 165.

Table 16: Comparison of Interprovider and Carrier-of-Carriers VPNs

Feature	ISP Customer	VPN Service Provider Customer
Customer edge device	AS border router	PE router
IBGP sessions	Carry IPv4 routes	Carry external VPN-IPv4 routes with associated labels
Forwarding within the customer network	MPLS is optional	MPLS is required

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Interprovider and Carrier-of-Carriers VPNs on page 166
- *Interprovider VPNs*

Interprovider and Carrier-of-Carriers VPNs

All interprovider and carrier-of-carriers VPNs share the following characteristics:

- Each interprovider or carrier-of-carriers VPN customer must distinguish between internal and external customer routes.
- Internal customer routes must be maintained by the VPN service provider in its PE routers.
- External customer routes are carried only by the customer's routing platforms, not by the VPN service provider's routing platforms.

The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same AS or to separate ASs:

- *Interprovider VPNs*—The customer sites belong to different ASs. You need to configure EBGP to exchange the customer's external routes.
- "Carrier-of-Carriers VPNs" on page 164—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.

In general, each service provider in a VPN hierarchy is required to maintain its own internal routes in its P routers, and the internal routes of its customers in its PE routers. By recursively applying this rule, it is possible to create a hierarchy of VPNs.

The following are definitions of the types of PE routers specific to interprovider and carrier-of-carriers VPNs:

- The AS border router is located at the AS border and handles traffic leaving and entering the AS.
- The end PE router is the PE router in the customer VPN; it is connected to the CE router at the end customer's site.

Related Documentation

- Carrier-of-Carriers VPNs on page 164
- *Interprovider VPNs*
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138

Fast Reroute Overview

Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 8 on page 167 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers (or switches) that are not shown in the figure.

Fast reroute protects traffic against any single point of failure between the ingress and egress routers (or switches). If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

Figure 8: Detours Established for an LSP Using Fast Reroute

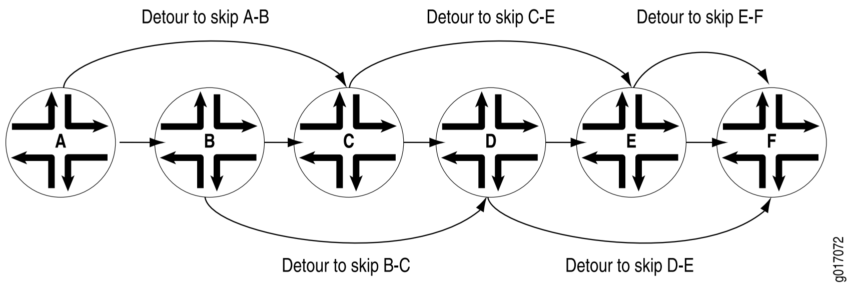

If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 9 on page 167 illustrates the detour taken when the link between Router B and Router C fails.

Figure 9: Detour After the Link from Router B to Router C Fails

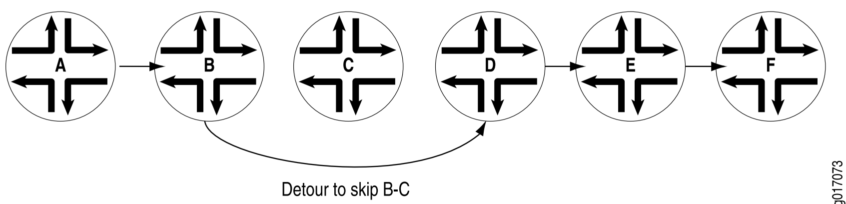

If the network topology is not rich enough (there are not enough routers with sufficient links to other routers), some of the detours might not succeed. For example, the detour from Router A to Router C in Figure 8 on page 167 cannot traverse link A-B and Router B. If such a path is not possible, the detour does not occur.

Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.

NOTE: If you issue show commands after the node has switched traffic to the initial detour, the node might indicate that the traffic is still flowing over the original LSP. This situation is temporary and should correct itself quickly.

The time required for a fast-rerouting detour to take effect depends on two independent time intervals:

- Amount of time to detect that there is a link or node failure—This interval depends greatly on the link layer in use and the nature of the failure. For example, failure detection on an SONET/SDH link typically is much faster than on a Gigabit Ethernet link, and both are much faster than detection of a router failure.
- Amount of time required to splice the traffic onto the detour—This operation is performed by the Packet Forwarding Engine, which requires little time to splice traffic onto the detour. The time needed can vary depending on the number of LSPs being switched to detours.

Fast reroute is a short-term patch to reduce packet loss. Because detour computation might not reserve adequate bandwidth, the detours might introduce congestion on the alternate links. The ingress router is the only router that is fully aware of LSP policy constraints and, therefore, is the only router able to come up with adequate long-term alternate paths.

Detours are created by use of RSVP and, like all RSVP sessions, they require extra state and overhead in the network. For this reason, each node establishes at most one detour for each LSP that has fast reroute enabled. Creating more than one detour for each LSP increases the overhead, but serves no practical purpose.

To reduce network overhead further, each detour attempts to merge back into the LSP as soon as possible after the failed node or link. If you can consider an LSP that travels through *n* router nodes, it is possible to create *n* – 1 detours. For instance, in Figure 10 on page 168, the detour tries to merge back into the LSP at Router D instead of at Router E or Router F. Merging back into the LSP makes the detour scalability problem more manageable. If topology limitations prevent the detour from quickly merging back into the LSP, detours merge with other detours automatically.

Figure 10: Detours Merging into Other Detours

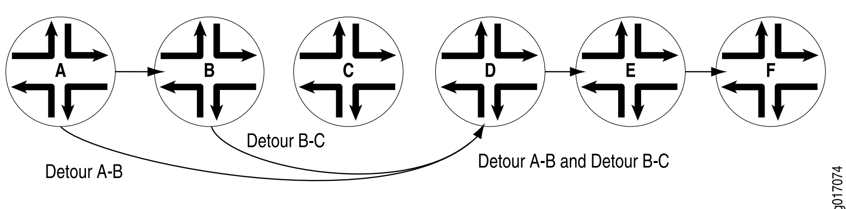

Related Documentation

- fast-reroute on page 265
- *Configuring Fast Reroute*
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Interprovider and Carrier-of-Carriers VPNs on page 166

Types of LSPs

There are three types of LSPs:

- Static LSPs—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- LDP-signaled LSPs—See “LDP Introduction” on page 4.
- RSVP-signaled LSPs—For signaled paths, RSVP is used to set up the path and dynamically assign labels. (RSVP signaling messages are used to set up signaled paths.) You configure only the ingress router. The transit and egress routers accept signaling information from the ingress router, and they set up and maintain the LSP cooperatively. Any errors encountered while establishing an LSP are reported to the ingress router for diagnostics. For signaled LSPs to work, a version of RSVP that supports tunnel extensions must be enabled on all routers.

There are two types of RSVP-signaled LSPs:

- Explicit-path LSPs—All intermediate hops of the LSP are manually configured. The intermediate hops can be strict, loose, or any combination of the two. Explicit path LSPs provide you with complete control over how the path is set up. They are similar to static LSPs but require much less configuration.
- Constrained-path LSPs—The intermediate hops of the LSP are automatically computed by the software. The computation takes into account information provided by the topology information from the IS-IS or OSPF link-state routing protocol, the current network resource utilization determined by RSVP, and the resource requirements and constraints of the LSP. For signaled constrained-path LSPs to work, either the IS-IS or OSPF protocol and the IS-IS or OSPF traffic engineering extensions must be enabled on all routers.

Configuring CoS Bits for an MPLS Network

When traffic enters a labeled-switch path (LSP) tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service Feature Guide for Routing Devices* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP:

1. Specify a class of service value for the LSP:

   NOTE: **The CoS value set using the class-of-service statement at the [edit protocols mpls] hierarchy level supersedes the CoS value set at the [edit class-of-service] hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.**

   ```
   [edit protocols mpls]
   user@switch# set class-of-service cos-value
   ```

Related Documentation

- *Understanding CoS Classifiers*
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- Configuring a Global MPLS EXP Classifier on page 175
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187
- *Defining CoS Rewrite Rules*

Configuring Ethernet over MPLS (L2 Circuit)

To implement Ethernet over MPLS, you must configure a Layer 2 circuit on the provider edge (PE) switches. No special configuration is required on the customer edge (CE) switches. The provider switches require MPLS and LDP to be configured on the interfaces that will be receiving and transmitting MPLS packets.

NOTE: **A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two PE switches. In contrast, each CCC requires a dedicated LSP.**

This topic describes how to configure the PE switches to support Ethernet over MPLS. You must configure interfaces and protocols on both the local PE (PE1) and the remote PE (PE2) switches. The interface configuration varies depending upon whether the Layer 2 circuit is port-based or VLAN-based.

Figure 11 on page 171 shows an example of a Layer 2 circuit configuration.

Figure 11: Ethernet over MPLS Layer 2 Circuit

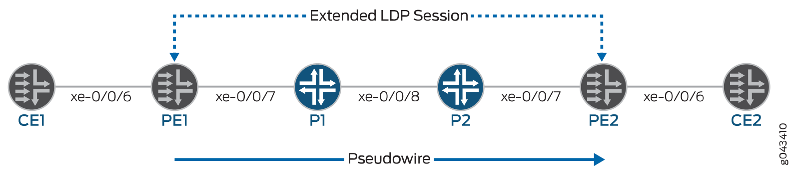

NOTE: This topic refers to the local PE switch as PE1 and the remote PE switch as PE2. It also uses interface names rather than variables to help clarify the connections between the switches. The loopback addresses of the switches are configured as follows:

- PE1: 1.1.1.1
- PE2: 4.4.4.4

Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

CAUTION: Configure MPLS networks with an MTU (maximum transmission unit) that is at least 12 bytes larger than the largest frame size that will be transported by the LSPs. If the size of a encapsulated packet on the ingress LSR exceeds the LSP MTU, that packet is dropped. If an egress LSR receives a packet on a VC LSP with a length (after the label stack and sequencing control word have been popped) that exceeds the MTU of the destination layer 2 interface, that packet is also dropped.

To configure the local PE switch (PE1) for a port-based layer 2 circuit (pseudo-wire):

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```

NOTE: On QFX Series switches, the L2 circuit CE facing interface does not support Aggregated Ethernet (AE) interfaces.

2. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch#set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch#set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

To configure the remote PE switch (PE2) for a port-based layer 2 circuit:

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```

NOTE: On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch#set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch#set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```


Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit

To configure the local PE switch (PE1) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```

NOTE: On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch#set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch#set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit

To configure the remote PE switch (PE2) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```

NOTE: On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch#set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch#set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

Related Documentation

- Understanding Ethernet-over-MPLS (L2 Circuit) on page 159

Configuring a Global MPLS EXP Classifier

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure as many EXP classifiers as you want, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

If an EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

NOTE: There is no default MPLS EXP classifier. If you want to use an MPLS EXP classifier, you must configure it. The MPLS EXP classifier is global and applies only to all **family mpls** interfaces on the switch. You can configure as many MPLS EXP classifiers as you want, but you can only use one MPLS EXP classifier on switch interfaces at any time.

To configure a unicast MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

   `[edit class-of-service classifiers]`
   `user@switch#` **set (dscp | ieee-802.1 | exp)** ***classifier-name*** **forwarding-class** ***forwarding-class-name*** **loss-priority** ***level*** **code-points [*aliases*] [*bit-patterns*]**
2. Apply the EXP classifier to the switch interfaces:

   `[edit class-of-service]`
   `user@switch#` **set system-defaults classifiers exp** ***classifier-name***

Related Documentation

- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- *Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces*
- *Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)*
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187

Configuring MPLS Firewall Filters and Policers

You can configure firewall filters to filter MPLS traffic. To use an MPLS firewall filter, you must first configure the filter and then apply it to an interface you have configured for forwarding MPLS traffic. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached.

NOTE: You can configure ingress MPLS firewall filters only. Egress MPLS firewall filters are not supported. You cannot apply MPLS firewall filters to loopback interfaces.

When you configure an MPLS firewall filter, you define filtering criteria (terms, with match conditions) for the packets and an action (action, or action modifier) for the switch to take if the packets match the filtering criteria.

- Table 17 on page 176 describes the match conditions you can configure for MPLS firewall filters at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level.

  NOTE: If a packet has multiple MPLS labels, the filter applies the match conditions to only the bottom label in the label stack.

Table 17: Supported Match Conditions for MPLS Firewall Filters

Match Condition	Description
exp ***number***	Experimental (EXP) bit number or range of bit numbers in the MPLS header of a packet. For *number*, you can specify one or more values from 0 through 7 in binary, decimal or hexadecimal format, as given below: • A single EXP bit—for example, **exp 3** • Several EXP bits—for example, **exp 0,4** • A range of EXP bits—for example, **exp [0-5]**
label ***number***	MPLS label value or range of label values in the MPLS header of a packet. For *number*, you can specify one or more values from 0 through 1048575 in decimal or hexadecimal format, as given below: • A single label—for example, **label 3** • Several labels—for example, **label 0,4** • A range of labels—for example, **label [0-5]**

- Table 18 on page 176 describes the actions you can configure for MPLS firewall filters at the **[edit firewall family mpls filter *filter-name* term *term-name* then]** hierarchy level.

Table 18: Supported Actions for MPLS Firewall Filters

Action	Description
accept	Accept a packet

Table 18: Supported Actions for MPLS Firewall Filters *(continued)*

Action	Description
count ***counter-name***	Count the number of packets that pass this filter or term. NOTE: We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.
discard	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message
policer	Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer.
three-color-policer	Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a three-color policer.

- Configuring an MPLS Firewall Filter on page 177
- Applying an MPLS Firewall Filter to an MPLS Interface on page 177
- Configuring Policers for LSPs on page 178

Configuring an MPLS Firewall Filter

To configure an MPLS firewall filter:

1. Configure the filter name, term name, and at least one match condition—for example, match on MPLS packets with EXP bits set to either 0 or 4:

   ```
   [edit firewall family mpls]
   user@switch# set filter ingress-exp-filter term term-one from exp 0,4
   ```

2. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term—for example, count MPLS packets with EXP bits set to either 0 or 4:

   ```
   [edit firewall family mpls filter ingress-exp-filter term term-one then]
   user@switch# set count counter0
   user@switch# set accept
   ```

Applying an MPLS Firewall Filter to an MPLS Interface

To apply the MPLS firewall filter to an interface you have configured for forwarding MPLS traffic (using the **family mpls** statement at the **[edit interfaces** ***interface-name*** **unit** ***unit-number*****]** hierarchy level):

NOTE: You can apply firewall filters only to filter MPLS packets that enter an interface.

1. Apply the firewall filter to an MPLS interface—for example, apply the firewall filter to interface xe-0/0/5:


```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls filter input ingress-exp-filter
```

2. Review your configuration and issue the **commit** command:

```
[edit interfaces]
user@switch# commit
commit complete
```

Configuring Policers for LSPs

Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer or three-color policer. MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Supported MPLS Scaling Values on page 218

- *Overview of Policers*

Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable or disable MPLS statistics collection, include the **statistics** statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
}
```

You can configure these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The default interval is 300 seconds.

If you configure the **file** option, the statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP. Feature parity for the display of packet and byte statistics for sub-LSPs of a point-to-multipoint LSP on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. Sample output follows:

```
lsp6                    0 pkt            0 Byte      0 pps        0 Bps     0
lsp5                    0 pkt            0 Byte      0 pps        0 Bps     0
lsp6.1              34845 pkt      2926980 Byte   1049 pps    88179 Bps   132
lsp5.1                  0 pkt            0 Byte      0 pps        0 Bps     0
lsp4                    0 pkt            0 Byte      0 pps        0 Bps     0
Dec  7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored
```

Related Documentation

- *Configuring Automatic Bandwidth Allocation for LSPs*

Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. Configuring the Ingress PE Switch on page 180
2. Configuring the Egress PE Switch on page 181

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```

NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:

NOTE: You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
user@switch# set traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure MPLS traffic engineering.

```
[edit protocols mpls]
user@switch# set traffic-engineering
```

6. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16
```

9. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0
```

10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1
```

11. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf
```

12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve
```

Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24
```

NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:

NOTE: You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0
```

```
user@switch# set rsvp interface xe-0/0/6.0
```

4. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch#  set static route 121.121.121.0/24 next-hop 192.168.12.1
user@switch# set static route 121.121.121.0/24 resolve
```

Related Documentation

- MPLS Configuration Guidelines on page 217
- Configuring MPLS on Provider Switches on page 183
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding MPLS Components on page 144
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156

Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:

   NOTE: You can use the switch address as an alternative to the loopback interface.

   ```
   [edit protocols ospf]
   user@switch# set area 0.0.0.0 interface lo0.0
   user@switch# set area 0.0.0.0 interface xe-0/0/5.0
   user@switch# set area 0.0.0.0 interface xe-0/0/6.0
   user@switch# set area 0.0.0.0 interface ae0
   ```

   NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

   ```
   [edit protocols mpls]
   user@switch# set interface xe-0/0/5.0
   user@switch# set interface xe-0/0/6.0
   user@switch# set interface ae0
   ```
3. Configure RSVP on the loopback interface and the core interfaces:

   ```
   [edit protocols rsvp]
   user@switch# set interface lo0.0
   user@switch# set interface xe-0/0/5.0
   user@switch# set interface xe-0/0/6.0
   user@switch# set interface ae0
   ```
4. Configure an IP address for the loopback interface and the core interfaces:

   ```
   [edit interfaces]
   user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
   user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
   user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
   user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
   ```
5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

   ```
   [edit interfaces]
   user@switch# set xe-0/0/5 unit 0 family mpls
   user@switch# set xe-0/0/6 unit 0 family mpls
   user@switch# set ae0 unit 0 family mpls
   ```

NOTE: You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- MPLS Configuration Guidelines on page 217
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding MPLS Components on page 144
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156

Configuring Static Label Switched Paths for MPLS

Configuring static label-switched paths (LSPs) for MPLS is similar to configuring static routes on individual switches. As with static routes, there is no error reporting, liveliness detection, or statistics reporting.

To configure static LSPs, configure the ingress PE switch and each provider switch along the path up to and including the egress PE switch.

For the ingress PE switch, configure which packets to tag (based on the packet's destination IP address), configure the next switch in the LSP, and the tag to apply to the packet. Manually assigned labels can have values from 0 through 1,048,575.

For the transit switches in the path, configure the next switch in the path and the tag to apply to the packet. Manually assigned labels can have values from 1,000,000 through 1,048,575.

The egress PE switch removes the label and forwards the packet to the IP destination. However, if the previous switch removed the label, the egress switch examines the packet's IP header and forwards the packet toward its IP destination.

Before you configure a static LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See "Configuring MPLS on Provider Edge Switches" on page 179.

NOTE: Do not configure LSPs at the **[edit protocols mpls label-switched-path]** hierarchy level on the PE switches.

- Configure one or more provider switches. See “Configuring MPLS on Provider Switches” on page 183.

This topic describes how to configure an ingress PE switch, one or more provider switches, and an egress PE switch for static LSP:

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for every core interface:

   `[edit interfaces]`
   `user@switch#` **set *interface-name* unit *logical-unit-number* family inet address *address***

   NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure the name associated with the static LSP:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name***
3. Configure the next hop switch for the LSP:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* ingress next-hop *address-of-next-hop***
4. Specify the address of the egress switch for the LSP:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* ingress to *address-of-egress-switch***
5. Configure the new label that you want to add to the top of the label stack:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* ingress push *out-label***

Configuring the Provider and the Egress PE Switch

To configure a static LSP for MPLS on the provider and egress PE switch:

1. Configure a transit static LSP:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* transit *incoming-label***
2. Configure the next hop switch for the LSP:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* transit *incoming-label* next-hop *address-of-next-hop***
3. Only for provider switches, remove the label at the top of the label stack and replace it with the specified label:

   `[edit protocols mpls]`
   `user@switch#` **set static-label-switched-path *lsp-name* transit *incoming-label* swap *out-label***
4. Only for the egress PE switch, remove the label at the top of the label stack:

NOTE: If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name transit incoming-label pop
```

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183
- Understanding MPLS Label Operations on page 148

Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.

NOTE: There are no default rewrite rules.

You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

NOTE: To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

   `user@switch#` **set class-of-service rewrite-rules exp *rewrite-rule-name* forwarding-class *forwarding-class-name* loss-priority *level* code-points [*aliases*] [*bit-patterns*]**

   For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

   `user@switch#` **set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1 loss-priority low code-points 001**

2. Apply the rewrite rule to a logical interface:

   `user@switch #` **set class-of-service interfaces *interface-name* unit *logical-unit* rewrite-rules exp *rewrite-rule-name***

For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```

NOTE: In this example, all forwarding classes assigned to port xe-0/0/10 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

Related Documentation

- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- *Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces*
- *Monitoring CoS Rewrite Rules*
- *Defining CoS Rewrite Rules*
- Configuring a Global MPLS EXP Classifier on page 175

Example: Configuring MPLS-Based Layer 3 VPNs

You can implement an MPLS-based Layer 3 virtual private network (VPN) on QFX switches to interconnect sites for customers who want the service provider to handle all the Layer 3 routing functions. To support an MPLS-based Layer 3 VPN, you need to add components of the Layer 3 VPN to the configuration of the two provider edge (PE) switches. You do not need to change the configuration of the provider switches.

This example shows how to configure an MPLS-based Layer 3 VPN spanning two corporate sites:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.3 or later for the QFX Series
- Three QFX switches

Before you configure the Layer 3 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See “Configuring MPLS on Provider Edge Switches” on page 179.
- Configure one or more provider switches. See “Configuring MPLS on Provider Switches” on page 183.

Overview and Topology

Layer 3 VPNs allow customers to leverage the service provider’s technical expertise to ensure efficient site-to-site routing. The customer’s customer edge (CE) switch uses a routing protocol such as BGP or OSPF to communicate with the service provider’s provider edge (PE) switch to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use only IP over MPLS; other protocol packets are not supported. This example includes two PE switches, PE1 and PE2.

In the basic MPLS configuration of the PE switches using IP over MPLS, the PE switches were configured to use OSPF as the routing protocol between the MPLS switches and RSVP as the signaling protocol. Traffic engineering was enabled. A label-switched path (LSP) was configured.

The following components must be added to the PE switches for an MPLS-based Layer 3 VPN:

- BGP group with **family inet-vpn unicast**
- Routing instance with instance type **vrf**

Figure 12 on page 189 illustrates the topology of this MPLS-based Layer 3 VPN.

Figure 12: MPLS-Based Layer 3 VPN

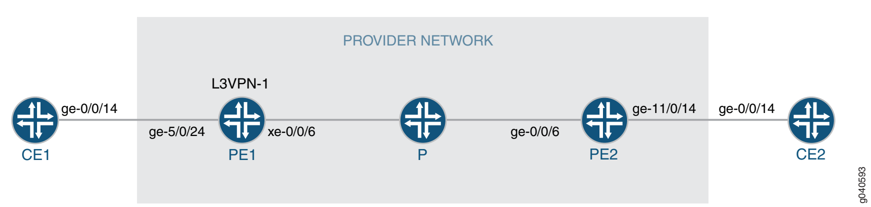

Table 19 on page 190 shows the settings of the customer edge interface on the local CE switch.

Table 19: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Local CE switch hardware	QFX switch	CE1
Customer edge interface	**ge-0/0/14 unit 0** **family inet** **address 51.51.0.14/16**	Interface that connects CE1 to PE1.

Table 20 on page 190 shows the settings of the customer edge interface on the remote CE switch.

Table 20: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Remote CE switch hardware	QFX switch	CE2
Customer edge interface	**ge-0/0/14 unit 0** **family inet** **address 11.22.26.1/16**	Interface that connects CE2 to PE2.

Table 21 on page 190 shows the Layer 3 VPN components of the local PE switch.

Table 21: Layer 3 VPN Components of the Local PE Switch

Property	Settings	Description
Local PE switch hardware	QFX switch	PE1
Customer edge interface	**ge-5/0/24 unit 0** **family inet** **address 51.51.0.1/16**	Connects PE1 to CE1. NOTE: The **family inet** configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration.
Core interface	**xe-0/0/6 unit 0** **family inet address 60.0.0.60/16** **family mpls**	Connects PE1 to P. NOTE: This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.

Table 21: Layer 3 VPN Components of the Local PE Switch *(continued)*

Property	Settings	Description
Loopback interface	**lo0 unit 0** **family inet address 21.21.21.21/32**	NOTE: This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.
BGP	**bgp**	Added for the Layer 3 VPN configuration.
Routing instance	**L3VPN-1**	Added for the Layer 3 VPN configuration.

Table 22 on page 191 shows the Layer 3 VPN components of the remote PE switch.

Table 22: Layer 3 VPN Components of the Remote PE Switch

Property	Settings	Description
Remote PE switch hardware	QFX switch	PE2
Customer edge interface	**ge-0/0/14 unit 0** **family inet** **address 11.22.26.14/16**	Connects PE2 to CE2. For the Layer 3 VPN configuration, added **family mpls**. NOTE: The **family inet** configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration.
Core interface	**xe-0/0/6 unit 0** **family inet address 60.2.0.60/16** **family mpls**	Connects PE1 to P. NOTE: This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.
Loopback interface	**lo0 unit 0** **family inet address 22.22.22.22/32**	NOTE: This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.
BGP	**bgp**	Added for the Layer 3 VPN configuration.
Routing instances	**L3VPN-1**	Added for the Layer 3 VPN configuration.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure the Layer 3 VPN components on the local PE switch, copy the following commands and paste them into the switch terminal window of PE1:

```
[edit]
set protocols bgp local-address 21.21.21.21 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 22.22.22.22
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-0/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label
set routing-options router-id 21.21.21.21
set routing-options autonomous-system 10
```

Step-by-Step Procedure

To configure the Layer 3 VPN components on the local PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:

   ```
   [edit protocols bgp]
   user@switchPE1# set local-address 21.21.21.21 family inet-vpn unicast
   ```
2. Configure the BGP group, specifying the group name and type:

   ```
   [edit protocols bgp]
   user@switchPE1# set group PE1-PE2 type internal
   ```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:

   ```
   [edit protocols bgp]
   user@switchPE1# set neighbor 22.22.22.22
   ```
4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 instance-type vrf
   ```
5. Configure a description for this routing instance:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
   ```
6. Configure the routing instance to use a route distinguisher:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 route-distinguisher 21:21
   ```

   NOTE: Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances require a route distinguisher to allow BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

7. Configure the VPN routing and forwarding (VRF) target of the routing instance:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 vrf-target target:21:21
   ```

   NOTE: You can create more complex policies by explicitly configuring VRF import and export policies using the import and export options. See the *Junos OS VPNs Library for Routing Devices*.

8. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 vrf-table-label
   ```

9. Configure the router ID and autonomous system (AS):

   NOTE: We recommend that you explicitly configure the router identifier under the **[edit routing-options]** hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

   ```
   [edit routing-options]
   user@switchPE1# set router-id 21.21.21.21 autonomous-system 10
   ```

Results Display the results of the configuration:

```
user@switchPE1> show configuration
```

```
interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 51.51.0.1/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      family inet {
        address 60.0.0.60/16;
      }
      family mpls;
    }
  }
```

```
protocols {
   mpls {
      label-switched-path 21-22 {
         from 21.21.21.21;
         to 22.22.22.22;
         no-cspf;
      }
      interface xe-0/0/6.0;
      interface lo0.0;
   bgp {
      local-address 21.21.21.21;
      family inet-vpn {
         unicast;
      }
      group PE1-PE2 {
         type internal;
         neighbor 22.22.22.22;
      }
   }
   ospf
      traffic-engineering;
      area 0.0.0.0 {
         interface lo0.0;
         interface xe-0/0/6.0;
      }
   }
}
routing-instances {
   L3VPN-1 {
      instance-type vrf;
      description "BETWEEN PE1 AND PE2";
      route-distinguisher 21:21;
      vrf-target target:21:21;
      vrf-table-label;
   }
routing-options {
   router-id 21.21.21.21;
   autonomous-system 10;
```

Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure the Layer 3 VPN components on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE2:

```
[edit]
set protocols bgp local-address 22.22.22.22 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 21.21.21.21
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-11/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label;
set routing-options router-id 22.22.22.22
set routing-options autonomous-system 10
```


Step-by-Step Procedure

To configure Layer 3 VPN components on the remote PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:

   ```
   [edit protocols bgp]
   user@switchPE2# set local-address 22.22.22.22 family inet-vpn unicast
   ```
2. Configure the BGP group, specifying the group name and type:

   ```
   [edit protocols bgp]
   user@switchPE2# set group PE1-PE2 type internal
   ```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:

   ```
   [edit protocols bgp]
   user@switchPE2# set neighbor 21.21.21.21
   ```
4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:

   ```
   [edit routing-instances]
   user@switchPE2# set L3VPN-1 instance-type vrf
   ```
5. Configure a description for this routing instance:

   ```
   [edit routing-instances]
   user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
   ```
6. Configure the routing instance to apply to the customer edge interface:

   ```
   [edit routing-instances]
   user@switchPE2# set L3VPN-1 interface ge-0/0/14.0
   ```
7. Configure the routing instance to use a route distinguisher, using the format *ip-address:number*:

   ```
   [edit routing-instances]
   user@switchPE2# set L3VPN-1 route-distinguisher 21:21
   ```
8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

   ```
   [edit routing-instances]
   user@switchPE2# set L3VPN-1 vrf-target target:21:21
   ```
9. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header.

   ```
   [edit routing-instances]
   user@switchPE2# set L3VPN-1 vrf-tabel-label
   ```
10. Configure the router ID and autonomous system (AS):

    ```
    [edit routing-options]
    user@switchPE2# set router-id 22.22.22.22 autonomous-system 10
    ```

Results

Display the results of the configuration:

```
user@switchPE2> show configuration
```

```
interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 11.22.26.14/16;
      }
    }
```

```
}
lo0 {
    unit 0 {
        family inet {
            address 22.22.22.22/32;
        }
    }
}
xe-0/0/6 {
    unit 0 {
        family inet {
            address 60.2.0.60/16;
        }
        family mpls;
    }
}
protocols {
    mpls {
        label-switched-path 22-21 {
            from 22.22.22.22;
            to 21.21.21.21;
            no-cspf;
        }
        interface xe-0/0/6.0;
        interface lo0.0;
    bgp {
        local-address 22.22.22.22;
        family inet-vpn {
            unicast;
        }
        group PE1-PE2 {
            type internal;
            neighbor 21.21.21.21;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ge-0/0/14.0;
            interface lo0.0;
            interface xe-0/0/6.0;
        }
    }
}
routing-instances {
    L3VPN-1 {
        instance-type vrf;
        description"BETWEEN PE1 AND PE2";
        route-distinguisher 21:21;
        vrf-target target:21:21;
        vrf-table-label;
    }
routing-options {
    router-id 22.22.22.22;
    autonomous-system 10;
```

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183

Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network

Starting with Junos OS Release 14.1X53-D40, QFX5100 switches support integrated routing and bridging (IRB) interfaces over an MPLS core network on QFX5100 switches. An IRB is a logical Layer 3 VLAN interface used to route traffic between VLANs.

By definition, VLANs divide a LAN's broadcast environment into isolated virtual broadcast domains, thereby limiting the amount of traffic flowing across the entire LAN and reducing the possible number of collisions and packet retransmissions within the LAN. To forward packets between different VLANs, you normally need a router that connects the VLANs. Now you can accomplish this forwarding without using a router by simply configuring an IRB interface on the switch.

The IRB interface functions as a logical switch on which you can configure a Layer 3 logical interface for each VLAN. The switch relies on its Layer 3 capabilities to provide this basic routing between VLANs. With IRB, you can configure label-switched paths (LSPs) to enable the switch to recognize which packets are being sent to local addresses, so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated.

This example shows how to configure integrated routing and bridging (IRB) interfaces over an MPLS core network on QFX5100 switches.

- Requirements on page 197
- Overview and Topology on page 198
- Configuration on page 198

Requirements

This example uses the following hardware and software components:

- Three QFX5100 switches.
- Junos OS Release 14.1X53-D40 or later

Before you begin, be sure you have:

- An understanding of IRB concepts. See Understanding Integrated Routing and Bridging for an overview of IRB.
- The required ternary content addressable memory (TCAM) space available on the switch. TCAM rules must be observed while configuring and implementing IRBs. For detailed information, see MPLS Limitations on QFX Series and EX4600 Switches.

Overview and Topology

Figure 13 on page 198 illustrates a sample topology for configuring IRB over an MPLS core network. In this example, an LSP is established between the ingress provider edge switch (PE1) and the provider edge egress switch (PE2). An IRB Layer 3 interface (irb.0) is configured between the provider switch (P) and PE2 under vlan 100. In this configuration, the P switch replaces (swaps) the label at the top of the label stack with a new label, adds the vlan identifier 100 to the MPLS packet, and then sends the packet out the IRB interface. PE2 receives this vlan-tagged MPLS packet, removes (pops) the label from the top of the label stack, performs a regular IP route lookup, and then forwards the packet with its IP header to the next-hop address.

Figure 13: IRB Topology over an MPLS Core Network

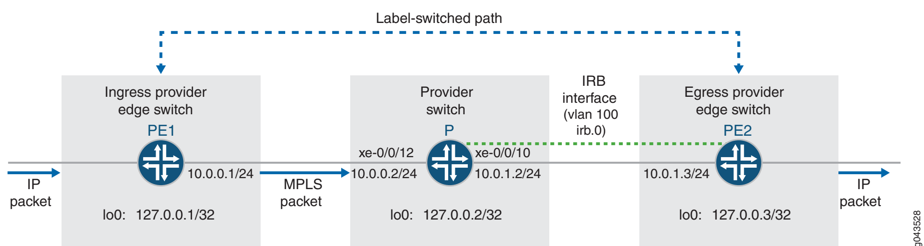

Configuration

To configure the topology in this example, perform these tasks:

- Configuring the Local Ingress PE Switch on page 198
- Configuring the Provider Switch on page 200
- Configuring the Remote Egress PE Switch on page 203

Configuring the Local Ingress PE Switch

CLI Quick Configuration

To quickly configure the local ingress PE switch (PE1), copy and paste the following commands into the switch terminal window of switch PE1:

```
set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.1/24
set interfaces xe-0/0/12 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set routing-options router-id 127.0.0.1
set routing-options autonomous-system 65550
set routing-options forwarding-table export pplb
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface em0.0 disable
set protocols ldp interface xe-0/0/12.0
set protocols ldp interface lo0.0
set policy-options policy-statement pplb then load-balance per-packet
```

Step-by-Step Procedure

To configure the ingress PE switch (PE1):

1. Configure the interfaces.

   [edit interfaces]
   user@switchPE1# **set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.1/24**
   user@switchPE1# **set interfaces xe-0/0/12 unit 0 family mpls**
   user@switchPE1# **set interfaces lo0 unit 0 family inet address 127.0.0.1/32**

2. Configure the router identifier and autonomous system (AS) number.

   NOTE: We recommend that you explicitly configure the router identifier under the [edit routing-options] hierarchy level to prevent unpredictable behavior if the interface address on a loopback interface changes.

   [edit routing-options]
   user@switchPE1# **router-id 127.0.0.1/32**
   user@switchPE1# **set autonomous-system 65550**

3. Apply an export routing policy to the forwarding table for per-packet load balancing.

   [edit routing-options]
   user@switchPE1# **set routing-options forwarding-table export pplb**

4. Configure traffic engineering on all the interfaces.

   [edit protocols mpls]
   user@switchPE1# **set mpls interface all**

5. Configure OSPF with traffic engineering enabled.

   [edit protocols ospf]
   user@switchPE1# **set ospf traffic-engineering**

6. Create an OSPF area and set the loopback address to be passive.

   [edit protocols ospf}]
   user@switchPE1# **set area 0.0.0.0 interface all**
   user@switchPE1# **set area 0.0.0.0 interface lo0.0 passive**
   user@switchPE1# **set area 0.0.0.0 interface em0.0 disable**

7. Configure LDP on the interfaces.

   [edit protocols ldp]
   user@switchPE1# **set ldp interface xe-0/0/12.0**
   user@switchPE1# **set ldp interface lo0.0**

8. Set the policy to perform per-packet load balancing.

   [edit policy-options]
   user@switchPE1# **set policy-options policy-statement pplb then load-balance per-packet**

Results

Display the results of the PE1 switch configuration:

```
user@switchPE1# show configuration
interfaces {
  xe-0/0/12 {
```

```
        unit 0 {
            family inet {
                address 10.0.1/24;
            }
            family mpls;
        }
    }
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            }
        }
    }
}
routing-options {
    router-id 127.0.0.1
    autonomous-system 65550
    forwarding-table {
        export pplb;
    }
}
protocols {
    mpls {
        interface all;
        interface lo0.0;
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface lo0.0 {
                passive;
            }
            interface em0.0 {
                disable;
            }
        }
    }
    ldp {
        interface xe-0/0/12.0
        interface lo0.0;
    }
}
```

Configuring the Provider Switch

CLI Quick Configuration

To quickly configure the provider switch (P), copy and paste the following commands into the switch terminal window of the P switch:

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members v100
set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.2/24
set interfaces xe-0/0/12 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.2/32
set interfaces irb unit 0 family inet address 10.0.1.2/24
```

```
set interfaces irb unit 0 family mpls
set routing-options router-id 127.0.0.2
set routing-options autonomous-system 65550
set routing-options forwarding-table export pplb
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface em0.0 disable
set protocols ldp interface all
set policy-options policy-statement pplb then load-balance per-packet
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.0
```

Step-by-Step Procedure

To configure the provider switch (P):

1. Configure the interfaces.

   ```
   [edit interfaces]
   user@switchP# set interfaces xe-0/0/10 unit 0 family ethernet-switching
     interface-mode trunk
   user@switchP# set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan
     members v100
   user@switchP#  set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.2/24
   user@switchP#  set interfaces xe-0/0/12 unit 0 family mpls
   ```

2. Configure an IRB interface.

   ```
   [edit interfaces]
   user@switchP# set interfaces irb unit 0 family inet address 10.0.1.2/24
   user@switchP# set interfaces irb unit 0 family mpls
   user@switchP# set interfaces lo0 unit 0 family inet address 127.0.0.2/32
   ```

3. Configure the router identifier and AS number.

   NOTE: We recommend that you explicitly configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

   ```
   [edit routing-options]
   user@switchPE1# router-id 127.0.0.2
   user@switchPE1# set autonomous-system 65550
   ```

4. Apply an export routing policy to the forwarding table for per-packet load balancing.

   ```
   [edit routing-options]
   user@switchPE1# set routing-options forwarding-table export pplb
   ```

5. Configure OSPF with traffic engineering enabled.

   ```
   [edit protocols ospf]
   user@switchPE1# set ospf traffic-engineering
   ```

6. Create an OSPF area and set the loopback address to passive.

   ```
   [edit protocols ospf]
   ```

```
user@switchPE1# set area 0.0.0.0 interface all
user@switchPE1# set area 0.0.0.0 interface lo0.0 passive
user@switchPE1# set area 0.0.0.0 interface em0.0 disable
```

7. Configure LDP to include all provider interfaces.

```
[edit protocols ldp]
user@switchP# set protocols ldp interface all
```

8. Create the VLAN by assigning it a name and a VLAN ID.

```
[edit interfaces]
user@switchP# set vlans v100 vlan-id 100
```

9. Link the VLAN to the logical Layer 3 IRB interface.

```
[edit interfaces]
user@switchP# set vlans v100 l3-interface irb.0
```

NOTE: Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is switched, while across VLANs, traffic is routed.

Results Display the results of the provider switch configuration:

```
user@switchP# show configuration
interfaces {
  xe-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
  xe-0/0/12 {
    unit 0
    family inet {
      address 10.0.0.2/24;
    }
    family mpls;
  }
}
irb {
  unit 0 {
    family inet {
      address 10.0.1.2/24;
    }
    family mpls;
  }
}
lo0 {
```



```
  unit 0 {
    family inet {
      address 127.0.0.2/32;
    }
  }
}
routing-options {
  router-id 127.0.0.2;
  autonomous-system 65550;
  forwarding-table {
    export pplb;
  }
}
protocols {
  mpls {
    interface all;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface lo0.0 {
        passive;
      }
      interface em0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
  }
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}
vlans {
  v100 {
    vlan-id 100;
    l3-interface irb.0;
  }
}
```

Configuring the Remote Egress PE Switch

CLI Quick Configuration

To quickly configure the remote egress PE switch (PE2), copy and paste the following commands into the switch terminal window of PE2:

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk
```

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 0 family inet address 10.0.1.3/24
set interfaces irb unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.3/32
set routing-options router-id 127.0.0.3
set routing-options autonomous-system 65550
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface em0.0 disable
set protocols ldp interface irb.0
set protocols ldp interface lo0.0
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.0
```

Step-by-Step Procedure

To configure the remote PE switch (PE2):

1. Configure the interfaces.

   ```
   [edit interfaces]
   user@switchPE2# set interfaces xe-0/0/10 unit 0 family ethernet-switching
     interface-mode trunk
   user@switchPE2# set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan
     members v100
   ```

2. Configure an IRB interface.

   ```
   [edit interfaces]
   set interfaces irb unit 0 family inet address 10.0.1.3/24
   set interfaces irb unit 0 family mpls
   set interfaces lo0 unit 0 family inet address 127.0.0.3/32
   ```

3. Configure the routing options.

   ```
   [edit routing-options]
   user@switchPE2# set routing-options router-id 127.0.0.3/32
   user@switchPE2# set routing-options autonomous-system 65550
   user@switchPE2# set routing-options forwarding-table export pplb
   ```

4. Enable OSPF traffic engineering.

   ```
   [edit protocols ospf]
   user@switchPE2# set protocols ospf traffic-engineering
   user@switchPE2# set protocols ospf area 0.0.0.0 interface all
   user@switchPE2# set protocols ospf area 0.0.0.0 interface lo0.0 passive
   user@switchPE2# set protocols ospf area 0.0.0.0 interface em0.0 disable
   ```

5. Configure LDP to include all provider interfaces.

   ```
   [edit protocols ldp]
   user@switchPE2# set protocols ldp interface irb.0
   user@switchPE2# set protocols ldp interface lo0.0
   ```

6. Create the VLAN by assigning it a name and a VLAN ID.

   ```
   [edit interfaces]
   user@switchPE2# set vlans v100 vlan-id 100
   ```


7. Link the Layer 2 VLAN to the logical Layer 3 IRB interface.

```
[edit interfaces]
user@switchPE2# set vlans v100 l3-interface irb.0
```

Results Display the results of the PE2 switch configuration:

```
user@switchPE2# show configuration
interfaces {
   xe-0/0/10 {
      unit 0 {
         family ethernet-switching {
            interface-mode trunk;
            vlan {
               members v100;
            }
         }
      }
}
irb {
   unit 0 {
      family inet {
         address 10.0.1.3/24;
      }
      family mpls;
   }
}
lo0 {
   unit 0 {
      family inet {
         address 127.0.0.3/32;
      }
   }
}

routing-options {
   router-id 127.0.0.3;
   autonomous-system 65550;
   forwarding-table {
      export pplb;
   }
}

protocols {
   mpls {
      interface all;
   }
   ospf {
      traffic-engineering;
      area 0.0.0.0 {
         interface all;
         interface lo0.0 {
            passive;
         }
         interface em0.0 {
            disable;
```

```
            }
          }
        }
        ldp {
          interface all;
        }
      }
      policy-options {
        policy-statement pplb {
          then {
            load-balance per-packet;
          }
        }
      }
      vlans {
        v100 {
          vlan-id 100;
          l3-interface irb.0;
        }
      }
```

Release History Table

Release	Description
14.1X53-D40	This example shows how to configure integrated routing and bridging (IRB) interfaces over an MPLS core network on QFX5100 switches.

Related Documentation

- Configuring IRB Interfaces
- Understanding Integrated Routing and Bridging
- Understanding Bridging and VLANs
- Understanding Layer 3 Logical Interfaces

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

This example shows how to configure Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. External BGP (EBGP) is used between the customer edge (CE) and provider edge (PE) devices. The remote CE devices have different AS numbers for loop detection.

Requirements

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

Overview

Detailed information about the Juniper Networks implementation of IPv6 over MPLS is described in the following Internet drafts:

- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

These Internet drafts are available on the IETF website at **http://www.ietf.org/**.

This example shows you how to interconnect a two IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the routers in your core network. Multiprotocol Border Gateway Protocol (MP-BGP) is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

In Figure 14 on page 207, PE1 and PE2 are dual-stack BGP routers or switches, meaning they have both IPv4 and IPv6 stacks. The PE devices link the IPv6 networks through the customer edge (CE) routers or switches to the IPv4 core network. The CE devices and the PE devices connect through a link layer that can carry IPv6 traffic. The PE devices use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.

Figure 14: IPv6 Networks Linked by MPLS IPv4 Tunnels

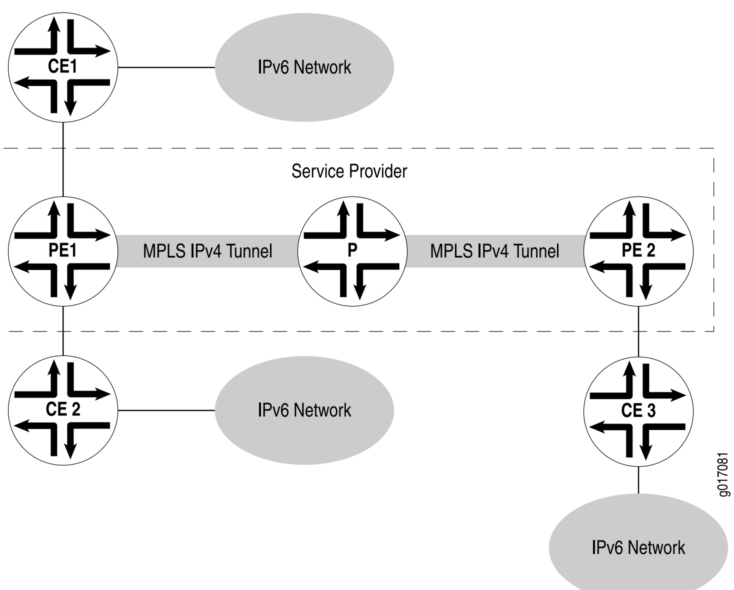

The two PE devices are linked through an MP-BGP session using IPv4 addresses. They use the session to exchange IPv6 routes with an IPv6 (value 2) address family indicator (AFI) and a subsequent AFI (SAFI) (value 4). Each PE router sets the next hop for the IPv6 routes advertised on this session to its own IPv4 address. Because MP-BGP requires the BGP next hop to correspond to the same address family as the network layer reachability information (NLRI), this IPv4 address needs to be embedded within an IPv6 format.

The PE devices can learn the IPv6 routes from the CE devices connected to them using MP-BGP or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE device and CE device could occur over an IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGP, and between BGP and any other protocol.

The PE routers have MPLS LSPs routed to each others' IPv4 addresses. IPv4 provides signaling for the LSPs by means of RSVP. These LSPs are used to resolve the next-hop addresses of the IPv6 routes learned from MP-BGP. The next hops use IPv4-mapped IPv6 addresses, while the LSPs use IPv4 addresses.

The PE devices always advertise IPv6 routes to each other using a label value of 2, the explicit null label for IPv6 as defined in RFC 3032, *MPLS Label Stack Encoding*. As a consequence, each of the forwarding next hops for the IPv6 routes learned from remote PE routers normally push two labels. The inner label is 2 (this label could be different if the advertising PE device is not a Juniper Networks routing or switching platform), and the outer label is the LSP label. If the LSP is a single-hop LSP, then only Label 2 is pushed.

It is also possible for the PE devices to exchange plain IPv6 routes using SAFI unicast. However, there is one major advantage in exchanging labeled IPv6 routes. The penultimate-hop router for an MPLS LSP can pop the outer label and then send the packet with the inner label as an MPLS packet. Without the inner label, the penultimate-hop router would need to discover whether the packet is an IPv4 or IPv6 packet to set the protocol field in the Layer 2 header correctly.

When the PE1 device in Figure 14 on page 207 receives an IPv6 packet from the CE1 device, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 device, then no labels need to be pushed and the packet is simply sent to the CE2 device. If the destination matches a prefix that was learned from the PE2 device, then the PE1 router pushes two labels onto the packet and sends it to the Provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

Each provider router in the service provider's network handles the packet as it would any MPLS packet, swapping labels as it passes from provider router to provider router. The penultimate-hop provider router for the LSP pops the outer label and sends the packet to the PE2 router. When the PE2 router receives the packet, it recognizes the IPv6 explicit null label on the packet (Label 2). It pops this label and treats it as an IPv6 packet, performing a lookup in the IPv6 forwarding table and forwarding the packet to the CE3 router.

This example includes the following settings:

- In addition to configuring the **family inet6** statement on all the CE router–facing interfaces, you must also configure the statement on all the core-facing interfaces running MPLS. Both configurations are necessary because the router must be able to process any IPv6 packets it receives on these interfaces. You should not see any regular

IPv6 traffic arrive on these interfaces, but you will receive MPLS packets tagged with Label 2. Even though Label 2 MPLS packets are sent in IPv4, these packets are treated as native IPv6 packets.

- You enable IPv6 tunneling by including the **ipv6-tunneling** statement in the configuration for the PE routers. This statement allows IPv6 routes to be resolved over an MPLS network by converting all routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.

NOTE: BGP automatically runs its import policy even when copying routes from a primary routing table group to a secondary routing table group. If IPv4 labeled routes arrive from a BGP session (for example, when you have configured the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level on the PE router), the BGP neighbor's import policy also accepts IPv6 routes, since the neighbor's import policy is run while doing the copy operation to the inet6.3 routing table.

- When you configure MP-BGP to carry IPv6 traffic, the IPv4 MPLS label is removed at the destination PE router. The remaining IPv6 packet without a label can then be forwarded to the IPv6 network. To enable this, include the **explicit-null** statement in the BGP configuration.

Configuration

CLI Quick Configuration

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

Device PE1

```
set interfaces xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
set interfaces xe-0/0/5 unit 2 family mpls
set interfaces xe-0/0/6 unit 5 family inet address 10.1.1.5/30
set interfaces xe-0/0/6 unit 5 family inet6
set interfaces xe-0/0/6 unit 5 family mpls
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.2
set protocols mpls interface xe-0/0/6.5
set protocols bgp group toCE1 type external
set protocols bgp group toCE1 local-address ::10.1.1.2
set protocols bgp group toCE1 family inet6 unicast
set protocols bgp group toCE1 export send-bgp6
set protocols bgp group toCE1 peer-as 1
set protocols bgp group toCE1 neighbor ::10.1.1.1
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 1.1.1.2
set protocols bgp group toPE2 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE2 export next-hop-self
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface xe-0/0/6.5
```

```
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols rsvp interface xe-0/0/6.5
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 2
```

Device PE2

```
set interfaces xe-0/0/5 unit 10 family inet address 10.1.1.10/30
set interfaces xe-0/0/5 unit 10 family inet6
set interfaces xe-0/0/5 unit 10 family mpls
set interfaces xe-0/0/6 unit 13 family inet6 address ::10.1.1.13/126
set interfaces xe-0/0/6 unit 13 family mpls
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.10
set protocols mpls interface xe-0/0/6.13
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 local-address 1.1.1.4
set protocols bgp group toPE1 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE1 export next-hop-self
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 neighbor 1.1.1.2
set protocols bgp group toCE3 type external
set protocols bgp group toCE3 local-address ::10.1.1.13
set protocols bgp group toCE3 family inet6 unicast
set protocols bgp group toCE3 export send-bgp6
set protocols bgp group toCE3 peer-as 3
set protocols bgp group toCE3 neighbor ::10.1.1.14
set protocols ospf area 0.0.0.0 interface xe-0/0/5.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols rsvp interface xe-0/0/5.10
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 2
```

Device P

```
set interfaces xe-0/0/5 unit 6 family inet address 10.1.1.6/30
set interfaces xe-0/0/5 unit 6 family inet6
set interfaces xe-0/0/5 unit 6 family mpls
set interfaces xe-0/0/6 unit 9 family inet address 10.1.1.9/30
set interfaces xe-0/0/6 unit 9 family inet6
set interfaces xe-0/0/6 unit 9 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
```



```
set protocols mpls interface xe-0/0/5.6
set protocols mpls interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface xe-0/0/5.6
set protocols ospf area 0.0.0.0 interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols rsvp interface xe-0/0/5.6
set protocols rsvp interface xe-0/0/6.9
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
```

Device CE1

```
set interfaces xe-0/0/5 unit 1 family inet6 address ::10.1.1.1/126
set interfaces xe-0/0/5 unit 1 family mpls
set interfaces lo0 unit 1 family inet6 address ::1.1.1.1/128
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 local-address ::10.1.1.1
set protocols bgp group toPE1 family inet6 unicast
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor ::10.1.1.2
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1
```

Device CE3

```
set interfaces xe-0/0/5 unit 14 family inet6 address ::10.1.1.14/126
set interfaces xe-0/0/5 unit 14 family mpls
set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
set protocols bgp group toPE2 type external
set protocols bgp group toPE2 local-address ::10.1.1.14
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 peer-as 2
set protocols bgp group toPE2 neighbor ::10.1.1.13
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 3
```

Configuring Device PE1

Step-by-Step Procedure

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

To configure Device PE1:

1. Configure the interfaces.

```
[edit interfaces]
user@PE1# set xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set xe-0/0/5 unit 2 family mpls

user@PE1# set xe-0/0/6 unit 5 family inet address 10.1.1.5/30
```

```
user@PE1# set xe-0/0/6 unit 5 family inet6
user@PE1# set xe-0/0/6 unit 5 family mpls

user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32
```

2. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface xe-0/0/5.2
user@PE1# set interface xe-0/0/6.5
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE1# set group toCE1 type external
user@PE1# set group toCE1 local-address ::10.1.1.2
user@PE1# set group toCE1 family inet6 unicast
user@PE1# set group toCE1 export send-bgp6
user@PE1# set group toCE1 peer-as 1
user@PE1# set group toCE1 neighbor ::10.1.1.1

user@PE1# set group toPE2 type internal
user@PE1# set group toPE2 local-address 1.1.1.2
user@PE1# set group toPE2 family inet6 labeled-unicast explicit-null
user@PE1# set group toPE2 export next-hop-self
user@PE1# set group toPE2 export send-v6
user@PE1# set group toPE2 neighbor 1.1.1.4
```

4. Configure OSPF

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface xe-0/0/6.5
user@PE1# set interface lo0.2 passive
```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set rsvp interface xe-0/0/6.5
```

6. Configure the routing policies.

```
[edit policy-options]
user@PE1# set policy-statement next-hop-self then next-hop self

user@PE1# set policy-statement send-bgp6 from family inet6
user@PE1# set policy-statement send-bgp6 from protocol bgp
user@PE1# set policy-statement send-bgp6 then accept

user@PE1# set policy-statement send-v6 from family inet6
user@PE1# set policy-statement send-v6 from protocol bgp
user@PE1# set policy-statement send-v6 from protocol direct
user@PE1# set policy-statement send-v6 then accept
```

7. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 2
```


Results From configuration mode, confirm your configuration by entering 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.

```
user@R1# show interfaces
xe-0/0/5 {
  unit 2 {
    family inet6 {
      address ::10.1.1.2/126;
    }
    family mpls;
  }
}
xe-0/0/6 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family inet6;
    family mpls;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 1.1.1.2/32;
    }
  }
}

user@R1# show policy-options
policy-statement next-hop-self {
  then {
    next-hop self;
  }
}
policy-statement send-bgp6 {
  from {
    family inet6;
    protocol bgp;
  }
  then accept;
}
policy-statement send-v6 {
  from {
    family inet6;
    protocol [ bgp direct ];
  }
  then accept;
}

user@R1# show protocols
mpls {
  ipv6-tunneling;
  interface xe-0/0/5.2;
```

```
    interface xe-0/0/6.5;
  }
  bgp {
    group toCE1 {
      type external;
      local-address ::10.1.1.2;
      family inet6 {
        unicast;
      }
      export send-bgp6;
      peer-as 1;
      neighbor ::10.1.1.1;
    }
    group toPE2 {
      type internal;
      local-address 1.1.1.2;
      family inet6 {
        labeled-unicast {
          explicit-null;
        }
      }
      export [ next-hop-self send-v6 ];
      neighbor 1.1.1.4;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface xe-0/0/6.5;
      interface lo0.2 {
        passive;
      }
    }
  }
  rsvp {
    interface xe-0/0/6.5;
  }
```

```
user@R1# show routing-options
router-id 1.1.1.2;
autonomous-system 2;
```

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

Configure the other devices in the topology, as shown in "CLI Quick Configuration" on page 209.

Verification

Confirm that the configuration is working properly.

Verifying That the CE Devices Have Connectivity

Purpose Make sure that the tunnel is operating.

Action From operational mode, enter the **ping** command.

```
user@CE1> ping ::10.1.1.14
PING6(56=40+8+8 bytes) ::10.1.1.1 --> ::10.1.1.14
16 bytes from ::10.1.1.14, icmp_seq=0 hlim=61 time=10.687 ms
16 bytes from ::10.1.1.14, icmp_seq=1 hlim=61 time=9.239 ms
16 bytes from ::10.1.1.14, icmp_seq=2 hlim=61 time=1.842 ms

user@CE3> ping ::10.1.1.1
PING6(56=40+8+8 bytes) ::10.1.1.14 --> ::10.1.1.1
16 bytes from ::10.1.1.1, icmp_seq=0 hlim=61 time=1.484 ms
16 bytes from ::10.1.1.1, icmp_seq=1 hlim=61 time=1.338 ms
16 bytes from ::10.1.1.1, icmp_seq=2 hlim=61 time=1.351 ms
```

Meaning The IPv6 CE devices can communicate over the core IPv4 network.

Related Documentation

Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action

```
user@switch> show interfaces xe-* terse

Interface               Admin Link Proto    Local                Remote
xe-0/0/0                up    up
xe-0/0/0.0              up    up
xe-0/0/1.0              up    up
xe-0/0/2.0              up    up
xe-0/0/3.0              up    up   inet      2.2.2.1/16
xe-0/0/4.0              up    up
xe-0/0/5.0              up    up   inet      10.1.5.1/24
                                   mpls
xe-0/0/6.0              up    up   inet      10.1.6.1/24
                                   mpls
```

Meaning The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

Action `user@switch>` **show ospf neighbor**

```
Address          Interface              State     ID               Pri  Dead
127.1.1.1        xe-0/0/5               Full      10.10.10.10      128    39
```

Meaning The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the *Junos OS Routing Protocols and Policies Command Reference* .

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

Action `user@switch>` **show mpls interface**

```
Interface        State       Administrative groups
ge-0/0/5            Up          <none>
ge-0/0/6            Up          <none>
```

Meaning The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

Verifying RSVP

Purpose Verify the state of the RSVP session. You should perform this verification task on each of the switches.

Action

```
user@switch> show rsvp session

Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.3       127.1.1.1        Up       0  1 FF       -    300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.1       127.1.1.3        Up       0  1 FF   299968        - lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are up.

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183

MPLS Configuration Guidelines

When configuring MPLS on QFX Series devices or on EX4600, note that the number of IP prefixes supported depends on the specific platform being used. See the scale specifications in the data sheet of your device for additional information.

- We recommend the following:
  - If your ingress provider edge (PE) switch needs to support more than 8000 external IP prefixes, use a larger capacity device as an ingress PE switch.
  - If you use a switch as a route reflector for BGP labeled routes, use it as a dedicated route reflector (that is, the switch must not participate in managing data traffic).
  - If you use a switch as a PE switch or as a route reflector for BGP labeled routes, configure routing policies on the PE switch and the route reflector to filter external IP routes from the routing table.

    The configuration example for a routing policy named fib_policy (at the **[edit policy-options** and **[edit routing-options** hierarchy levels) to filter BGP labeled routes from the inet.0 routing table is given below:

```
user@switch# show policy-options
policy-statement fib_policy {
  from {
    protocol bgp;
    rib inet.0;
  }
  then reject;
}

user@switch# show routing-options
forwarding-table {
```

```
    export fib_policy;
}
```

- Packet fragmentation using the **allow-fragmentation** statement at the **[edit protocols mpls path-mtu]** hierarchy level is not supported on QFX Series devices or on the EX4600 switch. Therefore, you must ensure that the maximum transmission unit (MTU) values configured on every MPLS interface is sufficient to handle MPLS packets. The packets whose size exceeds the MTU value of an interface will be dropped.

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183
- Configuring a Global MPLS EXP Classifier on page 175
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138

Supported MPLS Scaling Values

This topic lists the MPLS scaling values supported on QFX Series switches.

Table 23 on page 218 lists the MPLS scaling values supported on Juniper QFX switches and on the EX4600 switch.

Table 23: MPLS Scaling Values

Feature	QFX3500 Scaling Value	QFX5100 and EX4600 Scaling Value
Maximum number of MPLS labels in a packet's label stack	3 labels for Push operations 2 labels for Pop operations 1 label for Swap operations	3 labels for Push operations 2 labels for Pop operations 1 label for Swap operations
Maximum number of MPLS labels on provider switches	4096	16386
Maximum number of tunnel (combination of routes and LSPs) initiations	Ingress LSPs: 1024 Transit LSPs: 4000	Ingress LSPs: 1024 Transit LSPs: 16386
Maximum number of unique next-hops on egress provider edge (PE) switches	512	512
Maximum number of MPLS firewall filters	768	1536
Virtual Routing and Forwarding (VRF)	1K	1K

Table 23: MPLS Scaling Values *(continued)*

Feature	QFX3500 Scaling Value	QFX5100 and EX4600 Scaling Value
Layer 3 Host	IPV4: 8K	See *Understanding the Unified Forwarding Table*.
Layer 3 Longest Prefix Match (LPM)	IPV4: 16K IPV6: 4K	See *Understanding the Unified Forwarding Table*.

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- MPLS Configuration Guidelines on page 217

Using MPLS Stitching with BGP to Connect Virtual Machines

The stitching feature of Junos OS provides connectivity between virtual machines VMs) that reside either on opposite sides of data center routers or in different data centers. An external controller, programmed in the data-plane, assigns MPLS labels to both VMs and servers. Then, the signaled MPLS labels are used between the data center routers, generating static label-switched paths (LSPs), resolved over either BGP labeled unicast, RSVP or LDP, to provide the routes dictated by the labels.

- Requirements on page 219
- Overview on page 219
- Configuration on page 221

Requirements

This example uses the following hardware and software components:

- Two QFX5100 top-of-rack switches named Baby Dutchmen
- Two MX480 routers named Coleman and Taxes
- Two VMs configured on servers located on either side of the topology
- Junos OS Release 14.1X53-D25

Before you begin, be sure you have:

- Configured VMs, one on each of two servers. The servers are located on either side of the topology as shown in Figure 15 on page 221.
- Installed two routers and two top-of-rack switches as shown in the topology shown in Figure 15 on page 221.

Overview

This example shows MPLS stitching resolved over BGP from top-of-rack switch Baby to the top-of-rack switch Dutchmen through the routers named Coleman and Taxes.

Table 24: Topology Components for Top-of-Rack Named Baby

Interface	Description
ge-8/1/2	Port on the switch Baby connected to the router Taxes
ge-8/1/3	Port on the switch Baby connected to the router Coleman
ge-8/2/4	Port on the switch Baby connected to a VM

Table 25: Topology Components for Top-of-Rack Switch Dutchmen

Interface	Description
ge-0/0/0	Port on the top-of-rack switch Dutchmen connected to a VM
ge-0/1/0	Port on the top-of-rack switch Dutchmen connected to the router named R2
ge-0/2/0	Port on the top-of-rack switch Dutchmen connected to the router named R2

Table 26: Topology Components for MX Series Router Named Coleman

Interface	Description
ge-2/0/0	Port on router named Coleman connected to top-of-rack switch named Baby
ge-2/0/1	Port on router named Coleman connected to top-of-rack switch named Baby
ge-2/0/5	Port on router named Coleman connected to top-of-rack switch named Dutchmen
ge-2/0/6	Port on router named Coleman connected to top-of-rack switch named Dutchmen

Table 27: Topology Components for MX Series Router Named Taxes

Interface	Description
ge-0/0/3	Port on router named Taxes connected to router named Coleman
ge-1/0/0	Port on router named Taxes connected to router named Coleman
ge-1/0/2	Port on router named Taxes connected to top-of-rack switch named Dutchmen
ge-1/2/0	Port on router named Taxes connected to top-of-rack switch named Dutchmen

Topology

Figure 15: MPLS Stitching to Connect Virtual Machines

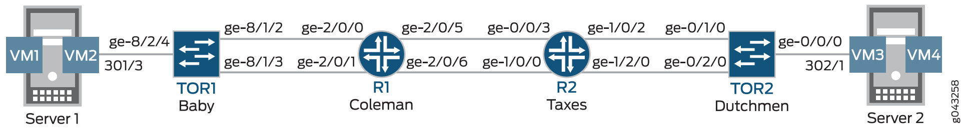

Configuration

To configure MPLS stitching, perform these tasks:

- Configuring Stitching on Top-of-Rack Switch Named Baby on page 224
- Configuring Stitching on Top-of-Rack Switch Named Dutchmen on page 225
- [xref target has no title]
- Configuring MX Series Router Named Taxes on page 227
- Results on page 228

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.

Configure the top-of-rack switch named Baby:

```
set interfaces ge-8/1/2 unit 0 family inet address 2.1.1.1/24
set interfaces ge-8/1/2 unit 0 family mpls
set interfaces ge-8/1/3 unit 0 family inet address 2.1.2.1/24
set interfaces ge-8/1/3 unit 0 family mpls
set interfaces ge-8/2/4 unit 0 family inet address 1.1.1.1/24
set interfaces ge-8/2/4 unit 0 family mpls
set routing-options interface-routes rib-group inet inet0toinet3
set routing-options rib-groups inet0toinet3 import-rib inet.0
set routing-options rib-groups inet0toinet3 import-rib inet.3
set routing-options router-id 10.9.82.49
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 next-hop
  10.9.82.47
set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 stitch
set protocols mpls interface all set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp-R1 type internal
set protocols bgp group ibgp-R1 local-address 10.9.82.49
set protocols bgp group ibgp-R1 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp-R1 export advLo
set protocols bgp group ibgp-R1 neighbor 10.9.82.46
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement advLo term lo0 from protocol direct
set policy-options policy-statement advLo term lo0 from route-filter 10.9.82.49/32 exact
```

```
set policy-options policy-statement advLo term lo0 then next-hop 2.1.1.1
set policy-options policy-statement advLo term lo0 then accept
set policy-options policy-statement advLo term other then reject
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
```

Configure the top-of-rack switch Dutchmen:

```
set interfaces ge-0/0/0 unit 0 family inet address 5.1.1.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/1/0 unit 0 family inet address 4.1.1.2/24
set interfaces ge-0/1/0 unit 0 family mpls
set interfaces ge-0/2/0 unit 0 family inet address 4.1.2.2/24
set interfaces ge-0/2/0 unit 0 family mpls
set routing-options interface-routes rib-group inet inet0toinet3
set routing-options rib-groups inet0toinet3 import-rib inet.0
set routing-options rib-groups inet0toinet3 import-rib inet.3
set routing-options router-id 10.9.82.47
set routing-options autonomous-system 102
set routing-options forwarding-table export pplb
set protocols mpls static-label-switched-path to-server transit 1000001 next-hop 5.1.1.2
set protocols mpls static-label-switched-path to-server transit 1000001 pop
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp-R2 type internal
set protocols bgp group ibgp-R2 local-address 10.9.82.47
set protocols bgp group ibgp-R2 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp-R2 export advLo
set protocols bgp group ibgp-R2 neighbor 10.9.82.48
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement advLo term lo0 from protocol direct
set policy-options policy-statement advLo term lo0 from route-filter 10.9.82.47/32 exact
set policy-options policy-statement advLo term lo0 then next-hop 4.1.1.2
set policy-options policy-statement advLo term lo0 then accept
set policy-options policy-statement advLo term other then reject
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
```

Configure the MX Series router named Coleman:

```
set interfaces ge-2/0/0 unit 0 family inet address 2.1.1.2/24
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 family inet address 2.1.2.2/24
set interfaces ge-2/0/5 unit 0 family inet address 3.1.1.1/24
set interfaces ge-2/0/5 unit 0 family mpls
set interfaces ge-2/0/6 unit 0 family inet address 3.1.2.1/24
set interfaces ge-2/0/6 unit 0 family mpls
set routing-options interface-routes rib-group inet inet0toinet3
set routing-options rib-groups inet0toinet3 import-rib inet.0
set routing-options rib-groups inet0toinet3 import-rib inet.3
set routing-options router-id 10.9.82.46
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```



```
set protocols bgp group ibgp-ToR1 type internal
set protocols bgp group ibgp-ToR1 accept-remote-nexthop
set protocols bgp group ibgp-ToR1 local-address 10.9.82.46
set protocols bgp group ibgp-ToR1 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp-ToR1 export advToR2
set protocols bgp group ibgp-ToR1 export advLo
set protocols bgp group ibgp-ToR1 neighbor 10.9.82.49
set protocols bgp group ebgp-R2 type external
set protocols bgp group ebgp-R2 local-address 3.1.2.1
set protocols bgp group ebgp-R2 family inet labeled-unicast rib inet.3
set protocols bgp group ebgp-R2 peer-as 102
set protocols bgp group ebgp-R2 local-as 100
set protocols bgp group ebgp-R2 neighbor 3.1.2.2 as-override
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

Configure the MX Series router named Taxes:

```
set interfaces ge-0/0/3 unit 0 family inet address 3.1.1.2/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-1/0/0 unit 0 family inet address 3.1.2.2/24
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/2 unit 0 family inet address 4.1.1.1/24
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-1/2/0 unit 0 family inet address 4.1.2.1/24
set interfaces ge-1/2/0 unit 0 family mpls
set routing-options interface-routes rib-group inet inet0toinet3
set routing-options rib-groups inet0toinet3 import-rib inet.0
set routing-options rib-groups inet0toinet3 import-rib inet.3
set routing-options router-id 10.9.82.48 set routing-options autonomous-system 102
set routing-options forwarding-table export pplb
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ebgp-R1 type external
set protocols bgp group ebgp-R1 local-address 3.1.2.2
set protocols bgp group ebgp-R1 family inet labeled-unicast rib inet.3
set protocols bgp group ebgp-R1 peer-as 100
set protocols bgp group ebgp-R1 local-as 102
set protocols bgp group ebgp-R1 neighbor 3.1.2.1 as-override
set protocols bgp group ibgp-ToR2 type internal
set protocols bgp group ibgp-ToR2 local-address 10.9.82.48
set protocols bgp group ibgp-ToR2 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp-ToR2 export advToR2
set protocols bgp group ibgp-ToR2 export advLo
set protocols bgp group ibgp-ToR2 neighbor 10.9.82.47
set protocols ospf traffic-engineering set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement advLo term lo0 from protocol direct
set policy-options policy-statement advLo term lo0 from route-filter 10.9.82.48/32 exact
set policy-options policy-statement advLo term lo0 then next-hop 4.1.1.1
set policy-options policy-statement advLo term lo0 then accept
set policy-options policy-statement advLo term other then reject
set policy-options policy-statement advToR2 term ToR2lo from protocol bgp
set policy-options policy-statement advToR2 term ToR2lo from route-filter 10.9.82.48/32
  exact
```

```
set policy-options policy-statement advToR2 term ToR2lo then accept
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement pplb then accept
```

Configuring Stitching on Top-of-Rack Switch Named Baby

Step-by-Step Procedure

To configure stitching on the top-of-rack switch named Baby, follow these steps:

1. Configure three MPLS interfaces on top-of-rack switch named Baby.

```
user@baby# [edit interfaces]
user@baby# set ge-8/1/2 unit 0 family inet address 2.1.1.1/24
user@baby# set ge-8/1/2 unit 0 family mpls
user@baby# set ge-8/1/3 unit 0 family inet address 2.1.2.1/24
user@baby# set ge-8/1/3 unit 0 family mpls
user@baby# set ge-8/2/4 unit 0 family inet address 1.1.1.1/24
user@baby# set ge-8/2/4 unit 0 family mpls
```

2. Configure routing options on top-of-rack switch named Baby.

```
user@baby# [edit routing-options]
user@baby# set interface-routes rib-group inet inet0toinet3
user@baby# set rib-groups inet0toinet3 import-rib inet.0
user@baby# set rib-groups inet0toinet3 import-rib inet.3
user@baby# set router-id 10.9.82.49
user@baby# set autonomous-system 100
user@baby# set forwarding-table export pplb
```

3. Configure the protocols MPLS, BGP, and OSPF on top-of-rack switch named Baby.

```
user@baby# edit [protocols mpls]
user@baby# set static-label-switched-path static-to-ToR2 transit 1000000
  next-hop 10.9.82.47
user@baby# set static-label-switched-path static-to-ToR2 transit 1000000
  next-hop 10.9.82.47
user@baby# set static-label-switched-path static-to-ToR2 transit 1000000 stitch
set interface all set mpls interface fxp0.0 disable
user@baby# edit [protocols bgp group ibgp-R1]
user@baby# set type internal
user@baby# set local-address 10.9.82.49
user@baby# set family inet labeled-unicast rib inet.3
user@baby# set export advLo
user@baby# set neighbor 10.9.82.46
user@baby# edit [protocols ospf]
user@baby#set traffic-engineering
user@baby# set area 0.0.0.0 interface all
user@baby# set area 0.0.0.0 interface fxp0.0 disable
```

4. Configure policy options on top-of-rack switch named Baby.

```
user@baby# [edit policy-options policy-statement advLo]
user@baby# set term lo0 from protocol direct
user@baby# set term lo0 from route-filter 10.9.82.49/32 exact
user@baby# set term lo0 then next-hop 2.1.1.1
user@baby# set term lo0 then accept
user@baby# set term other then reject
user@baby# [edit policy-options policy-statement pplb]
user@baby# set then load-balance per-packet
```



```
user@baby# set then accept
```

Configuring Stitching on Top-of-Rack Switch Named Dutchmen

Step-by-Step Procedure

To configure stitching on the top-of-rack switch named Dutchmen, follow these steps:

1. Configure three MPLS interfaces on the top-of-rack switch Dutchmen.

   ```
   user@dutchmen# [edit interfaces]
   user@dutchmen# set ge-0/0/0 unit 0 family inet address 5.1.1.1/24
   user@dutchmen# set ge-0/0/0 unit 0 family mpls
   user@dutchmen# set ge-0/1/0 unit 0 family inet address 4.1.1.2/24
   user@dutchmen# set ge-0/1/0 unit 0 family mpls
   user@dutchmen# set ge-0/2/0 unit 0 family inet address 4.1.2.2/24
   user@dutchmen# set ge-0/2/0 unit 0 family mpls
   ```

2. Configure routing options on the top-of-rack switch Dutchmen.

   ```
   user@dutchmen# [edit routing-options]
   user@dutchmen# set interface-routes rib-group inet inet0toinet3
   user@dutchmen# set rib-groups inet0toinet3 import-rib inet.0
   user@dutchmen# set rib-groups inet0toinet3 import-rib inet.3
   user@dutchmen# set router-id 10.9.82.47
   user@dutchmen# set autonomous-system 102
   user@dutchmen# set forwarding-table export pplb
   ```

3. Configure the protocols MPLS, BGP, and OSPF on the top-of-rack switch Dutchmen.

   ```
   user@dutchmen# edit [protocols mpls]
   user@dutchmen# set static-label-switched-path to-server transit 1000001 next-hop
     5.1.1.2
   user@dutchmen# set static-label-switched-path to-server transit 1000001 pop
   user@dutchmen# set interface all
   user@dutchmen# set interface fxp0.0 disable
   user@dutchmen# edit [protocols bgp group ibgp-R2]
   user@dutchmen# set type internal
   user@dutchmen# set local-address 10.9.82.47
   user@dutchmen# set family inet labeled-unicast rib inet.3
   user@dutchmen# set export advLo
   user@dutchmen# set neighbor 10.9.82.48
   user@dutchmen# edit [protocols ospf]
   user@dutchmen# set traffic-engineering
   user@dutchmen# set area 0.0.0.0 interface all
   user@dutchmen# set area 0.0.0.0 interface fxp0.0 disable
   ```

4. Configure policy options on the top-of-rack switch Dutchmen.

   ```
   user@dutchmen# edit [policy-options policy-statement advLo]
   user@dutchmen# set term lo0 from protocol direct
   user@dutchmen# set term lo0 from route-filter 10.9.82.47/32 exact
   user@dutchmen# set term lo0 then next-hop 4.1.1.2
   user@dutchmen# set term lo0 then accept
   user@dutchmen# set term other then reject
   user@dutchmen# edit [policy-options policy-statement pplb]
   set then load-balance per-packet
   set then accept
   ```

Step-by-Step Procedure

To configure the MX Series router named Coleman, follow these steps:

1. Configure four interfaces on the router named Coleman.

```
[edit interfaces]
user@coleman# set ge-2/0/0 unit 0 family inet address 2.1.1.2/24
user@coleman# set ge-2/0/0 unit 0 family mpls
user@coleman# set ge-2/0/1 unit 0 family inet address 2.1.2.2/24
user@coleman# set ge-2/0/5 unit 0 family inet address 3.1.1.1/24
user@coleman# set ge-2/0/5 unit 0 family mpls
user@coleman# set ge-2/0/6 unit 0 family inet address 3.1.2.1/24
user@coleman# set ge-2/0/6 unit 0 family mpls
```

2. Configure routing options on the router named Coleman.

```
[edit routing-options]
user@coleman# set interface-routes rib-group inet inet0toinet3
user@coleman# set rib-groups inet0toinet3 import-rib inet.0
user@coleman# set rib-groups inet0toinet3 import-rib inet.3
user@coleman# set router-id 10.9.82.46
user@coleman# set autonomous-system 100 set routing-options forwarding-table
  export pplb
```

3. Configure the protocols MPLS, BGP, and OSPF on on the router named Coleman.

```
[edit protocols mpls]
user@coleman# set interface all
user@coleman# set interface fxp0.0 disable
[edit protocols bgp group ibgp-ToR1]
user@coleman# set type internal
user@coleman# set accept-remote-nexthop
user@coleman# set local-address 10.9.82.46
user@coleman# set family inet labeled-unicast rib inet.3
user@coleman# set export advToR2
user@coleman# set export advLo
user@coleman# set neighbor 10.9.82.49
[edit protocols bgp group ebgp-R2]
user@coleman# set type external
user@coleman# set local-address 3.1.2.1
user@coleman# set family inet labeled-unicast rib inet.3
user@coleman# set peer-as 102
user@coleman# set local-as 100
user@coleman# set neighbor 3.1.2.2 as-override
[edit protocols ospf]
user@coleman# set traffic-engineering
user@coleman# set area 0.0.0.0 interface all
user@coleman# set area 0.0.0.0 interface fxp0.0 disable
```

4. Configure the policy options on the router named Coleman.

```
[edit policy-options policy-statement advLo]
user@coleman# set term lo0 from protocol direct
user@coleman# set term lo0 from route-filter 10.9.82.46/32 exact
user@coleman# set term lo0 then next-hop 2.1.1.2
user@coleman# set term lo0 then accept
user@coleman# set term other then reject


[edit policy-options policy-statement advToR2 ]
```



```
user@coleman# set term ToR2lo from protocol bgp
user@coleman# set term ToR2lo from protocol bgp then accept
user@coleman# set term ToR2lo from route-filter 10.9.82.47/32 exact
user@coleman# set term ToR2lo then accept

[edit policy-options policy-statement pplb]
user@coleman# set then load-balance per-packet
user@coleman# set then accept
```

Configuring MX Series Router Named Taxes

Step-by-Step Procedure

To configure the MX Series router named Taxes, follow these steps:

1. Configure four MPLS interfaces on the router named Taxes.

```
user@taxes# [edit interfaces]
user@taxes# set ge-0/0/3 unit 0 family inet address 3.1.1.2/24
user@taxes# set ge-0/0/3 unit 0 family mpls
user@taxes# set ge-1/0/0 unit 0 family inet address 3.1.2.2/24
user@taxes# set ge-1/0/0 unit 0 family mpls
user@taxes# set ge-1/0/2 unit 0 family inet address 4.1.1.1/24
user@taxes# set ge-1/0/2 unit 0 family mpls
user@taxes# set ge-1/2/0 unit 0 family inet address 4.1.2.1/24
user@taxes# set ge-1/2/0 unit 0 family mpls
```

2. Configure routing options on the router named Taxes.

```
[edit routing-options]
user@taxes#set interface-routes rib-group inet inet0toinet3
user@taxes#set rib-groups inet0toinet3 import-rib inet.0
user@taxes#set rib-groups inet0toinet3 import-rib inet.3
user@taxes#set router-id 10.9.82.48
user@taxes#set autonomous-system 102
user@taxes#set forwarding-table export pplb
```

3. Configure the protocols MPLS, BGP, and OSPF on the router named Taxes.

```
user@taxes# [edit protocols mpls]
user@taxes# set interface all
user@taxes# set interface fxp0.0 disable
user@taxes# [edit protocols bgp group ebgp-R1]
user@taxes# set type external
user@taxes# set local-address 3.1.2.2
user@taxes# set family inet labeled-unicast rib inet.3
user@taxes# set peer-as 100
user@taxes# set local-as 102
user@taxes# set neighbor 3.1.2.1 as-override
user@taxes# [edit protocols bgp group ibgp-ToR2]
user@taxes# set type internal
user@taxes# set local-address 10.9.82.48
user@taxes# set family inet labeled-unicast rib inet.3
user@taxes# set export advToR2
user@taxes# set export advLo
user@taxes# set neighbor 10.9.82.47
user@taxes# [edit protocols ospf]
user@taxes# set traffic-engineering
user@taxes# set area 0.0.0.0 interface all
```

user@taxes# **set area 0.0.0.0 interface fxp0.0 disable**

4. Configure policy options on the router named Taxes.

   user@taxes# [edit policy-options policy-statement advLo]
   user@taxes# **set term lo0 from protocol direct**
   user@taxes# **set term lo0 from route-filter 10.9.82.48/32 exact**
   user@taxes# **set term lo0 then next-hop 4.1.1.1**
   user@taxes# **set term lo0 then accept**
   user@taxes# **set term other then reject**
   user@taxes# [edit policy-options policy-statement pplb]
   user@taxes# **set then load-balance per-packet**
   user@taxes# **set then accept**

Results

Check the MPLS configuration on the top-of-rack switch Baby by using the **show mpls protocol** command.

```
user@baby# show mpls protocol
interfaces {
    ge-8/1/2 {
        unit 0 {
            family inet {
                address 2.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-8/1/3 {
        unit 0 {
            family inet {
                address 2.1.2.1/24;
            }
            family mpls;
        }
    }
    ge-8/2/4 {
        unit 0 {
            family inet {
                address 1.1.1.1/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
        }
    }
    router-id 10.9.82.49;
    autonomous-system 100;
    forwarding-table {
        export pplb;
    }
```



```
}
protocols {
    mpls {
        static-label-switched-path static-to-ToR2 {
            transit 1000000 {
                next-hop 10.9.82.47;
                stitch;
            }
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group ibgp-R1 {
            type internal;
            local-address 10.9.82.49;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export advLo;
            neighbor 10.9.82.46;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
policy-options {
    policy-statement advLo {
        term lo0 {
            from {
                protocol direct;
                route-filter 10.9.82.49/32 exact;
            }
            then {
                next-hop 2.1.1.1;
                accept;
            }
        }
        term other {
            then reject;
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
            accept;
        }
```

```
    }
}
```

Check the MPLS configuration on the top-of-rack switch Coleman by using the **show mpls protocol** command.

```
user@coleman# show mpls protocol
interfaces {
    ge-2/0/0 {
        unit 0 {
            family inet {
                address 2.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-2/0/1 {
        unit 0 {
            family inet {
                address 2.1.2.2/24;
            }
        }
    }
    ge-2/0/5 {
        unit 0 {
            family inet {
                address 3.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-2/0/6 {
        unit 0 {
            family inet {
                address 3.1.2.1/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
        }
    }
    router-id 10.9.82.46;
    autonomous-system 100;
    forwarding-table {
        export pplb;
    }
}
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
```



```
    }
    bgp {
        group ibgp-ToR1 {
            type internal;
            accept-remote-nexthop;
            local-address 10.9.82.46;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export [ advToR2 advLo ];
            neighbor 10.9.82.49;
        }
        group ebgp-R2 {
            type external;
            local-address 3.1.2.1;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            peer-as 102;
            local-as 100;
            neighbor 3.1.2.2 {
                as-override;
            }
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
policy-options {
    policy-statement advLo {
        term lo0 {
            from {
                protocol direct;
                route-filter 10.9.82.46/32 exact;
            }
            then {
                next-hop 2.1.1.2;
                accept;
            }
        }
        term other {
            then reject;
        }
    }
    policy-statement advToR2 {
        term ToR2lo {
```

```
            from {
                protocol bgp;
                route-filter 10.9.82.47/32 exact;
            }
            then accept;
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
            accept;
        }
    }
}
```

Check the MPLS configuration on the router named Taxes by using the **show mpls protocols** command:

```
user@taxes# show mpls protocol
interfaces {
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 3.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-1/0/0 {
        unit 0 {
            family inet {
                address 3.1.2.2/24;
            }
            family mpls;
        }
    }
    ge-1/0/2 {
        unit 0 {
            family inet {
                address 4.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-1/2/0 {
        unit 0 {
            family inet {
                address 4.1.2.1/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
```



```
            }
        }
        router-id 10.9.82.48;
        autonomous-system 102;
        forwarding-table {
            export pplb;
        }
    }
    protocols {
        mpls {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
            group ebgp-R1 {
                type external;
                local-address 3.1.2.2;
                family inet {
                    labeled-unicast {
                        rib {
                            inet.3;
                        }
                    }
                }
                peer-as 100;
                local-as 102;
                neighbor 3.1.2.1 {
                    as-override;
                }
            }
            group ibgp-ToR2 {
                type internal;
                local-address 10.9.82.48;
                family inet {
                    labeled-unicast {
                        rib {
                            inet.3;
                        }
                    }
                }
                export [ advToR2 advLo ];
                neighbor 10.9.82.47;
            }
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface all;
                interface fxp0.0 {
                    disable;
                }
            }
        }
    }
    policy-options {
        policy-statement advLo {
            term lo0 {
                from {
                    protocol direct;
```

```
                route-filter 10.9.82.48/32 exact;
            }
            then {
                next-hop 4.1.1.1;
                accept;
            }
        }
        term other {
            then reject;
        }
    }
    policy-statement advToR2 {
        term ToR2lo {
            from {
                protocol bgp;
                route-filter 10.9.82.47/32 exact;
            }
            then accept;
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
            accept;
        }
    }
}
```

Check the MPLS configuration on the router named Dutchmen by using the **show mpls protocols** command:

```
user@dutchmen# show mpls protocol
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 5.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-0/1/0 {
        unit 0 {
            family inet {
                address 4.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-0/2/0 {
        unit 0 {
            family inet {
                address 4.1.2.2/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
```



```
        }
        rib-groups {
            inet0toinet3 {
                import-rib [ inet.0 inet.3 ];
            }
        }
        router-id 10.9.82.47;
        autonomous-system 102;
        forwarding-table {
            export pplb;
        }
    }
    protocols {
        mpls {
            static-label-switched-path to-server {
                transit 1000001 {
                    next-hop 5.1.1.2;
                    pop;
                }
            }
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
            group ibgp-R2 {
                type internal;
                local-address 10.9.82.47;
                family inet {
                    labeled-unicast {
                        rib {
                            inet.3;
                        }
                    }
                }
                export advLo;
                neighbor 10.9.82.48;
            }
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface all;
                interface fxp0.0 {
                    disable;
                }
            }
        }
    }
    policy-options {
        policy-statement advLo {
            term lo0 {
                from {
                    protocol direct;
                    route-filter 10.9.82.47/32 exact;
                }
                then {
                    next-hop 4.1.1.2;
                    accept;
                }
```

```
        }
        term other {
            then reject;
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
            accept;
        }
    }
}
```

Check the MPLS configuration on the top-of-rack switch Coleman by using the **show mpls command**.

```
user@coleman# show mpls protocol
interfaces {
    ge-2/0/0 {
        unit 0 {
            family inet {
                address 2.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-2/0/1 {
        unit 0 {
            family inet {
                address 2.1.2.2/24;
            }
        }
    }
    ge-2/0/5 {
        unit 0 {
            family inet {
                address 3.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-2/0/6 {
        unit 0 {
            family inet {
                address 3.1.2.1/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
        }
    }
    router-id 10.9.82.46;
    autonomous-system 100;
```



```
    forwarding-table {
        export pplb;
    }
}
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group ibgp-ToR1 {
            type internal;
            accept-remote-nexthop;
            local-address 10.9.82.46;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export [ advToR2 advLo ];
            neighbor 10.9.82.49;
        }
        group ebgp-R2 {
            type external;
            local-address 3.1.2.1;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            peer-as 102;
            local-as 100;
            neighbor 3.1.2.2 {
                as-override;
            }
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
policy-options {
    policy-statement advLo {
        term lo0 {
            from {
                protocol direct;
                route-filter 10.9.82.46/32 exact;
            }
            then {
```

```
                    next-hop 2.1.1.2;
                    accept;
                }
            }
            term other {
                then reject;
            }
        }
        policy-statement advToR2 {
            term ToR2lo {
                from {
                    protocol bgp;
                    route-filter 10.9.82.47/32 exact;
                }
                then accept;
            }
        }
        policy-statement pplb {
            then {
                load-balance per-packet;
                accept;
            }
        }
    }
```

Check the MPLS configuration on the router Taxes by using the **show mpls protocols** command:

```
user@taxes# show mpls protocol
interfaces {
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 3.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-1/0/0 {
        unit 0 {
            family inet {
                address 3.1.2.2/24;
            }
            family mpls;
        }
    }
    ge-1/0/2 {
        unit 0 {
            family inet {
                address 4.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-1/2/0 {
        unit 0 {
            family inet {
                address 4.1.2.1/24;
            }
            family mpls;
```



```
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
        }
    }
    router-id 10.9.82.48;
    autonomous-system 102;
    forwarding-table {
        export pplb;
    }
}
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group ebgp-R1 {
            type external;
            local-address 3.1.2.2;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            peer-as 100;
            local-as 102;
            neighbor 3.1.2.1 {
                as-override;
            }
        }
        group ibgp-ToR2 {
            type internal;
            local-address 10.9.82.48;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export [ advToR2 advLo ];
            neighbor 10.9.82.47;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
```

```
                disable;
            }
        }
    }
}
policy-options {
    policy-statement advLo {
        term lo0 {
            from {
                protocol direct;
                route-filter 10.9.82.48/32 exact;
            }
            then {
                next-hop 4.1.1.1;
                accept;
            }
        }
        term other {
            then reject;
        }
    }
    policy-statement advToR2 {
        term ToR2lo {
            from {
                protocol bgp;
                route-filter 10.9.82.47/32 exact;
            }
            then accept;
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
            accept;
        }
    }
}
```

Check the MPLS configuration on the router Dutchmen by using the **show mpls protocols** command:

```
show mpls protocol??
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 5.1.1.1/24;
            }
            family mpls;
        }
    }
    ge-0/1/0 {
        unit 0 {
            family inet {
                address 4.1.1.2/24;
            }
            family mpls;
        }
    }
    ge-0/2/0 {
        unit 0 {
```



```
            family inet {
                address 4.1.2.2/24;
            }
            family mpls;
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet inet0toinet3;
    }
    rib-groups {
        inet0toinet3 {
            import-rib [ inet.0 inet.3 ];
        }
    }
    router-id 10.9.82.47;
    autonomous-system 102;
    forwarding-table {
        export pplb;
    }
}
protocols {
    mpls {
        static-label-switched-path to-server {
            transit 1000001 {
                next-hop 5.1.1.2;
                pop;
            }
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group ibgp-R2 {
            type internal;
            local-address 10.9.82.47;
            family inet {
                labeled-unicast {
                    rib {
                        inet.3;
                    }
                }
            }
            export advLo;
            neighbor 10.9.82.48;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
policy-options {
```



```
        policy-statement advLo {
            term lo0 {
                from {
                    protocol direct;
                    route-filter 10.9.82.47/32 exact;
                }
                then {
                    next-hop 4.1.1.2;
                    accept;
                }
            }
            term other {
                then reject;
            }
        }
        policy-statement pplb {
            then {
                load-balance per-packet;
                accept;
            }
        }
    }
```

Related Documentation

- MPLS Stitching For Virtual Machine Connection on page 161

Layer 2 Circuit Overview

A Layer 2 circuit is a point-to-point Layer 2 connection transported using Multiprotocol Label Switching (MPLS) or other tunneling technology on the service provider's network. A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple virtual circuits (VCs) are transported over a single shared label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a separate dedicated LSP.

To establish a Layer 2 circuit, the Link Integrity Protocol (LIP) is used as the signaling protocol to advertise the ingress label to the remote PE routers. For this purpose, a targeted remote LDP neighbor session is established using the extended discovery mechanism described in LDP, and the session is brought up on the remote PE loopback IP address. Because LDP looks at the Layer 2 circuit configuration and initiates extended neighbor discovery for all the Layer 2 circuit neighbors (the remote PEs), no new configuration is necessary in LDP. Each Layer 2 circuit is represented by the logical interface connecting the local PE router to the local customer edge (CE) router. Note that LDP must be enabled on the lo0.0 interface for extended neighbor discovery to function correctly.

Packets are sent to remote CE routers over an egress VPN label advertised by the remote PE router, using a targeted LDP session. The VPN label is sent over an LDP LSP to the remote PE router connected to the remote CE router. Return traffic from the remote CE router destined to the local CE router is sent using an ingress VPN label advertised by the local PE router, which is also sent over the LDP LSP to the local PE router from the remote PE router.

Related Documentation

- *Layer 3 VPN Overview*
- *Layer 2 VPN Overview*
- *Layer 2 VPN Applications*
- *Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit*
- *Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN*
- *Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit*
- *Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN*
- *Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN*

CHAPTER 5

Configuration Statements for MPLS

adaptive

Syntax adaptive;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.

Default The configured object is disabled.

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

Related Documentation
- *Configuring Adaptive LSPs*

adjust-interval

Syntax adjust-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Specify the bandwidth reallocation interval.

Options ***seconds***—Bandwidth reallocation interval, in seconds.

Range: 300 through 315,360,000 seconds

Default: 86,400 seconds

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

adjust-threshold

Syntax adjust-threshold *percent*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.

Options ***percent***—Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

adjust-threshold-overflow-limit

Syntax adjust-threshold-overflow-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced in Junos OS Release 7.5.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.

Options *number*—Number of consecutive bandwidth overflow samples.

Range: 1 through 65,535

Default: This feature is disabled by default.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

adjust-threshold-underflow-limit

Syntax adjust-threshold-underflow-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced in Junos OS Release 11.3.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.

Options *number*—Number of consecutive bandwidth underflow samples.

Range: 1 through 65,535

Default: This feature is disabled by default.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

admin-down

Syntax admin-down;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.

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

Related Documentation

- *Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the Junos OS*

advertisement-hold-time

Syntax advertisement-hold-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.

Options ***seconds***—Hold time, in seconds.

Range: 0 through 65,535 seconds

Default: 5 seconds

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

Related Documentation

- *Damping Advertisement of LSP State Changes*

associate-backup-pe-groups

Syntax associate-backup-pe-groups;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Enable an LSP to monitor the status of its destination PE router. You can configure multiple backup PE router groups using the same router's address. Backup PE router groups provide ingress PE router redundancy when point-to-multipoint LSPs are configured for multicast distribution. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. This statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to the destination address.

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

Related Documentation

- *Enabling Point-to-Point LSPs to Monitor Egress PE Routers*

auto-bandwidth

Syntax

```
auto-bandwidth {
  adjust-interval seconds;
  adjust-threshold percent;
  adjust-threshold-activate-bandwidthbps
  adjust-threshold-overflow-limit number;
  adjust-threshold-underflow-limit number;
  maximum-bandwidth bps;
  minimum-bandwidth bps;
  minimum-bandwidth-adjust-interval
  minimum-bandwidth-adjust-threshold-change
  minimum-bandwidth-adjust-threshold-value
  monitor-bandwidth;
}
```

Hierarchy Level

[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Automatic Bandwidth Allocation for LSPs*
- request mpls lsp adjust-autobandwidth on page 340

backup-pe-group

Syntax

backup-pe-group *group-name* {
  backups [*addresses*];
  local-address *address*;
}

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

backups *addresses*—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

local-address *address*—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

pe-group-name—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.

Required Privilege Level

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

Related Documentation

- *Example: Configuring Ingress PE Redundancy*
- *Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs*

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

Syntax

bandwidth *bps* {
  ct0 *bps*;
  ct1 *bps*;
  ct2 *bps*;
  ct3 *bps*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description

When configuring an LSP, specify the traffic rate associated with the LSP.

When configuring fast reroute, allocate bandwidth for the reroute path. By default, no bandwidth is reserved for the rerouted path. The fast reroute bandwidth does not need to be identical to that allocated for the LSP itself.

When configuring a multiclass LSP, use the **ct*number* *bandwidth*** statements to specify the bandwidth to be allocated for each class type.

Options

bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

NOTE: On the ACX Series, *bps* is the only supported option.

ct*number* ***bps***—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion [also called a thousand million]).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

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

Related Documentation

- *Configuring Fast Reroute*
- *Configuring the Bandwidth Value for LSPs*
- *Configuring LSPs for DiffServ-Aware Traffic Engineering*
- *Configuring Multiclass LSPs*

bandwidth-model

Syntax

```
bandwidth-model {
  extended-mam;
  mam;
  rdm;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.

Options **extended-mam**—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.

mam—The MAM is defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

rdm—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.

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

Related Documentation

- *Configuring Routers for DiffServ-Aware Traffic Engineering*

bypass (Static LSP)

Syntax

```
bypass bypass-name {
  bandwidth bps;
  description string;
  next-hop (address | interface-name | address/interface-name);
  push out-label;
  to address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 10.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description

Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Static LSPs*

class-of-service (Protocols MPLS)

Syntax class-of-service *cos-value*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Class-of-service (CoS) value given to all packets in the LSP.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options *cos-value*—CoS value. A higher value typically corresponds to a higher level of service.

Range: 0 through 7

Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

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

Related Documentation

- *Configuring Class of Service for MPLS LSPs*
- *Configuring the Ingress Router for Static LSPs*
- *Configuring Static LSPs*

corouted-bidirectional

Syntax corouted-bidirectional;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*], [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the **corouted-bidirectional-passive** statement.

Default This statement is disabled by default.

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

Related Documentation

- *Configuring Corouted Bidirectional LSPs*
- corouted-bidirectional-passive on page 258

corouted-bidirectional-passive

Syntax corouted-bidirectional-passive;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*], [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX VC/VCF Series.

Description Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the **corouted-bidirectional** statement.

Default This statement is disabled by default.

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

Related Documentation

- *Configuring Corouted Bidirectional LSPs*
- corouted-bidirectional on page 258

description (Protocols MPLS)

Syntax description *text*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* bypass],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* ingress],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* transit *incoming-label*],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* bypass],
[edit protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show mpls lsp detail** command and has no effect on the operation of the LSP.

Options *text*—Provide a textual description of the LSP. The description text can be no more than 80 characters in length.

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

Related Documentation

- *Configuring a Text Description for LSPs*

diffserv-te

Syntax

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
  }
  te-class-matrix {
    tenumber {
      priority priority;
      traffic-class {
        ctnumber priority priority;
      }
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify properties for differentiated services in traffic engineering.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Routers for DiffServ-Aware Traffic Engineering*

disable (Protocols MPLS)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls interface *interface-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls],
[edit protocols mpls interface *interface-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Disable the functionality of the configured object.

Default The configured object is enabled (operational) unless explicitly disabled.

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

Related Documentation

- *Minimum MPLS Configuration*

exclude (for Fast Reroute)

Syntax (exclude [*group-names*] | no-exclude);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX3500 and QFX 5100 in Junos OS Release 14.1X53-D10.

Description Control exclusion of administrative groups:

- **exclude**—Define the administrative groups to exclude for fast reroute.
- **no-exclude**—Disable administrative group exclusion.

Options ***group-names***—Names of one or more groups defined with the **admin-groups** statement.

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

Related Documentation

- *Configuring Fast Reroute*
- *admin-groups*

exclude-srlg

Syntax exclude-srlg;

Hierarchy Level [edit protocols mpls],
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls label-switched-path *path-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *destination*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *destination*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Exclude Shared Risk Link Group (SRLG) links for the secondary path for critical links where it is imperative to keep the secondary and primary label-switched paths completely disjoint from any common SRLG.

When specified, the Constrained Shortest Path First (CSPF) algorithm excludes any link belonging to the set of SRLGs in the primary path. When not specified and if a link belongs to the set of SRLGs in the primary path, CSPF adds the SRLG cost to the metric, but still accepts the link for computing the path.

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

Related Documentation

- *Example: Excluding SRLG Links Completely for the Secondary LSP*

explicit-null (Protocols MPLS)

Syntax explicit-null;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Advertise label 0 to the egress router of an LSP.

Default If you do not include the **explicit-null** statement in the MPLS configuration, label 3 (implicit null) is advertised.

NOTE: Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces.

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

Related Documentation

- *Configuring RSVP to Pop the Label on the Ultimate-Hop Router*

fast-reroute (Protocols MPLS)

Syntax

```
fast-reroute {
  (bandwidth bps | bandwidth-percent percentage);
  (exclude  [ group-names ] | no-exclude );
  hop-limit number;
  (include-all [ group-names ] | no-include-all);
  (include-any [ group-names ] | no-include-any);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Fast Reroute*
- Fast Reroute Overview on page 166
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Interprovider and Carrier-of-Carriers VPNs on page 166

forwarding-table

Syntax

```
forwarding-table {
  chained-composite-next-hop;
  export [ policy-name ];
  (indirect-next-hop | no-indirect-next-hop);
  (indirect-next-hop-change-acknowledgements |
    no-indirect-next-hop-change-acknowledgements;)
  krt-nexthop-ack-timeout interval;
  unicast-reverse-path (active-paths | feasible-paths);
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series.

Description

Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Example: Load Balancing BGP Traffic*

from (Protocols MPLS)

Syntax from *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*], [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the source address to use for the LSP.

The address you specify does not affect the outgoing interface used by the LSP.

Default If you do not include this statement, the software automatically selects the loopback interface as the address.

Options *address*—IP address.

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

Related Documentation

- *Configuring the Ingress and Egress Router Addresses for LSPs*

gpid

Syntax

gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* lsp-attributes],
[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]

Release Information

Statement introduced before Junos OS Release 7.4.
pos-scrambling-crc-16, **pos-no-scrambling-crc-16**, **pos-scrambling-crc-32**, and **pos-no-scrambling-crc-32** options added in Junos OS Release 8.0.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the type of payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet (GPID value: 33)
- **hdlc**—High-level Data Link Control (HDLC) (GPID value: 44)
- **ipv4**—IP version 4 (GPID value: 0x0800)
- **pos-no-scrambling-crc-16**—for interoperability with other vendors' equipment (GPID value: 29)
- **pos-no-scrambling-crc-32**—for interoperability with other vendors' equipment (GPID value: 30)
- **pos-scrambling-crc-16**—for interoperability with other vendors' equipment (GPID value: 31)
- **pos-scrambling-crc-32**—for interoperability with other vendors' equipment (GPID value: 32)
- **ppp**—Point-to-Point Protocol (PPP) (GPID value: 50)

Default

ipv4

Required Privilege Level

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

Related Documentation

- *Configuring MPLS LSPs for GMPLS*

hop-limit

Syntax

hop-limit *number*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:

- LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.
- Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
- Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.

Options

number—Maximum number of hops.

Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)

Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)

Required Privilege Level

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

Related Documentation

- *Configuring Fast Reroute*
- *Limiting the Number of Hops in LSPs*

- *Configuring Link Protection on Interfaces Used by LSPs*

include-all (for Fast Reroute)

Syntax (include-all [*group-names*] | no-include-all);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description Control inclusion of administrative groups:

- **include-all**—Define the administrative groups that must all be included for fast reroute.
- **no-include-all**—Disable administrative group inclusion.

Options *group-names*—One or more names of groups defined with the **admin-groups** statement.

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

Related Documentation

- *Configuring Fast Reroute*

include-any (for Fast Reroute)

Syntax (include-any [*group-names*] | no-include-any);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description Control inclusion of administrative groups:

- **include-any**—Define the administrative groups to include for fast reroute.
- **no-include-any**—Disable administrative group inclusion.

Options *group-names*—One or more names of groups defined with the **admin-groups** statement.

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

Related Documentation

- *Configuring Fast Reroute*

ingress (LSP)

Syntax

```
ingress {
  bandwidth bps;
  class-of-service cos-value;
  description string;
  entropy-label;
  install {
    destination-prefix <active>;
  }
  link-protection bypass-name name;
  metric metric;
  next-hop (address | interface-name | address/interface-name);
  node-protection bypass-name name next-next-label label;
  no-install-to-address;
  policing {
    filter filter-name;
    no-auto-policing;
  }
  preference preference;
  push out-label;
  to address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name*]

Release Information

Statement introduced in Junos OS Release 10.1.
entropy-label option introduced in Junos OS Release 14.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Configure an ingress LSR for a static LSP.

The remaining statements are explained separately

Required Privilege Level

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

Related Documentation

- *Configuring Static LSPs*

install (Protocols MPLS)

Syntax

install {
  destination-prefix <active>;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 or inet6.3 routing table.

Options

active—(Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a **ping** or **traceroute** command on this address.

destination-prefix—IPv4 or IPv6 address to associate with the LSP.

Required Privilege Level

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

Related Documentation

- *Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table*

interface (Protocols MPLS)

Syntax

interface (*interface-name* | all) {
  disable;
  admin-group [*group-names*];
  srlg *srlg-name*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Enable MPLS on one or more interfaces.

Options

interface-name—Name of the interface on which to configure MPLS. To configure all interfaces, specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

srlg *srlg-name*—Name of the SRLG to associate with an interface.

The remaining options are explained separately.

Required Privilege Level

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

Related Documentation

- *Minimum MPLS Configuration*
- *Configuring Static LSPs*
- *Example: Configuring SRLG*

ipv6-tunneling

Syntax ipv6-tunneling;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.

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

Related Documentation
- *Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks*

ldp-tunneling

Syntax ldp-tunneling;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric

Description Enable the LSP to be used for LDP tunneling.

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

Related Documentation
- Enabling LDP over RSVP-Established LSPs on page 34

link-protection (Static LSPs)

Syntax link-protection bypass-name *name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* ingress],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* transit *incoming-label*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable link protection on the specified static LSP. Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

Default Link protection is disabled.

Options **bypass-name *name***—Bypass LSP name.

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

Related Documentation

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

log-updown (Protocols MPLS)

Syntax

```
log-updown {
  no-trap {
    mpls-lsp-traps;
    rfc3812-traps;
  }
  (syslog | no-syslog);
  trap;
  trap-path-down;
  trap-path-up;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
The **mpls-lsp-traps** and **rfc-3812-traps** options added in Junos OS Release 9.0.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.

NOTE: System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the **no-syslog** option under the **log-updown** statement.

Default

There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.

Options

no-syslog—Do not log a message to the system log file.

no-trap—Do not send an SNMP trap.

syslog—Log a message to the system log file.

trap—Send an SNMP trap.

trap-path-down—Send an SNMP trap when an LSP path goes down.

trap-path-up—Send an SNMP trap when an LSP path comes up.

The **no-trap** statement is explained separately.

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

Related Documentation
- *Configuring System Log Messages and SNMP Traps for LSPs*
- *Network Management Administration Guide for Routing Devices*
- no-trap on page 285
- traceoptions (Protocols MPLS) on page 319

lsp-attributes

Syntax

```
lsp-attributes {
  encoding-type (ethernet | packet | pdh | sonet-sdh);
  gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
    pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
  signal-bandwidth type;
  switching-type (fiber | lambda | psc-1 | tdm);
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
pos-scrambling-crc-16, **pos-no-scrambling-crc-16**, **pos-scrambling-crc-32**, and **pos-no-scrambling-crc-32** options added in Junos OS Release 8.0.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.

The options are explained separately.

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

Related Documentation
- *Configuring MPLS LSPs for GMPLS*

maximum-bandwidth (Protocols MPLS)

Syntax maximum-bandwidth *bps*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Specify the maximum amount of bandwidth in bits per second (bps).

Options ***bps***—Maximum amount of bandwidth.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

metric (Protocols MPLS)

Syntax metric *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.

Options ***metric***—LSP metric value.

Default: No metric assigned (dynamic)

Range: 1 through 16,777,215

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

Related Documentation
- *Configuring LSP Metrics*

minimum-bandwidth

Syntax minimum-bandwidth *bps*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.

Options *bps*—Miniminum bandwidth for the LSP.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

monitor-bandwidth

Syntax monitor-bandwidth;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth],
[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.

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

Related Documentation
- *Configuring Automatic Bandwidth Allocation for LSPs*

mtu-signaling

Syntax mtu-signaling;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls path-mtu rsvp],
[edit protocols mpls path-mtu rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Enable MTU signaling in RSVP.

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

Related Documentation

- *Configuring MTU Signaling in RSVP*

no-cspf

Syntax

no-cspf;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mplslabel-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable constrained-path LSP computation.

An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.

A constrained-path LSP relies on an ingress router to compute the complete path. The ingress router takes into account the following information during the computation:

- Interior gateway protocol (IGP) topology database
- Link utilization information from extensions in the IGP link-state database
- Administrative group information from extensions in the IGP link-state database
- LSP requirements, including bandwidth, hop count, and administrative group

Constrained-path LSPs can generally avoid link failures and congested links. They also permit recomputation (therefore, a new path) during topology changes or unsuccessful setup.

Default

Constrained-path LSP computation enabled.

Required Privilege Level

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

Related Documentation

- *Disabling Constrained-Path LSP Computation*
- *Configuring Explicit-Path LSPs*

no-decrement-ttl

Syntax no-decrement-ttl;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Disable normal time-to-live (TTL) decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.

Default Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

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

Related Documentation

- *Disabling Normal TTL Decrementing*
- no-propagate-ttl on page 284

no-install-to-address

Syntax no-install-to-address;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Prevent the egress router address configured using the **to** statement from being installed into the inet.3 and inet.0 routing tables.

Default The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.

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

Related Documentation

- *Preventing the Addition of Egress Router Addresses to Routing Tables*
- to on page 318

no-propagate-ttl

Syntax no-propagate-ttl;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Disable normal time-to-live (TTL) decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.

When you add the **no-propagate-ttl** statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.

Default Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

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

Related Documentation

- *Disabling Normal TTL Decrementing*
- *Example: Disabling Normal TTL Decrementing in a VRF Routing Instance* (on *Layer 3 VPNs Feature Guide for Routing Devices* or in the *Junos VPNs Configuration Guide*)
- no-decrement-ttl on page 282

no-trap

Syntax

```
no-trap {
  mpls-lsp-traps;
  rfc-3812-traps;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls log-updown],
[edit protocols mpls log-updown]

Release Information

Statement introduced before Junos OS Release 7.4.
The **mpls-lsp-traps** and **rfc-3812-traps** options added in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Prevent the transmission of SNMP traps.

Options

mpls-lsp-traps—Block the MPLS LSP traps defined in the**rfc-3812-traps**, but allows the **rfc3812.mib** traps.

rfc-3812-traps—Block the traps defined in the **rfc3812.mib**, but allows the MPLS LSP traps defined in the **jnx-mpls.mib**.

Required Privilege Level

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

Related Documentation

- *Configuring System Log Messages and SNMP Traps for LSPs*
- *Network Management Administration Guide for Routing Devices*
- traceoptions (Protocols MPLS) on page 319

node-link-protection (Protocols MPLS)

Syntax node-link-protection;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the **link-protection** statement at the **[edit protocols rsvp interface *interface-name*]** hierarchy level.

Default Node and link protection is disabled.

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

Related Documentation

- *Configuring Node Protection or Link Protection for LSPs*
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Interprovider and Carrier-of-Carriers VPNs on page 166

oam (Protocols MPLS)

Syntax

```
oam {
  bfd-liveness-detection{
    failure-action teardown;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
  }
  lsp-ping-interval seconds;
  mpls-tp-mode;
}
```

Hierarchy Level

[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*]
[edit protocols mpls label-switched-path *lsp-name* primary *path-name*]

Release Information

Statement introduced in Junos OS Release 7.6.
lsp-ping-interval option introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.

Options

lsp-ping-interval ***seconds***—Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the **ping mpls rsvp** command.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring BFD for MPLS IPv4 LSPs*

optimize-aggressive

Syntax optimize-aggressive;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default.

Default Aggressive optimization is disabled.

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

Related Documentation

- *Optimizing Signaled LSPs*

optimize-hold-dead-delay

Syntax optimized-hold-dead-delay *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switch-path *lsp-name*],
[edit protocols mpls],
[edit protocols mpls label-switch-path *lsp-name*]

Description Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths. This delay timer starts when the timer specified by the **optimize-switchover-dealy** statement has elapsed.

Options *seconds*—Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization.

Default: 60 seconds

Range: 0 through 65,535 seconds

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

Related Documentation

- *Optimizing Signaled LSPs*
- optimize-switchover-delay on page 290
- optimize-timer on page 291

optimize-switchover-delay

Syntax optimize-switchover-delay *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced in Junos OS Release 11.1R1.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths.

Options *seconds*—Configure the time in seconds to wait before switching LSPs to newly optimized paths.

Default: 1 second

Range: 1 through 900 seconds

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

Related Documentation

- *Optimizing Signaled LSPs*
- optimize-hold-dead-delay on page 289
- optimize-timer on page 291

optimize-timer (Protocols MPLS)

Syntax optimize-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path. This feature is useful only on LSPs for which constrained-path computation is enabled; that is, for which the **no-cspf** statement is not configured. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).

To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.

Default The optimize timer is disabled.

Options ***seconds***—Length of the optimize timer, in seconds.

Range: 0 through 65,535 seconds

Default: 0 seconds (the optimize timer is disabled)

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

Related Documentation

- *Optimizing Signaled LSPs*

p2mp (Protocols MPLS)

Syntax p2mp *p2mp-lsp-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.

Options ***p2mp-lsp-name***—Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

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

Related Documentation
- *Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs*

path (Protocols MPLS)

Syntax

path *path-name* {
  (*address* | *hostname*) <strict | loose>;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Create a named path and optionally specify the sequence of explicit routers that form the path.

You must include this statement when configuring explicit LSPs.

Options

address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is **loose**. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

hostname—See ***address***.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

loose—(Optional) Indicate that the next address in the **path** statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.

Default: **strict**

path-name—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

strict—(Optional) Indicate that the LSP must go to the next address specified in the **path** statement without traversing other nodes. This is the default.

Required Privilege Level

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

Related Documentation

- *Creating Named Paths*

path-mtu

Syntax

```
path-mtu {
  allow-fragmentation;
  rsvp {
    mtu-signaling;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring MTU Signaling in RSVP*

policing (Protocols MPLS)

Syntax

```
policing {
   filter filter-name;
   no-auto-policing;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls
   static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the policing filter for the LSP.

Options

filter *filter-name*—Specify the name of the policing filter.

no-auto-policing—Disable automatic policing on this LSP.

Required Privilege Level

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

Related Documentation

- *Configuring MPLS Firewall Filters and Policers*
- *auto-policing*

policy-statement

Syntax

```
policy-statement policy-name {
  term term-name {
    from {
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
inet-mdt option introduced in Junos OS Release 10.0R2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
route-target option introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement** ***policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

Options ***actions***—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family ***family-name***—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify**route-target**.

NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

policy ***subroutine-policy-name***—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policy names cannot take the form **__.*-internal__**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

prefix-list ***prefix-list-name***—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter ***prefix-list-name***—Name of a prefix list to evaluate using qualifiers; ***match-type*** is the type of match (see *Configuring Prefix List Filters*), and ***actions*** is the action to take if the prefixes match.

route-filter ***destination-prefix match-type <actions>***—(Optional) List of routes on which to perform an immediate match; ***destination-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***destination-prefix*** matches.

source-address-filter ***source-prefix match-type <actions>***—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. ***source-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***source-prefix*** matches.

term *term-name*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

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

Related Documentation
- *dynamic-db*

pop

Syntax
pop;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*],
[edit protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

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

Related Documentation
- *Configuring the Intermediate and Egress Routers for Static LSPs*
- swap on page 315

preference (Protocols MPLS)

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Preference for the route.

You can optionally configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for LSPs is lower (more preferred) than all learned routes except direct interface routes.

Options *preference*—Preference to assign to the route. A route with a lower preference value is preferred.

Range: 1 through 255

Default: 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs

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

Related Documentation

- *Configuring Preference Values for LSPs*
- *Configuring Static LSPs*
- *Configuring Static LSPs*

primary (Protocols MPLS)

Syntax

```
primary path-name {
  adaptive;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth bps;
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  no-decrement-ttl;
  optimize-timer seconds;
  preference preference;
  priority setup-priority reservation-priority;
  (record | no-record);
  select (manual | unconditional);
  standby;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the primary path to use for an LSP. You can configure only one primary path.

You can optionally specify preference, CoS, and bandwidth values for the primary path, which override any equivalent values that you configure for the LSP (at the **[edit mpls label-switched-path *lsp-name*]** hierarchy level).

Options

path-name—Name of a path that you created with the **path** statement.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Primary and Secondary LSPs*

push

Syntax

push *out-label*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* bypass],
[edit logical-systems *logical-system-name* protocols mpls
static-label-switched-path *lsp-name* ingress],
[edit protocols mpls static-label-switched-path *lsp-name* bypass],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Add a new label to the top of the label stack. This statement is used to configure static LSPs at ingress routers and to configure bypass LSPs for static LSPs.

Options

out-label—Manually assigned outgoing label value.

Range: 0 through 1,048,575.

Required Privilege Level

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

Related Documentation

- pop on page 299
- swap on page 315

- *Configuring Static LSPs*

record

Syntax

(record | no-record);

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.

Default

Record routes.

Required Privilege Level

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

Related Documentation

- *Disabling Path Route Recording*

retry-limit

Syntax retry-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name*],

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection.

Options *number*—Maximum number of tries to establish the primary path.

Range: 0 through 10,000

Default: 0 (The ingress node never stops trying to establish the primary path.)

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

Related Documentation

- *Configuring the Connection Between Ingress and Egress Routers*

revert-timer

Syntax

revert-timer *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
BFD behavior modified in Junos OS Release 9.0.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.

If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter.

If you have configured a value of 0 seconds for the **revert-timer** statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.

Options

seconds—Time in seconds.

Range: 0 through 65,535 seconds

Default: 60 seconds

Required Privilege Level

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

Related Documentation

- *Configuring Primary and Secondary LSPs*

rsvp-error-hold-time

Syntax rsvp-error-hold-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations. The more time you configure, the more time a source node (ingress of an RSVP LSP) can have to learn about the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes.

Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the database and the network.

Options *seconds*—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.

Range: 0 through 240 seconds

Default: 25 seconds

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

Related Documentation

- *Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages*

secondary (Protocols MPLS)

Syntax

```
secondary path-name {
  adaptive;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth bps;
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  no-decrement-ttl;
  optimize-timer seconds;
  preference preference;
  priority setup-priority reservation-priority;
  (record | no-record);
  retry-limit number;
  retry-timer seconds;
  select (manual | unconditional);
  standby;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*], [edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

You can specify secondary paths even if you have not specified any primary paths.

Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the **[edit mpls label-switched-path]** hierarchy level).

Options

path-name—Name of a path that you created with the **path** statement.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Primary and Secondary LSPs*

select

Syntax select (manual | unconditional);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the conditions under which the path is selected to carry traffic. The **manual** and **unconditional** options are mutually exclusive.

Options **manual**—The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors).

unconditional—The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).

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

Related Documentation
- *Configuring Primary and Secondary LSPs*

signal-bandwidth

Syntax signal-bandwidth *type*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* lsp-attributes],
[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the bandwidth encoding of the signal used for path computation and admission control.

Options ***type***—Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: **10gigether**, **ds1**, **ds3**, **e1**, **e3**, **ethernet**, **fastether**, **gigether**, **stm-1**, **stm-4**, **stm-16**, **stm-64**, **stm-256, sts-1**, **vt1-5**, or **vt2**.

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

Related Documentation
- *Configuring MPLS LSPs for GMPLS*

smart-optimize-timer

Syntax smart-optimize-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Enable the smart optimization timer. When you enable the smart optimization timer on a router, the Junos OS operates on the assumption that the original LSP path is preferable to any alternate or secondary path. When you enable the smart optimization timer and an LSP fails and its traffic is switched to an alternate path, the smart optimization timer starts and waits 3 minutes (this time is configurable). After 3 minutes have passed, the LSP is switched back to the original path. If the original path fails again and the LSP is switched to an alternate path again, the router waits 1 hour before attempting to switch the LSP back to its original path.

If you want to disable the smart optimizer, you can set it to zero. The **smart-optimize-timer** value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the optimize-timer, which is usually set to a high value. Some ISPs have the **optimize-timer** set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the **smart-optimize-timer** is disabled automatically, and the **optimize-timer** (regular path optimization) goes into effect. This is to protect against a flapping link being used.

Default The smart optimization timer is enabled by default.

Options ***seconds***—(Optional) Specify the number of seconds to wait before switching an LSP back to its original path. If you do not specify the number of seconds, the default value is used.

Range: 0 through 65,535 seconds

Default: 180 seconds

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

Related Documentation

- *Configuring the Smart Optimize Timer*
- *Optimizing Signaled LSPs*
- optimize-aggressive on page 288
- optimize-timer on page 291

standby

Syntax
standby;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Have the path remain up at all times to provide instant switchover if connectivity problems occur.

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

Related Documentation
- *Configuring Hot Standby of Secondary Paths*

static-label-switched-path

Syntax

```
static-label-switched-path lsp-name {
  bypass bypass-name {
    bandwidth bps;
    description string;
    next-hop (address | interface-name | address/interface-name);
    push out-label;
    to address;
  }
  ingress {
    bandwidth bps;
    class-of-service cos-value;
    description string;
    install {
      destination-prefix <active>;
    }
    link-protection bypass-name name;
    metric metric;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    no-install-to-address;
    policing {
      filter filter-name;
      no-auto-policing;
    }
    preference preference;
    push out-label;
    to address;
  }
  transit incoming-label {
    bandwidth bps;
    description string;
    link-protection bypass-name name;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    pop;
    swap out-label;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information

Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Configure a static LSP.

Options

lsp-name—Name of the path.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Static LSPs*

statistics (Protocols MPLS)

Syntax

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Enable MPLS statistics collection and reporting.

Options **file *filename***—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file mpls-stat in the /var/log directory.

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

Range: 2 or more

Default: 2 files

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

interval *seconds*—Interval at which to periodically collect statistics.

Range: 1 through 65,535

Default: 300 seconds

no-world-readable—(Optional) Prevent users from reading the log file.

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

Syntax: 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: 1 MB

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

world-readable—(Optional) Enable users to read the log file.

The other statements are explained separately.

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

Related Documentation

- Configuring MPLS to Gather Statistics on page 178
- *Configuring Automatic Bandwidth Allocation for LSPs*

swap

Syntax swap *out-label*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*],
[edit protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Remove the label at the top of the label stack and replace it with the specified label. Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. This statement is used to configure static LSPs at transit routers.

Options
out-label—Manually assigned outgoing label value.

Range: 0 through 1,048,575

Default: If you do not define the ***out-label*** option, the original label value remains unchanged.

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

Related Documentation

- pop on page 299
- push on page 302
- *Configuring Static LSPs*

switching-type

Syntax
switching-type (fiber | lambda | psc-1 | tdm);

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* lsp-attributes],
[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Specify the switching method for the LSP. The switching method can be one of the following values:

- **fiber**—Fiber switching
- **lambda**—Lambda switching
- **psc-1**—Packet switching
- **tdm**—Time-division multiplexing (TDM) switching

Default
psc-1

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

Related Documentation
- *Configuring MPLS LSPs for GMPLS*

te-class-matrix

Syntax

```
te-class-matrix {
  tenumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.

Default

The default traffic engineering class matrix is:

```
te-class-matrix {
  te0 traffic-class ct0 priority 7;
  te1 traffic-class ct1 priority 7;
  te2 traffic-class ct2 priority 7;
  te3 traffic-class ct3 priority 7;
  te4 traffic-class ct0 priority 0;
  te5 traffic-class ct1 priority 0;
  te6 traffic-class ct2 priority 0;
  te7 traffic-class ct3 priority 0;
}
```

If you define any of the traffic engineering classes, all the default values are dropped.

Options

ct*number*—Specify the number of the class type. It can be one of four values: **ct0**, **ct1**, **ct2**, or **ct3**.

priority *priority*—Specify the priority of the class type. It can be one of eight values from 0 through 7.

te*number*—Specify the number of the traffic engineering class. It can be one of eight values: **te0**, **te1**, **te2**, **te3**, **te4**, **te5**, **te6**, or **te7**. You must configure the traffic engineering classes in order, starting with **te0**.

traffic-class—Specify the traffic class for the traffic engineering class.

Required Privilege Level

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

Related Documentation

- *Configuring Routers for DiffServ-Aware Traffic Engineering*

to

Syntax

to *address*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* bypass],
[edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* bypass],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the egress router of a dynamic LSP.

Options

address—Address of the egress router.

Required Privilege Level

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

Related Documentation

- *Configuring the Ingress and Egress Router Addresses for LSPs*

traceoptions (Protocols MPLS)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Configure MPLS tracing options at the protocol level or for a label-switched path.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default MPLS protocol-level tracing options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-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.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MPLS Tracing Flags

- **all**—Trace all operations
- **autobw-state**—Automatic bandwidth events.
- **connection**—All circuit cross-connect (CCC) activity
- **connection-detail**—Detailed CCC activity
- **cspf**—CSPF computations

- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events
- **lsping**—Trace lsping packets and return codes
- **nsr-synchronization**—Trace NSR synchronization events
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail
- **state**—All LSP state transitions
- **static**—Trace static label-switched path
- **timer**—Timer usage

no-world-readable—(Optional) Allow only certain users to read 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 this 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.

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: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

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

Required Privilege Level

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

Related Documentation

- *Tracing MPLS and LSP Packets and Operations*

traffic-engineering (Protocols MPLS)

Syntax traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.

Default bgp

Options **bgp**—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.

bgp-igp—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.

bgp-igp-both-ribs—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.

mpls-forwarding—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.

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

Related Documentation

- *Configuring Traffic Engineering for LSPs*
- *Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)*

transit-lsp-association

Syntax

```
transit-lsp-association transit-association-lsp-group-name {
   from-1 address-of-associated-lsp-1;
   from-2 address-of-associated-lsp-2;
   lsp-name-1 name-of-associated-lsp-1;
   lsp-name-2 name-of-associated-lsp-2;
}
```

Hierarchy Level [edit protocols mpls]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Associate two label-swiched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.

Options ***transit-association-lsp-group-name***—Name of the transit association LSP group.

from-1 *address-of-associated-lsp-1*—Address of the first associated LSP.

from-2 *address-of-associated-lsp-2*—Address of the second associated LSP.

lsp-name-1 *name-of-associated-lsp-1*—Name of the first associated LSP.

lsp-name-2 *name-of-associated-lsp-1*—Name of the second associated LSP.

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

Related Documentation

- *Example: Configuring the MPLS Transport Profile for OAM*

CHAPTER 6

Monitoring Commands for MPLS

- show mpls static-lsp
- show ted database
- show ted link
- show ted protocol

clear mpls lsp

List of Syntax

Syntax on page 325
Syntax (EX and QFX Series Switches) on page 325

Syntax

clear mpls lsp
<autobandwidth>
<logical-system (all | *logical-system-name*)>
<name *name*>
<optimize | optimize-aggressive>
<path *regular-expression*>
<statistics>

Syntax (EX and QFX Series Switches)

clear mpls lsp
<autobandwidth>
<name *name*>
<optimize | optimize-aggressive>
<path *regular-expression*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.

CAUTION: **This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.**

Options

none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

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

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

optimize | optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

path ***regular-expression***—(Optional) Clear the specific LSP path matching the specified regular expression.

statistics—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (**name** and **path** options) on transit routers.

Required Privilege Level
clear

Related Documentation
- show mpls lsp on page 397
- show rsvp session on page 506

List of Sample Output
clear mpls lsp on page 326

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

Sample Output

clear mpls lsp

```
user@host> clear mpls lsp
```

monitor label-switched-path

Syntax

monitor label-switched-path *lsp-name*
<logical-system (*logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Logical system support introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

Options

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

lsp-name—Name of the LSP.

Additional Information

You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the c key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in Table 28 on page 327. The keys are not case-sensitive.

Table 28: Output Control Keys for the monitor label-switched-path Command

Key	Action
c	Clears the screen and refreshes the display for this LSP.
f	Freezes the display, preventing new information from being displayed.
l	Monitors a different LSP. After you type l, you can type the new LSP name.
n	Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router.
p	Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws, or restarts, the data display for this LSP.

Required Privilege Level trace

List of Sample Output monitor label-switched-path on page 329

Output Fields Table 29 on page 328 describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

Table 29: monitor label-switched-path Output Fields

Field Name	Field Description
(1)	Displays the following information: • *hostname*—Name of the router. • **Seconds**—Time elapsed since this display was started. • **Time**—Current local time.
(2)	**Delay**—Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.
(3)	Displays the following: • **To**—Destination address of the LSP. • **From**—Originating address of the LSP. • **State**—Current state of the LSP: **Up** or **Down**.
(4)	Displays the following: • **LSPName**—Name of the LSP. • **Type**—Type of LSP: **Ingress**, **Egress**, or **Transit**.
(5)	Displays the following: • **Label in**—Incoming label of the LSP. • **Label out**—Outgoing label of the LSP.
(6)	**Port number**—Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always **0**.
(7/8)	**Record route**—All intermediate and egress router addresses for this LSP.
(9/10/11)	Displays traffic statistics: • **Output packets**—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago. • **Output bytes**—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.
(12)	Displays any errors the router encountered while attempting to retrieve information on the LSP.
(13)	Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see Table 28 on page 327.

Sample Output

monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host                           Seconds: 112              Time: 15:32:22
(2)                                                           Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4)   LSPname: k, type: Ingress
(5)   Label in: -, Label out: 126000
(6)   Port number: sender 1, receiver 45583, protocol 0
(7)   Record Route: <self> 192.168.224.196
(8)     192.168.224.202 192.168.224.179
(9)   Traffic statistics:                                        Current delta
(10)    Output packets:                      0                             [0]
(11)    Output bytes:                        0                             [0]
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
    LSP='l'
```

ping mpls l2circuit

Syntax

ping mpls l2circuit (interface *interface-name* | virtual-circuit *virtual-circuit-id* neighbor *address*)
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
reply-mode (application-level-control-channel | ip-udp | no-reply)
<size *bytes*>
<source *source-address*>
<sweep>
<v1>

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
The **size** and **sweep** options were introduced in Junos OS Release 9.6.
The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a **ping mpls l2circuit** command. You can also issue this command within logical systems.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

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

reply-mode—(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel—Reply using an application level control channel.

ip-udp—Reply using an IPv4 or IPv6 UDP packet.

no-reply—Do not reply to the ping request.

NOTE: **The reply-mode option and its suboptions application-level-control-channel, ip-udp, and no-reply are also available in Junos OS Release 10.2R4 and 10.3R2.**

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

v1—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

virtual-circuit *virtual-circuit-id* neighbor *address*—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

Additional Information

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls l2circuit interface on page 332
ping mpls l2circuit virtual-circuit detail on page 332
ping mpls l2circuit interface <interface-name> reply-mode on page 332

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


ping mpls bgp

Syntax

ping mpls bgp *fec*
<bottom-label-ttl>
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.

Options

bottom-label-ttl—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.

count ***count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination ***address***—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp ***forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance ***routing-instance-name***—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

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

size ***bytes***—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source ***source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls bgp fec count on page 334

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

Sample Output

ping mpls bgp fec count

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls l3vpn

Syntax

ping mpls l3vpn prefix *prefix-name*
<*l3vpn-name*>
<bottom-label-ttl>
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
The **size** and **sweep** options were introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Check the operability of an MPLS Layer 3 virtual private network (VPN) connection. Press Ctrl+c to interrupt a **ping mpls l3vpn** command.

Options

bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.

count ***count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination ***address***—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp ***forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

l3vpn-name—(Optional) Layer 3 VPN name.

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

prefix ***prefix-name***—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.

size ***bytes***—(Optional) Size of the label-switched path (LSP) ping request packet (**96** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes. The echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the **ping mpls l3vpn** command does not work.

Required Privilege Level

network

List of Sample Output

ping mpls l3vpn on page 336
ping mpls l3vpn detail on page 336

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. When an echo reply is received with an error code, the packets are not counted in the received packets count, and are counted separately.

Sample Output

ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
```



```
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls lsp-end-point

Syntax

ping mpls lsp-end-point *prefix-name*
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.
The **size** and **sweep** options were introduced in Junos OS Release 9.6.
The **instance** option was introduced in Junos OS Release 10.0.

Description

Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options

count ***count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination ***address***—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp ***forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

instance ***routing-instance-name***—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.

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

prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.

size ***bytes***—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is **88** bytes. If the endpoint is RSVP-based, the minimum size of the packet is **100** bytes. The maximum size in either case is **65468** bytes.

source ***source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls lsp-end-point detail on page 339

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls lsp-end-point detail

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


request mpls lsp adjust-autobandwidth

List of Syntax

Syntax on page 340
Syntax (EX and QFX Series Switches) on page 340

Syntax

request mpls lsp adjust-autobandwidth
<logical-system (all | *logical-system-name*)>
<name *lsp-name*>

Syntax (EX and QFX Series Switches)

request mpls lsp adjust-autobandwidth
<name *lsp-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).

Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.

During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as *make-before-break*.

Options

none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.

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

name *lsp-name*—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.

Additional Information

For this command to work properly, the following conditions must exist:

- Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the **request mpls lsp adjust-autobandwidth** command.
- The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.

Required Privilege Level

clear, maintenance

Related Documentation

- auto-bandwidth on page 252
- *Configuring Automatic Bandwidth Allocation for LSPs*

List of Sample Output request mpls lsp adjust-auto-bandwidth on page 341

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

Sample Output

request mpls lsp adjust-auto-bandwidth

`user@host>` **request mpls lsp adjust-auto-bandwidth**

show security keychain

Syntax

show security keychain
<brief | detail>

Release Information

Command introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.

Options

none—Display information about authentication keychains.

brief | detail—(Optional) Display the specified level of output.

Required Privilege Level

view

List of Sample Output

show security keychain brief on page 344
show security keychain detail on page 344

Output Fields

Table 30 on page 342 describes the output fields for the **show security keychain** command. Output fields are listed in the approximate order in which they appear.

Table 30: show security keychain Output Fields

Field Name	Field Description	Level of Output
keychain	The name of the keychain in operation.	All levels
Active-ID Send	Number of routing protocols packets sent with the active key.	All levels
Active-ID Receive	Number of routing protocols packets received with the active key.	All levels
Next-ID Send	Number of routing protocols packets sent with the next key.	All levels
Next-ID Receive	Number of routing protocols packets received with the next key.	All levels
Transition	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
Tolerance	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels

Table 30: show security keychain Output Fields *(continued)*

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>Id</td><td>Identification number configured for the current key.</td><td>detail</td></tr>
<tr><td>Algorithm</td><td>Authentication algorithm configured for the current key.</td><td>detail</td></tr>
<tr><td>State</td><td>State of the current key.

The value can be:

• receive
• send
• send-receive

For the active key, the State can be send-receive, send, or receive. For keys that have a future start time, the State is inactive. Compare the State field to the Mode field.</td><td>detail</td></tr>
<tr><td>Option</td><td>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.

The values can be:

• basic—Based on RFC 5304.
• isis-enhanced—Based on RFC 5310.

The default value is basic. When you configure the isis-enhanced option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.

When you configure basic (or do not include the options statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.

Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</td><td>detail</td></tr>
<tr><td>Start-time</td><td>Time that the current key became active.</td><td>detail</td></tr>
</table>

Table 30: show security keychain Output Fields *(continued)*

Field Name	Field Description	Level of Output
Mode	Mode of each key (Informational only.) The value can be • **receive** • **send** • **send-receive** The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the **Mode** can be **send-receive**, **send**, or **receive**, regardless of the configured start-time. Compare the **Mode** field to the **State** field.	**detail**

Sample Output

show security keychain brief

```
user@host> show security keychain brief
keychain                   Active-ID       Next-ID       Transition  Tolerance
                         Send  Receive   Send  Receive
 hakr                      3      3         1      1         1d 23:58    3600
```

show security keychain detail

```
user@host> show security keychain detail
keychain                   Active-ID       Next-ID       Transition  Tolerance
                         Send  Receive   Send  Receive
 hakr                      3      3         1      1         1d 23:58    3600
  Id 3, Algorithm hmac-md5, State send-receive, Option basic
  Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
  Id 1, Algorithm hmac-md5, State inactive, Option basic
  Start-time Fri Aug 20 11:30:57 2010, Mode send-receive
```


show link-management

Syntax show link-management

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show link-management peer on page 349
- show link-management routing on page 351
- show link-management statistics on page 355
- show link-management te-link on page 357

List of Sample Output show link-management on page 348

Output Fields Table 31 on page 345 describes the output fields for the **show link-management** command. Output fields are listed in the approximate order in which they appear.

Table 31: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.

Table 31: show link-management Output Fields *(continued)*

Field Name	Field Description
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: **Up**, **Down**, or **Init**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET, Ethernet**, **Packet**, and **PDH**.
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.

Table 31: show link-management Output Fields *(continued)*

Field Name	Field Description
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

Sample Output

show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
 State: Up, Control address: 10.255.245.4
   CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
         24547        24547 Up              1027       1026
  TE links:
   pro4-ba

 TE link name: pro4-ba, State: Init
  Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
  PSC-1,
  Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:

  155.52Mbps,
  Available bandwidth: 155.52Mbps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-1/0/2    Up       21271          0     155.52Mbps   No
```


show link-management peer

Syntax

show link-management peer
<name *peer-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) peer link information.

Options

none—Display all peer link information.

name ***peer-name***—(Optional) Display information for the specified peer only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management routing on page 351
- show link-management statistics on page 355
- show link-management te-link on page 357

List of Sample Output

show link-management peer on page 350

Output Fields

Table 32 on page 349 describes the output fields for the **show link-management peer** command. Output fields are listed in the approximate order in which they appear.

Table 32: show link-management peer Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.

Table 32: show link-management peer Output Fields *(continued)*

Field Name	Field Description
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

Sample Output

show link-management peer

```
user@host> show link-management peer
Peer name: sonet, System identifier: 41448
  State: Up, Control address: 70.70.70.70
  Hello interval: 10000, Hello dead interval: 30000
    CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
           3265            0 ConfSnd            1          0 R
   TE links:
    to-sonet
```


show link-management routing

Syntax

show link-management routing
<peer <name *name*> | te-link <name *name*>>
<resource <name *name*>>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.

Options

none—Display all peer and traffic-engineered link information.

peer <name *name*>—(Optional) Display information for all peers or for the specified peer only.

resource <name *name*>—(Optional) Display information for all resources or for the specified resource only.

te-link <name *name*>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management peer on page 349
- show link-management statistics on page 355
- show link-management te-link on page 357

List of Sample Output

show link-management routing on page 354

Output Fields

Table 33 on page 351 describes the output fields for the **show link-management routing** command. Output fields are listed in the approximate order in which they appear.

Table 33: show link-management routing Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.

Table 33: show link-management routing Output Fields *(continued)*

Field Name	Field Description
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: **Up** or **Down**.
System Identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.

Table 33: show link-management routing Output Fields *(continued)*

Field Name	Field Description
Traffic parameters	• **Encoding**—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**. • **Switching**—Type of switching that can be performed on the traffic-engineered link: **PSC-1** and **Packet**. • **Granularity**—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always **unknown**.

Sample Output

show link-management routing

```
user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
 State: Up, Control address: (null)
   Control-channel                    State
   fe-0/1/0.0                         Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
 State: Up, Control address: (null)
   Control-channel                    State
   fe-0/1/2.0                         Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
 State: Down, Control address: (null)
   Control-channel                    State
   so-0/2/0.0

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
 State: Down, Control address: (null)
   Control-channel                    State
   so-0/2/1.0

...

 TE link name: __rpd:fe-0/1/0.0, State: Up
  Local identifier: 2147483649, Remote identifier: 0,
  Local address: 192.168.37.66, Remote address: 192.168.37.66,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:fe-0/1/2.0, State: Up
  Local identifier: 2147483650, Remote identifier: 0,
  Local address: 192.168.37.73, Remote address: 192.168.37.73,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:so-0/2/0.0, State: Down
  Local identifier: 2147483651, Remote identifier: 0,
  Local address: 192.168.37.82, Remote address: 192.168.37.95,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
  Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn  System identifier: 2147483652,
Total bandwidth: 0bps,  Traffic parameters: Encoding: Packet, Switching: Packet,
 Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up  System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps,  Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown
```


show link-management statistics

Syntax

show link-management statistics
<peer <name *name*>>

Release Information

Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display statistical information for Link Management Protocol (LMP) packets.

Options

none—Display information for all peers.

peer <name ***name*****>**—(Optional) Display information for all peers or for the specified peer only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management peer on page 349
- show link-management routing on page 351
- show link-management te-link on page 357

List of Sample Output

show link-management statistics on page 356

Output Fields

Table 34 on page 355 describes the output fields for the **show link-management statistics** command. Output fields are listed in the approximate order in which they appear.

Table 34: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.

Table 34: show link-management statistics Output Fields *(continued)*

Field Name	Field Description
Stale acknowledgments	Number of **configAck** and **LinkSummaryAck** packets received that have a stale message ID.
Stale negative acknowledgments	Number of **configNack** and **LinkSummaryNack** packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

show link-management statistics

```
user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
    Received packets
        Config: 1
        Hello: 2572
    Small packets: 0
    Wrong protocol version: 0
    Messages for unknown peer: 0
    Messages for bad state: 0
    Stale acknowledgments: 0
    Stale negative acknowledgments: 0
    Sent packets
        Config: 2
        ConfigAck: 1
        Hello: 2572
    Retransmitted packets
        Config: 1
```


show link-management te-link

Syntax

show link-management te-link
<brief | detail>
<name *name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

Options

none—Display information for all traffic-engineered links.

brief | detail—(Optional) Display the specified level of output.

name ***name***—(Optional) Display information for the specified traffic-engineered link only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management peer on page 349
- show link-management routing on page 351
- show link-management statistics on page 355

List of Sample Output

show link-management te-link on page 358

Output Fields

Table 35 on page 357 describes the output fields for the **show link-management te-link** command. Output fields are listed in the approximate order in which they appear.

Table 35: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down.**
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.

Table 35: show link-management te-link Output Fields *(continued)*

Field Name	Field Description
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, **Packet**, and **PDH**.
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

Sample Output

show link-management te-link

```
user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-bd    Dn       43077          0           0bps No

 TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
   Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
 Available bandwidth: 8Mbps
```



```
Name         State Local ID  Remote ID      Bandwidth Used  LSP-name
falsp-be     Up        43076          0        10Mbps  Yes  e2elsp-bf
```

show mpls call-admission-control

List of Syntax
Syntax on page 360
Syntax (EX Series Switches) on page 360

Syntax
show mpls call-admission-control
<logical-system (all | *logical-system-name*)>
<*lsp-name*>

Syntax (EX Series Switches)
show mpls call-admission-control
<*lsp-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.

Options
none—Display CAC information for all LSPs.

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

lsp-name—(Optional) Display CAC information for the specified LSP only.

Additional Information
The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.

Required Privilege Level
view

List of Sample Output
show mpls call-admission-control on page 361

Output Fields
Table 36 on page 360 describes the output fields for the **show mpls call-admission-control** command. Output fields are listed in the approximate order in which they appear.

Table 36: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at **ct0**) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.

Table 36: show mpls call-admission-control Output Fields *(continued)*

Field Name	Field Description
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

Sample Output

show mpls call-admission-control

```
user@host# show mpls call-admission-control

LSP name: pro1-be
 *Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
 *Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
 *Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
 *Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
 *Primary   prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby   sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

show mpls cspf

List of Syntax

Syntax on page 362
Syntax (EX Series Switches) on page 362

Syntax

show mpls cspf
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show mpls cspf

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

Options

none—Display MPLS CSFP statistics.

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

Required Privilege Level

view

List of Sample Output

show mpls cspf on page 363

Output Fields

Table 37 on page 362 describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

Table 37: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of aborted computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.
successful	Number of path computations that were successfully completed.

Table 37: show mpls cspf Output Fields *(continued)*

Field Name	Field Description
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

Sample Output

show mpls cspf

```
user@host> show mpls cspf
CSPF statistics
Queue length   current       maximum      dequeued
                     0             0             0
Paths            total    successful      no route     sys error         CSPFs
                     0             0             0             0             0
Time (secs)      total         CSPFs  avg per CSPF      % of rpd
              0.000000      0.000000      0.000000        0.0000
```

show mpls diffserv-te

List of Syntax

Syntax on page 364
Syntax (EX Series Switches) on page 364

Syntax

show mpls diffserve-te
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show mpls diffserve-te

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.

Options

none—Display DiffServ classes and priorities used by MPLS LSPs.

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

Required Privilege Level

view

List of Sample Output

show mpls diffserv-te on page 365

Output Fields

Table 38 on page 364 describes the output fields for the **show mpls diffserv-te** command. Output fields are listed in the approximate order in which they appear.

Table 38: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	MPLS class type that corresponds to the DiffServ traffic engineering class: • **ct0**—Best effort • **ct1**—Assured forwarding • **ct2**—Expedited forwarding • **ct3**—Network control
Priority	MPLS preemption priority for this class type, a value from **0** through **7**. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

Sample Output

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class     Traffic class       Priority
te0          ct0                 3
te1          ct1                 2
```

show route forwarding-table

Syntax

show route forwarding-table
<detail | extensive | summary>
<ccc *ccc-interface-name*>
<*destination*>
<family *family-name*>
<label *label*>
<matching *ip_prefix*>
<multicast>
<vpn *vpn*>

Release Information

Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

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.

Options

none—Display the routes in the forwarding table.

detail | extensive | summary—(Optional) Display the specified level of output.

ccc—(Optional) Display the specified circuit cross-connect interface name for entries to match.

destination—(Optional) Display the destination prefix.

family *family-name*—(Optional) Display routing table entries for the specified family: **ethernet-switching**, **inet**, **inet6**, **iso**, **mpls**, **vlan classification**.

label *label*—(Optional) Display route entries for the specified label name.

matching *ip_prefix*—(Optional) Display route entries for the specified IP prefix.

multicast—(Optional) Display route entries for multicast routes.

vpn *vpn*—(Optional) Display route entries for the specified VPN.

Required Privilege Level

view

Related Documentation

- *Example: Configuring MPLS on EX8200 and EX4500 Switches*
- *Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)*

List of Sample Output

show route forwarding-table on page 368
show route forwarding-table summary on page 369

show route forwarding-table extensive on page 369
show route forwarding-table ccc on page 371
show route forwarding-table family (MPLS) on page 371
show route forwarding-table family (IPv6) on page 372
show route forwarding-table label on page 372
show route forwarding-table matching on page 372
show route forwarding-table multicast on page 372

Output Fields

Table 39 on page 367 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 39: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table (for example, **inet**, **inet6**, **mpls**).	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**).	All levels
Destination	Destination of the route.	**detail, extensive**
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **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**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface *interface-number*** —Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**

Table 39: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Nexthop	IP address of the next hop to the destination.	**detail, extensive**
Next hop type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)** —Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the 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.	**detail, extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail, extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	none **detail, extensive**
Next-hop interface (Netif)	Interface used to reach the next hop.	none **detail, 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 Interfaces	The next hop layer 2 interfaces. Seen with **multicast** option only.	**extensive**

Sample Output

show route forwarding-table

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


```
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:12:f2:21:cf:0    ucst   333     5 me0.0
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
2.2.2.0/24         intf     0                    rslv  1309     1 ae0.0
2.2.2.0/32         dest     0 2.2.2.0            recv  1307     1 ae0.0
2.2.2.1/32         dest     0 0:21:59:cc:89:c0   ucst  1320     1 ae0.0
2.2.2.2/32         intf     0 2.2.2.2            locl  1308     2
2.2.2.2/32         dest     0 2.2.2.2            locl  1308     2
2.2.2.255/32       dest     0 2.2.2.255          bcst  1306     1 ae0.0
3.3.3.0/24         intf     0                    rslv  1313     1 ae1.0
3.3.3.0/32         dest     0 3.3.3.0            recv  1311     1 ae1.0
3.3.3.1/32         intf     0 3.3.3.1            locl  1312     2
3.3.3.1/32         dest     0 3.3.3.1            locl  1312     2
3.3.3.2/32         dest     0 0:21:59:cc:89:c1   ucst  1321    24 ae1.0
3.3.3.255/32       dest     0 3.3.3.255          bcst  1310     1 ae1.0
4.4.4.0/24         user     0 3.3.3.2            ucst  1321    24 ae1.0
8.8.8.8/32         user     0 3.3.3.2            ucst  1321    24 ae1.0
9.9.9.9/32         intf     0 9.9.9.9            locl  1280     1
10.10.10.10/32     user     0 3.3.3.2            ucst  1321    24 ae1.0
10.93.8.0/21       intf     0                    rslv   323     1 me0.0
10.93.8.0/32       dest     0 10.93.8.0          recv   321     1 me0.0
10.93.13.238/32    intf     0 10.93.13.238       locl   322     2
10.93.13.238/32    dest     0 10.93.13.238       locl   322     2
10.93.15.254/32    dest     0 0:12:f2:21:cf:0    ucst   333     5 me0.0
10.93.15.255/32    dest     0 10.93.15.255       bcst   320     1 me0.0
14.14.14.0/24      ifdn     0                    rslv  1319     1 ge-0/0/25.0
14.14.14.0/32      iddn     0 14.14.14.0         recv  1317     1 ge-0/0/25.0
14.14.14.2/32      user     0                    rjct    36     2
14.14.14.2/32      intf     0 14.14.14.2         locl  1318     2
14.14.14.2/32      iddn     0 14.14.14.2         locl  1318     2
14.14.14.255/32    iddn     0 14.14.14.255       bcst  1316     1 ge-0/0/25.0
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
255.255.255.255/32 perm     0                    bcst    32     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 summary

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 mpls interface

List of Syntax

Syntax on page 374
Syntax (EX Series Switches) on page 374

Syntax

show mpls interface
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show mpls interface

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.

Options

none—Display information about MPLS-enabled interfaces.

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

Additional Information

MPLS is enabled on an interface when the interface is configured with both the **set protocol mpls interface *interface-name*** and **set interface *interface-name* unit 0 family mpls** statements.

Required Privilege Level

view

List of Sample Output

show mpls interface on page 375

Output Fields

Table 40 on page 374 describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 40: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **Up** or **Dn** (down).
Administrative groups	Administratively assigned colors of the link.
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the **maximum-labels** statement at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family mpls]** or the **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchy levels.

Table 40: show mpls interface Output Fields *(continued)*

Field Name	Field Description
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the **protection-revert-time** statement at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels.
Always mark connection protection tlv	**Enabled** or **Disabled**: **Enabled** indicates that the **always-mark-connection-protection-tlv** statement is configured at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the **switch-away-lsps** statement must be configured.
Switch away lsps	**Enabled** or **Disabled**: **Enabled** indicates that the **switch-away-lsps** statement is configured at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

```
user@host> show mpls interface

Interface: ge-0/2/1.57
  State: Up
  Administrative group: <none>
  Maximum labels: 5
  Static protection revert time: 5 seconds
  Always mark connection protection tlv: Disabled
  Switch away lsps : Disabled
```

show link-management statistics

Syntax

show link-management statistics
<peer <name *name*>>

Release Information

Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display statistical information for Link Management Protocol (LMP) packets.

Options

none—Display information for all peers.

peer <name *name*>—(Optional) Display information for all peers or for the specified peer only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management peer on page 349
- show link-management routing on page 351
- show link-management te-link on page 357

List of Sample Output

show link-management statistics on page 377

Output Fields

Table 34 on page 355 describes the output fields for the **show link-management statistics** command. Output fields are listed in the approximate order in which they appear.

Table 41: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.

Table 41: show link-management statistics Output Fields *(continued)*

Field Name	Field Description
Stale acknowledgments	Number of **configAck** and **LinkSummaryAck** packets received that have a stale message ID.
Stale negative acknowledgments	Number of **configNack** and **LinkSummaryNack** packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

show link-management statistics

```
user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
    Received packets
        Config: 1
        Hello: 2572
    Small packets: 0
    Wrong protocol version: 0
    Messages for unknown peer: 0
    Messages for bad state: 0
    Stale acknowledgments: 0
    Stale negative acknowledgments: 0
    Sent packets
        Config: 2
        ConfigAck: 1
        Hello: 2572
    Retransmitted packets
        Config: 1
```


show link-management te-link

Syntax

show link-management te-link
<brief | detail>
<name *name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

Options

none—Display information for all traffic-engineered links.

brief | detail—(Optional) Display the specified level of output.

name ***name***—(Optional) Display information for the specified traffic-engineered link only.

Required Privilege Level

view

Related Documentation

- show link-management on page 345
- show link-management peer on page 349
- show link-management routing on page 351
- show link-management statistics on page 355

List of Sample Output

show link-management te-link on page 379

Output Fields

Table 35 on page 357 describes the output fields for the **show link-management te-link** command. Output fields are listed in the approximate order in which they appear.

Table 42: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down.**
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.

Table 42: show link-management te-link Output Fields *(continued)*

Field Name	Field Description
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, **Packet**, and **PDH**.
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

Sample Output

show link-management te-link

```
user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-bd    Dn       43077          0           0bps No

 TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
   Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
 Available bandwidth: 8Mbps
```

```
Name         State Local ID  Remote ID      Bandwidth Used  LSP-name
falsp-be     Up        43076          0        10Mbps  Yes  e2elsp-bf
```


show mpls call-admission-control

List of Syntax

Syntax on page 381
Syntax (EX Series Switches) on page 381

Syntax

show mpls call-admission-control
<logical-system (all | *logical-system-name*)>
<*lsp-name*>

Syntax (EX Series Switches)

show mpls call-admission-control
<*lsp-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.

Options

none—Display CAC information for all LSPs.

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

lsp-name—(Optional) Display CAC information for the specified LSP only.

Additional Information

The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.

Required Privilege Level

view

List of Sample Output

show mpls call-admission-control on page 382

Output Fields

Table 36 on page 360 describes the output fields for the **show mpls call-admission-control** command. Output fields are listed in the approximate order in which they appear.

Table 43: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at **ct0**) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.

Table 43: show mpls call-admission-control Output Fields *(continued)*

Field Name	Field Description
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

Sample Output

show mpls call-admission-control

```
user@host# show mpls call-admission-control

LSP name: pro1-be
 *Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
 *Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
 *Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
 *Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
 *Primary   prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby   sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

show mpls cspf

List of Syntax

Syntax on page 383
Syntax (EX Series Switches) on page 383

Syntax

show mpls cspf
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show mpls cspf

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

Options

none—Display MPLS CSFP statistics.

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

Required Privilege Level

view

List of Sample Output

show mpls cspf on page 384

Output Fields

Table 37 on page 362 describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

Table 44: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of aborted computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.
successful	Number of path computations that were successfully completed.

Table 44: show mpls cspf Output Fields *(continued)*

Field Name	Field Description
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

Sample Output

show mpls cspf

```
user@host> show mpls cspf
CSPF statistics
Queue length   current       maximum      dequeued
                     0             0             0
Paths            total    successful      no route     sys error         CSPFs
                     0             0             0             0             0
Time (secs)      total         CSPFs  avg per CSPF      % of rpd
              0.000000      0.000000      0.000000        0.0000
```

show mpls diffserv-te

List of Syntax

Syntax on page 385
Syntax (EX Series Switches) on page 385

Syntax

show mpls diffserve-te
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show mpls diffserve-te

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.

Options

none—Display DiffServ classes and priorities used by MPLS LSPs.

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

Required Privilege Level

view

List of Sample Output

show mpls diffserv-te on page 386

Output Fields

Table 38 on page 364 describes the output fields for the **show mpls diffserv-te** command. Output fields are listed in the approximate order in which they appear.

Table 45: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	MPLS class type that corresponds to the DiffServ traffic engineering class: • **ct0**—Best effort • **ct1**—Assured forwarding • **ct2**—Expedited forwarding • **ct3**—Network control
Priority	MPLS preemption priority for this class type, a value from **0** through **7**. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

Sample Output

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class     Traffic class       Priority
te0          ct0                 3
te1          ct1                 2
```

show route forwarding-table

Syntax

show route forwarding-table
<detail | extensive | summary>
<ccc *ccc-interface-name*>
<*destination*>
<family *family-name*>
<label *label*>
<matching *ip_prefix*>
<multicast>
<vpn *vpn*>

Release Information

Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

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.

Options

none—Display the routes in the forwarding table.

detail | extensive | summary—(Optional) Display the specified level of output.

ccc—(Optional) Display the specified circuit cross-connect interface name for entries to match.

destination—(Optional) Display the destination prefix.

family *family-name*—(Optional) Display routing table entries for the specified family: **ethernet-switching**, **inet**, **inet6**, **iso**, **mpls**, **vlan classification**.

label *label*—(Optional) Display route entries for the specified label name.

matching *ip_prefix*—(Optional) Display route entries for the specified IP prefix.

multicast—(Optional) Display route entries for multicast routes.

vpn *vpn*—(Optional) Display route entries for the specified VPN.

Required Privilege Level

view

Related Documentation

- *Example: Configuring MPLS on EX8200 and EX4500 Switches*
- *Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)*

List of Sample Output

show route forwarding-table on page 389
show route forwarding-table summary on page 390

show route forwarding-table extensive on page 390
show route forwarding-table ccc on page 392
show route forwarding-table family (MPLS) on page 392
show route forwarding-table family (IPv6) on page 393
show route forwarding-table label on page 393
show route forwarding-table matching on page 393
show route forwarding-table multicast on page 393

Output Fields

Table 39 on page 367 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 46: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table (for example, **inet**, **inet6**, **mpls**).	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**).	All levels
Destination	Destination of the route.	**detail, extensive**
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **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**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number*** —Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**

Table 46: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Nexthop	IP address of the next hop to the destination.	**detail, extensive**
Next hop type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)** —Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the 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.	**detail, extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail, extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	none **detail, extensive**
Next-hop interface (Netif)	Interface used to reach the next hop.	none **detail, 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 Interfaces	The next hop layer 2 interfaces. Seen with **multicast** option only.	**extensive**

Sample Output

show route forwarding-table

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

```
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:12:f2:21:cf:0    ucst   333     5 me0.0
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
2.2.2.0/24         intf     0                    rslv  1309     1 ae0.0
2.2.2.0/32         dest     0 2.2.2.0            recv  1307     1 ae0.0
2.2.2.1/32         dest     0 0:21:59:cc:89:c0   ucst  1320     1 ae0.0
2.2.2.2/32         intf     0 2.2.2.2            locl  1308     2
2.2.2.2/32         dest     0 2.2.2.2            locl  1308     2
2.2.2.255/32       dest     0 2.2.2.255          bcst  1306     1 ae0.0
3.3.3.0/24         intf     0                    rslv  1313     1 ae1.0
3.3.3.0/32         dest     0 3.3.3.0            recv  1311     1 ae1.0
3.3.3.1/32         intf     0 3.3.3.1            locl  1312     2
3.3.3.1/32         dest     0 3.3.3.1            locl  1312     2
3.3.3.2/32         dest     0 0:21:59:cc:89:c1   ucst  1321    24 ae1.0
3.3.3.255/32       dest     0 3.3.3.255          bcst  1310     1 ae1.0
4.4.4.0/24         user     0 3.3.3.2            ucst  1321    24 ae1.0
8.8.8.8/32         user     0 3.3.3.2            ucst  1321    24 ae1.0
9.9.9.9/32         intf     0 9.9.9.9            locl  1280     1
10.10.10.10/32     user     0 3.3.3.2            ucst  1321    24 ae1.0
10.93.8.0/21       intf     0                    rslv   323     1 me0.0
10.93.8.0/32       dest     0 10.93.8.0          recv   321     1 me0.0
10.93.13.238/32    intf     0 10.93.13.238       locl   322     2
10.93.13.238/32    dest     0 10.93.13.238       locl   322     2
10.93.15.254/32    dest     0 0:12:f2:21:cf:0    ucst   333     5 me0.0
10.93.15.255/32    dest     0 10.93.15.255       bcst   320     1 me0.0
14.14.14.0/24      ifdn     0                    rslv  1319     1 ge-0/0/25.0
14.14.14.0/32      iddn     0 14.14.14.0         recv  1317     1 ge-0/0/25.0
14.14.14.2/32      user     0                    rjct    36     2
14.14.14.2/32      intf     0 14.14.14.2         locl  1318     2
14.14.14.2/32      iddn     0 14.14.14.2         locl  1318     2
14.14.14.255/32    iddn     0 14.14.14.255       bcst  1316     1 ge-0/0/25.0
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
255.255.255.255/32 perm     0                    bcst    32     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 summary**

```
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 mpls interface

List of Syntax
Syntax on page 395
Syntax (EX Series Switches) on page 395

Syntax
show mpls interface
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show mpls interface

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description
Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.

Options
none—Display information about MPLS-enabled interfaces.

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

Additional Information
MPLS is enabled on an interface when the interface is configured with both the **set protocol mpls interface *interface-name*** and **set interface *interface-name* unit 0 family mpls** statements.

Required Privilege Level
view

List of Sample Output
show mpls interface on page 396

Output Fields
Table 40 on page 374 describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 47: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **Up** or **Dn** (down).
Administrative groups	Administratively assigned colors of the link.
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the **maximum-labels** statement at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family mpls]** or the **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchy levels.

Table 47: show mpls interface Output Fields *(continued)*

Field Name	Field Description
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the **protection-revert-time** statement at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels.
Always mark connection protection tlv	**Enabled** or **Disabled**: **Enabled** indicates that the **always-mark-connection-protection-tlv** statement is configured at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the **switch-away-lsps** statement must be configured.
Switch away lsps	**Enabled** or **Disabled**: **Enabled** indicates that the **switch-away-lsps** statement is configured at the **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name* static]** or the **[edit protocols mpls interface *interface-name* static]** hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

```
user@host> show mpls interface

Interface: ge-0/2/1.57
  State: Up
  Administrative group: <none>
  Maximum labels: 5
  Static protection revert time: 5 seconds
  Always mark connection protection tlv: Disabled
  Switch away lsps : Disabled
```


show mpls lsp

List of Syntax
Syntax on page 397
Syntax (EX Series Switches) on page 397

Syntax

show mpls lsp
<brief | detail | extensive | terse>
<autobandwidth>
<bidirectional | unidirectional>
<bypass>
<count-active-routes>
<defaults>
<descriptions>
<down | up>
<externally-controlled>
<externally-provisioned>
<logical-system (all | *logical-system-name*)>
<*lsp-type*>
<name *name*>
<p2mp>
<statistics>
<transit>

Syntax (EX Series Switches)

show mpls lsp
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<descriptions>
<down | up>
<externally-controlled>
<externally-provisioned>
<*lsp-type*>
<name *name*>
<p2mp>
<statistics>
<transit>

Release Information

Command introduced before Junos OS Release 7.4.
defaults option added in Junos OS Release 8.5.
Command introduced in Junos OS Release 9.5 for EX Series switches.
autobandwidth option added in Junos OS Release 11.4.
externally-controlled option added in Junos OS Release 12.3.
externally-provisioned option added in Junos OS Release 13.3.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all configured and active dynamic MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

autobandwidth—(Optional) Display automatic bandwidth information. This option is explained separately (see show mpls lsp autobandwidth).

bidirectional | unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.

bypass—(Optional) Display LSPs used for protecting other LSPs.

count-active-routes—(Optional) Display active routes for LSPs.

defaults—(Optional) Display the MPLS LSP default settings.

descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

externally-controlled—(Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

externally-provisioned—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

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

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name ***name***—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

NOTE: If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.

transit—(Optional) Display LSPs transiting this routing device.

Required Privilege Level view

Related Documentation

- clear mpls lsp on page 325
- show mpls lsp autobandwidth on page 413

List of Sample Output

Output Fields Table 48 on page 399 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 48: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
P2MP name	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format ***identifier*:vpls:*router-id*:*routing-instance-name***. The ***identifier*** is automatically generated by Junos OS.	All levels

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
P2MP branch count	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
P	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
address	(**detail** and **extensive**) Destination (egress routing device) of the LSP.	**detail extensive**
To	Destination (egress routing device) of the session.	**brief**
From	Source (ingress routing device) of the session.	**brief detail**
State	State of the LSP handled by this RSVP session: **Up**, **Dn** (down), or **Restart**.	**brief detail**
Active Route	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (**mpls.0**).	**detail extensive**
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	**brief**
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	**brief**
ActivePath	(Ingress LSP) Name of the active path: **Primary** or **Secondary**.	**detail extensive**
LSPname	Name of the LSP.	**brief detail**
Statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	**extensive**
Aggregate statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the **clear mpls lsp statistics** command.	**extensive**
Packets	Displays the number of packets transmitted over the LSP.	**brief extensive**
Bytes	Displays the number of bytes transmitted over the LSP.	**brief extensive**
DiffServeInfo	Type of LSP: multiclass LSP **(multiclass diffServ-TE LSP)** or Differentiated-Services-aware traffic engineering LSP **(diffServ-TE LSP)**.	**detail**
LSPtype	Type of LSP: static **Static configured** or dynamic **Dynamic configured**. Also indicates if the LSP is a **Penultimate hop popping** LSP or an **Ultimate hop popping** LSP.	**detail extensive**
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. **LSPpath** can be either **primary** or **secondary** and can be displayed on the ingress, egress, and transit routing devices.	**detail**
Bidir	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
Bidirectional	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
FastReroute desired	Fast reroute has been requested by the ingress routing device.	**detail**
Link protection desired	Link protection has been requested by the ingress routing device.	**detail**
Node/Link protection desired	Link protection has been requested by the ingress routing device.	**detail extensive**
LoadBalance	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: **Most-fill**, **Least-fill**, or **Random**.	**detail extensive**
Signal type	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: **DS0**, **DS3**, **STS-1**, **STM-1**, or **STM-4**.	All levels
Encoding type	LSP encoding type: **Packet**, **Ethernet**, **PDH**, **SDH/SONET**, **Lambda**, or **Fiber**.	All levels
Switching type	Type of switching on the links needed for the LSP: **Fiber**, **Lamda**, **Packet**, **TDM**, or **PSC-1**.	All levels
GPID	Generalized Payload Identifier (identifier of the payload carried by an LSP): **HDLC**, **Ethernet**, **IPv4**, **PPP**, or **Unknown**.	All levels
Protection	Configured protection capability desired for the LSP: **Extra**, **Enhanced**, **none**, **One plus one**, **One to one**, or **Shared**.	All levels
Upstream label in	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
Upstream label out	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels
Suggested label received	(Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.	All levels
Suggested label sent	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
Autobandwidth	(Ingress LSP) The LSP is performing autobandwidth allocation.	**detail extensive**
MinBW	(Ingress LSP) Configured minimum value of the LSP, in bps.	**detail extensive**

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
MaxBW	(Ingress LSP) Configured maximum value of the LSP, in bps.	**detail extensive**
Dynamic MinBW	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	**detail extensive**
Adjustment Timer	(Ingress LSP) Configured value for the **adjust-timer** statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	**detail extensive**
Adjustment Threshold	(Ingress LSP) Configured value for the **adjust-threshold** statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	**detail extensive**
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	**detail extensive**
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	**detail extensive**
Max AvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	**detail extensive**
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	**detail extensive**
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	**detail extensive**
Bandwidth Adjustment in *nnn* second(s)	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	**detail extensive**
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	**detail extensive**
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	**detail extensive**
Underflow Max AvgBW	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	**detail extensive**
Active path indicator	(Ingress LSP) A value of ***** indicates that the path is active. The absence of ***** indicates that the path is not active. In the following example, “long” is the active path. ***Primary long** **Standby short**	**detail extensive**
Primary	(Ingress LSP) Name of the primary path.	**detail extensive**
Secondary	(Ingress LSP) Name of the secondary path.	**detail extensive**

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Standby	(Ingress LSP) Name of the path in standby mode.	**detail extensive**
State	(Ingress LSP) State of the path: **Up** or **Dn** (down).	**detail extensive**
COS	(Ingress LSP) Class-of-service value.	**detail extensive**
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	**detail extensive**
Priorities	(Ingress LSP) Configured value of the setup priority and the hold priority respecitively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	**detail extensive**
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	**detail extensive**
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	**detail extensive**
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	**detail extensive**
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (**S**) or loose (**L**).	**detail extensive**
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	**detail extensive**

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Received RRO	(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If **Received RRO** is different from **Computed ERO**, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node: • **0x01**—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the **SESSION_ATTRIBUTE** object of the corresponding Path message. • **0x02**—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • **0x03**—Combination of **0x01** and **0x02**. • **0x04**—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • **0x08**—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the **Local protection available** bit is set but the **Node protection** bit is cleared. • **0x09**—Detour is established. Combination of **0x01** and **0x08**. • **0x10**—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • **0x20**—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • **0xb**—Detour is in use. Combination of **0x01**, **0x02**, and **0x08**.	**detail extensive**
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	**extensive**
Date	(Ingress LSP) Date of the LSP event.	**extensive**
Time	(Ingress LSP) Time of the LSP event.	**extensive**
Event	(Ingress LSP) Description of the LSP event.	**extensive**
Created	(Ingress LSP) Date and time the LSP was created.	**extensive**
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	**brief detail extensive**
Labelin	Incoming label for this LSP.	**brief detail**
Labelout	Outgoing label for this LSP.	**brief detail**

Table 48: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
LSPname	Name of the LSP.	**brief detail**
Time left	Number of seconds remaining in the lifetime of the reservation.	**detail**
Since	Date and time when the RSVP session was initiated.	**detail**
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	**detail**
Port number	Protocol ID and sender or receiver port used in this RSVP session.	**detail**
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	**detail**
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	**detail**
RESV rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the **PATH rcvfrom** field, indicates that the RSVP negotiation is complete.	**detail**
Record route	Recorded route for the session, taken from the record route object.	**detail**
Soft preempt	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	**detail**
Soft preemption pending	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	**detail**
MPLS-TE LSP Defaults	Default settings for MPLS traffic engineered LSPs: • **LSP Holding Priority**—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • **LSP Setup Priority**—Determines whether a new LSP that preempts an existing LSP can be established. • **Hop Limit**—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • **Bandwidth**—Specifies the bandwidth in bits per second for the LSP. • **LSP Retry Timer**—Length of time in seconds that the ingress router waits between attempts to establish the primary path.	**defaults**

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth** . You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

Sample Output

show mpls lsp defaults

```
user@host> show mpls lsp defaults
   MPLS-TE LSP Defaults
      LSP Holding Priority      0
      LSP Setup Priority        7
      Hop Limit                 255
      Bandwidth                 0
      LSP Retry Timer           30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To              LSP name                             Description
10.0.0.195      to-sanjose                           to-sanjose-desc
10.0.0.195      to-sanjose-other-desc                other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp extensive

```
user@host> show mpls lsp extensive
Ingress LSP: 4 sessions

 1.1.1.1
   From: 3.3.3.3, State: Up, ActiveRoute: 0, LSPname: m120b-to-mx960
   ActivePath: DEFAULT (primary)
   FastReroute desired
   LSPtype: Static Configured, Penultimate hop popping
   LoadBalance: Random
   Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   DEFAULT          State: Up
     Priorities: 7 0
     SmartOptimizeTimer: 180
     Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 310)
 10.0.35.5 S 10.0.15.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          10.0.34.4(flag=1) 10.0.14.1
    50 Sep 13 16:08:19.712 Record Route:  10.0.35.5(flag=1) 10.0.15.1
    49 Sep 13 16:08:16.720 Record Route:  10.0.34.4(flag=1) 10.0.14.1
    48 Sep 13 16:08:16.699 Fast-reroute Detour Up
    47 Sep 13 16:08:13.702 Record Route:  10.0.34.4 10.0.14.1
    46 Sep 13 16:08:13.702 Up
    45 Sep 13 16:08:13.672 Originate make-before-break call
    44 Sep 13 16:08:13.672 CSPF: computation result accepted  10.0.34.4 10.0.14.1

    43 Sep 13 16:08:13.672 Selected as active path
    42 Sep 13 16:08:13.672 Make-before-break: Switched to new instance
    41 Sep 13 16:08:01.685 Pending path switchover, skip CSPF run[3 times]
    40 Sep 13 16:06:33.910 Deselected as active
    39 Sep 13 16:06:33.910 Pending path switchover, skip CSPF run

    38 Sep 13 16:06:19.521 Record Route:  10.0.35.5 10.0.15.1
    37 Sep 13 16:06:19.518 ResvTear received
    36 Sep 13 16:06:19.518 Fast-reroute Detour Down
    35 Sep 13 16:06:16.676 Record Route:  10.0.35.5(flag=1) 10.0.15.1
    34 Sep 13 16:06:13.670 Record Route:  10.0.35.5 10.0.15.1
    33 Sep 13 16:06:13.670 Up
    32 Sep 13 16:06:13.569 Pending path switchover, skip CSPF run

    31 Sep 13 16:06:13.569 CSPF: link down/deleted:
10.0.34.3(3.3.3.3:79)(m120-b-re1.00/3.3.3.3)->0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)

    30 Sep 13 16:06:13.552 Pending path switchover, skip CSPF run

    29 Sep 13 16:06:13.552 CSPF: link down/deleted:
0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)->0.0.0.0(4.4.4.4:0)(m10i-a-re0.00/4.4.4.4)

    28 Sep 13 16:06:13.549 Originate make-before-break call
    27 Sep 13 16:06:13.549 CSPF: computation result accepted  10.0.35.5 10.0.15.1

    26 Sep 13 16:06:13.548 Tunnel local repaired
    25 Sep 13 16:06:13.546 Record Route:  10.0.23.2 10.0.12.1
    24 Sep 13 16:06:13.546 10.0.34.3: Tunnel local repaired
    23 Sep 13 16:06:13.546 10.0.34.3: Down
    22 Sep 13 16:03:46.842 Fast-reroute Detour Up
```

```
    21 Sep 13 16:03:42.730 Record Route:  10.0.34.4(flag=1) 10.0.14.1
    20 Sep 13 16:03:39.836 Selected as active path
    19 Sep 13 16:03:39.834 Record Route:  10.0.34.4 10.0.14.1
    18 Sep 13 16:03:39.834 Up
    17 Sep 13 16:03:39.698 Originate Call
   16 Sep 13 16:03:39.698 CSPF: computation result accepted  10.0.34.4 10.0.14.1

    15 Sep 13 16:03:39.697 Clear Call
    14 Sep 13 16:03:39.696 Deselected as active
    13 Sep 13 16:03:37.837 Record Route:  10.0.34.4 10.0.14.1
    12 Sep 13 16:03:32.829 Fast-reroute Detour Down
    11 Sep 13 16:02:15.493 Record Route:  10.0.34.4(flag=1) 10.0.14.1
    10 Sep 13 16:02:15.486 Fast-reroute Detour Up
     9 Sep 13 16:02:12.468 Record Route:  10.0.34.4 10.0.14.1
     8 Sep 13 16:02:07.460 Fast-reroute Detour Down
     7 Sep 13 15:57:46.741 Fast-reroute Detour Up
     6 Sep 13 15:57:40.768 Record Route:  10.0.34.4(flag=1) 10.0.14.1
     5 Sep 13 15:57:37.761 Selected as active path
     4 Sep 13 15:57:37.760 Record Route:  10.0.34.4 10.0.14.1
     3 Sep 13 15:57:37.760 Up
     2 Sep 13 15:57:37.733 Originate Call
    1 Sep 13 15:57:37.733 CSPF: computation result accepted  10.0.34.4 10.0.14.1

   Created: Fri Sep 13 15:57:38 2013
 Total 1 displayed, Up 1, Down 0

 Egress LSP: 4 sessions, 6 detours
 Total 0 displayed, Up 0, Down 0

 Transit LSP: 6 sessions, 1 detours

 1.1.1.1
   From: 3.3.3.3, LSPstate: Up, ActiveRoute: 0
   LSPname: m120b-to-mx960, LSPpath: Primary
   Suggested label received: -, Suggested label sent: -
   Recovery label received: -, Recovery label sent: 302288
   Resv style: 1 FF, Label in: 300416, Label out: 302288
   Time left:  147, Since: Fri Sep 13 16:08:16 2013
   Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
   Port number: sender 4 receiver 13955 protocol 0
   Detour branch from 10.0.34.4, to skip 1.1.1.1, Up
     Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
     Adspec: received MTU 1500
     Path MTU: received 0
     PATH rcvfrom: 10.0.34.4 (ge-4/3/7.0) 7 pkts
     Adspec: received MTU 1500 sent MTU 1500
     PATH sentto: 10.0.35.5 (ge-3/1/0.0) 7 pkts
     RESV rcvfrom: 10.0.35.5 (ge-3/1/0.0) 7 pkts
     Explct route: 10.0.35.5 10.0.15.1
     Record route: 10.0.34.3 10.0.34.4 <self>10.0.35.5 10.0.15.1
  Label in: 300416, Label out: 302288
Total 1 displayed, Up 1, Down 0
```

show mpls lsp ingress extensive

```
user@host> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

50.0.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
  ActivePath:  (primary)
```



```
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    OptimizeTimer: 300
    SmartOptimizeTimer: 180
    Reoptimization in 240 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          1.1.1.2 4.4.4.1 5.5.5.2
   17 Aug  3 13:17:33.601 CSPF: computation result ignored, new path less avail
bw[3 times]
   16 Aug  3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
 times]
   15 Aug  3 12:54:36.678 Selected as active path
   14 Aug  3 12:54:36.676 Record Route:  1.1.1.2 4.4.4.1 5.5.5.2
   13 Aug  3 12:54:36.676 Up
   12 Aug  3 12:54:33.924 Deselected as active
   11 Aug  3 12:54:33.924 Originate Call
   10 Aug  3 12:54:33.923 Clear Call
    9 Aug  3 12:54:33.923 CSPF: computation result accepted  1.1.1.2 4.4.4.1
5.5.5.2
    8 Aug  3 12:54:33.922 2.2.2.2: No Route toward dest
    7 Aug  3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
 times]
    6 Aug  3 12:35:03.830 Selected as active path
    5 Aug  3 12:35:03.828 Record Route:  2.2.2.2 3.3.3.2
    4 Aug  3 12:35:03.827 Up
    3 Aug  3 12:35:03.814 Originate Call
    2 Aug  3 12:35:03.814 CSPF: computation result accepted  2.2.2.2 3.3.3.2
    1 Aug  3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
  Created: Tue Aug  3 12:34:35 2010
Total 1 displayed, Up 1, Down 0
```

show mpls lsp extensive (automatic bandwidth adjustment enabled)

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath:  (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
  Adjustment Timer: 300 secs AdjustThreshold: 25%
  Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
  Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps

  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
```

```
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
10.0.0.22(Label=3)
   12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
   11 Apr 30 10:25:16.023 Record Route:  192.168.0.6(flag=0x20)
10.0.0.18(Label=299792) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
   10 Apr 30 10:25:16.023 Up
    9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from
 300 bps to 1000 bps
    8 Apr 30 10:25:15.946 Originate make-before-break call
    7 Apr 30 10:25:15.946 CSPF: computation result accepted  10.0.0.18 10.0.0.22

    6 Apr 30 10:16:42.891 Selected as active path
    5 Apr 30 10:16:42.891 Record Route:  192.168.0.6(flag=0x20)
10.0.0.18(Label=299776) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    4 Apr 30 10:16:42.890 Up
    3 Apr 30 10:16:42.828 Originate Call
    2 Apr 30 10:16:42.828 CSPF: computation result accepted  10.0.0.18 10.0.0.22

    1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
  Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp

```
user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To               From             State Rt P ActivePath       LSPname
10.255.245.51    10.255.245.50    Up     0 * path1            p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To               From             State Rt P ActivePath       LSPname
10.255.245.51    10.255.245.50    Up     0 * path1            p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random
```



```
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          192.168.208.17
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail count-active-routes

```
user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
  From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  Adjustment Timer: 300 secs
  Max AvgBW util: 60.2599Mbps, Bandwidth Adjustment in 0 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
 10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp statistics extensive

```
user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions
```

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath:  (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          10.0.0.18 10.0.0.22
    6 Oct  3 11:18:28.281 Selected as active path
    5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
    4 Oct  3 11:18:28.280 Up
    3 Oct  3 11:18:27.995 Originate Call
    2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

    1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0
```


show mpls lsp autobandwidth

Syntax

show mpls lsp autobandwidth
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 11.4.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Display automatic bandwidth information for the LSP(s).

Options

brief | detail | extensive —(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

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

Required Privilege Level

view

Related Documentation

- show mpls lsp on page 397

List of Sample Output

show mpls lsp autobandwidth on page 414

Output Fields

Table 49 on page 413 describes the output fields for the **show mpls lsp autobandwidth** command. Output fields are listed in the approximate order in which they appear.

Table 49: show mpls lsp autobandwidth Output Fields

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All Levels
From	Source (ingress routing device) of the session.	All Levels
LSPname	Name of the LSP.	All Levels
Min BW	(Ingress LSP) Configured minimum value of the LSP, in bps.	**detail extensive**
Max BW	(Ingress LSP) Configured maximum value of the LSP, in bps.	**detail extensive**
Max AvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	**detail extensive**
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	**detail extensive**
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	**detail extensive**
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	**detail extensive**

Table 49: show mpls lsp autobandwidth Output Fields *(continued)*

Field Name	Field Description	Level of Output
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	**detail extensive**
Adjustment Timer	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	**detail extensive**
Adjustment Threshold	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	**detail extensive**
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	**detail extensive**
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	**detail extensive**
Last BW	Previous active bandwidth of the LSP.	**detail extensive**
Last Requested BW	Bandwidth requested in the previous automatic bandwidth adjustment.	**detail extensive**
Last Signaled BW	Bandwidth signaled in the previous automatic bandwidth adjustment.	**detail extensive**
Highest Watermark BW	Maximum bandwidth used by the LSP.	**detail extensive**
Total AutoBw Adjustments	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	**detail extensive**
Successful Adjustments	Number of successful automatic bandwidth adjustments.	**detail extensive**
Failed Adjustments	Number of failed automatic bandwidth adjustments.	**detail extensive**

Sample Output

show mpls lsp autobandwidth

```
user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun  3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0
```

show mpls path

List of Syntax
Syntax on page 416
Syntax (EX Series Switches) on page 416

Syntax

show mpls path
<logical-system (all | *logical-system-name*)>
<*path-name*>

Syntax (EX Series Switches)

show mpls path
<*path-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all MPLS LSPs.

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

path-name—(Optional) Display information about the specified LSP only.

Required Privilege Level

view

List of Sample Output

show mpls path on page 416

Output Fields

Table 50 on page 416 describes the output fields for the **show mpls path** command. Output fields are listed in the approximate order in which they appear.

Table 50: show mpls path Output Fields

Field Name	Field Description
Path name	Information about ingress LSPs. Each path has one line of output.
Address	Addresses of the routing devices that form the LSP.
Strict/loose address	Whether the address is a configured as a strict or loose address.

Sample Output

show mpls path

```
user@host> show mpls path
Path name        Address          Strict/loose address
p1               123.456.55.6     Strict
```

```
                 123.456.1.6      Loose
p2               191.456.1.4      Strict
```

show mpls static-lsp

Syntax

show mpls static-lsp
<brief | detail | extensive | terse>
<bypass>
<descriptions>
<down | up>
<ingress>
<logical-system (all | *logical-system-name*)>
<*lsp-type*>
<name *name*>
<statistics>
<transit>

Release Information

Command introduced in Junos OS Release 10.1.
Command updated in Junos OS Release 14.1X53-D25 to accommodate the stitching feature of MPLS.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all configured and active static MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

bypass—(Optional) Display LSPs used for protecting other static LSPs.

descriptions—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

down | up—(Optional) Display only static LSPs that are inactive or active, respectively.

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

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified static LSP or group of LSPs.

statistics—(Optional) Display accounting information about static LSPs.

transit—(Optional) Display static LSPs transiting this routing device.

Required Privilege Level view

List of Sample Output
show mpls static-lsp extensive on page 420
show mpls static-lsp statistics ingress on page 420
show mpls static-lsp (when MPLS stitching is used) on page 420

Output Fields Table 51 on page 419 describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 51: show mpls static-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSPs	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
Transit LSPs	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Bypass LSPs	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
LSPname	Name of the static LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
State	State of the static LSP handled by this RSVP session: **Up**, **Dn** (down), or **Restart**.	All levels
Packets	Number of packet transiting the static LSP (**statistics** option only).	All levels
Bytes	Number of bytes transiting the static LSP (**statistics** option only).	All levels
Nexthop	IP address for the next-hop router for the static LSP.	**detail**, **extensive**
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
Link protection desired	Link protection has been requested by the ingress routing device.	**detail**, **extensive**
LabelOperation	Label operation to perform: **Push**, **Pop**, **Swap**.	**detail**, **extensive**
Outgoing-label	Outgoing label to use for the MPLS packet in either push or swap label operations.	**detail**, **extensive**

Table 51: show mpls static-lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Created	(Ingress LSP) Date and time the static LSP was created.	**extensive**
Bandwidth	Bandwidth configured for the static LSP.	**detail**, **extensive**
Resv style	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	All levels

Sample Output

show mpls static-lsp extensive

```
user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
  State: Dn
  Nexthop: 192.168.10.1
  LabelOperation: Push, Outgoing-label: 1000001
  Created: Thu Jan 14 16:44:43 2010
  Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0
```

show mpls static-lsp statistics ingress

```
user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPname                     To              State     Packets              Bytes
alpha-to-beta               192.168.14.1    Dn             NA                 NA
Total 1, displayed 1, Up 0, Down 1
```

show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```
user@host> show mpls static-lsp
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0
Transit LSPs: LSPname                        Incoming-label  State
to-165                                  1000000         InProgress
```


show ted database

List of Syntax

Syntax on page 421
Syntax (EX Series Switches) on page 421

Syntax

show ted database
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*system-name*>

Syntax (EX Series Switches)

show ted database
<brief | detail | extensive>
<*system-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.

Options

none—Display standard information about all entries in the traffic engineering database.

brief | detail | extensive—(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.

system-name—(Optional) Display traffic engineering database information for a particular system.

Required Privilege Level

view

List of Sample Output

show ted database brief on page 423
show ted database detail on page 424
show ted database extensive on page 425

Output Fields

Table 52 on page 421 describes the output fields for the **show ted database** command. Output fields are listed in the approximate order in which they appear.

Table 52: show ted database Output Fields

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels

Table 52: show ted database Output Fields *(continued)*

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	**brief**
NodeID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**extensive**
Type	Type of node. It can be either **Rtr** (router) or **Net** (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: • **IS-IS(1)**—IS-IS Level 1. • **IS-IS(2)**—IS-IS Level 2. • **OSPF (area-number)**—OSPF from the specified area.	All levels
To	Address on the far end of a link.	**detail extensive**
Local	Address of the local interface being used to reach the remote node.	**detail extensive**
Remote	Address of the interface on the remote node.	**detail extensive**
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	**detail extensive**
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	**detail extensive**
Metric	Configured traffic engineering metric.	**extensive**
Static BW	Total interface bandwidth in bps.	**extensive**
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the **subscription** statement when configuring RSVP.	**extensive**
Available BW [priority]	(Must include **diffserv-te** statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	**extensive**
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	**extensive**

Table 52: show ted database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Available BW [TE-class]	(Must include the **diffserv-te** statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	**extensive**
Static BW [CT-class]	Total interface bandwidth used by an MPLS traffic class, in bps.	**extensive**

Interface Switching Capability Descriptor (*n*)	Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. ***n*** is the index number.
	• **Switching type**—Type of switching to be performed on a particular link:
	  ◦ **PSC-1**—Packet switch-capable 1
	  ◦ **PSC-2**—Packet switch-capable 2
	  ◦ **PSC-3**—Packet switch-capable 3
	  ◦ **PSC-4**—Packet switch-capable 4
	  ◦ **L2SC**—Layer-2-switch-capable
	  ◦ **TDM**—Time-division-multiplexing-capable
	  ◦ **LSC**—Lambda switch-capable
	  ◦ **FSC**—Fiber switch-capable
	• **Encoding type**—Encoding of the LSP being requested:
	  ◦ **Packet**
	  ◦ **Ethernet**
	  ◦ **ANSI/ETSI PDH**
	  ◦ **Reserved**
	  ◦ **SDH /SONET**
	  ◦ **Digital Wrapper**
	  ◦ **Lambda (photonic)**
	  ◦ **Fiber**
	  ◦ **FiberSDH/SONET**
	• **Maximum LSP BW [priority] bps**—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface.
	  ◦ **[*n*]**—Priority level. The range is from **0** (high) through **7** (low).
	  ◦ ***n* Mbps**—Amount of the maximum bandwidth.
	• **Minimum LSP BW**—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. **Minimum LSP BW** is displayed only when **switching type** is **PSC-1** or **TDM**.
	• **Interface MTU**—Displayed only when **switching type** is **TDM**.
	• **Interface supports standard SONET/SDH**—Displayed only when **switching type** is **TDM**.
		extensive

Sample Output

show ted database brief

`user@host>` **show ted database brief**

```
TED database: 12 ISIS nodes 0 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---    3178     2      0
Router-B.00                      ---    3152     2      0
Router-B.02                      Net     802     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---    3126     2      0
Router-C.02                      Net      38     0      2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---    3144     2      0
Router-D.02                      Net     723     0      2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net     607     0      2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---    3178     2      0
Router-E.02                      Net     131     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---    3153     2      0
Router-F.02                      Net     769     0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

show ted database detail

```
TED database: 12 ISIS nodes 0 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---    2913     2      0
Router-B.00                      ---    2887     2      0
Router-B.02                      Net     537     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---    2861     2      0
Router-C.02                      Net     597     0      2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
```



```
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---    2879     2      0
Router-D.02                      Net     458     0      2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net     342     0      2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---    2913     2      0
Router-E.02                      Net     640     0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---    2888     2      0
Router-F.02                      Net     504     0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```

show ted database extensive

```
user@host> show ted database extensive
TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
  Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
  Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
  Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
  Protocol: IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
NodeID: Router-C.00
  Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
```

```
    Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2
    Protocol: IS-IS(2)
      To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
          Maximum LSP BW [priority] bps:
            [0] 0bps        [1] 0bps        [2] 0bps        [3] 0bps
            [4] 0bps        [5] 0bps        [6] 0bps        [7] 0bps
      To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
          Maximum LSP BW [priority] bps:
            [0] 0bps        [1] 0bps        [2] 0bps        [3] 0bps
            [4] 0bps        [5] 0bps        [6] 0bps        [7] 0bps
  NodeID: Router-D.00
    Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
  NodeID: Router-D.02
    Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
    Protocol: IS-IS(2)
      To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
          Maximum LSP BW [priority] bps:
            [0] 0bps        [1] 0bps        [2] 0bps        [3] 0bps
            [4] 0bps        [5] 0bps        [6] 0bps        [7] 0bps
      To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
          Maximum LSP BW [priority] bps:
            [0] 0bps        [1] 0bps        [2] 0bps        [3] 0bps
            [4] 0bps        [5] 0bps        [6] 0bps        [7] 0bps
  NodeID: Router-D.03
    Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
    Protocol: IS-IS(2)
      To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
          Maximum LSP BW [priority] bps:
            [0] 0bps        [1] 0bps        [2] 0bps        [3] 0bps
            [4] 0bps        [5] 0bps        [6] 0bps        [7] 0bps
      To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        Interface Switching Capability Descriptor(1):
          Switching type: Packet
          Encoding type: Packet
```



```
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
NodeID: Router-E.00
  Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
  Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
  Protocol: IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
NodeID: Router-F.00
  Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
  Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
  Protocol: IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps         [1] 0bps        [2] 0bps        [3] 0bps
          [4] 0bps         [5] 0bps        [6] 0bps        [7] 0bps
```

show ted link

List of Syntax

Syntax on page 428
Syntax (EX Series Switches) on page 428

Syntax

```
show ted link
<brief | detail>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show ted link
<brief | detail>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.

Options

none—Display standard information about traffic engineering database link information.

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

List of Sample Output

show ted link brief on page 429
show ted link detail on page 429

Output Fields

Table 53 on page 428 describes the output fields for the **show ted link** command. Output fields are listed in the approximate order in which they appear.

Table 53: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**brief**
-->ID	Hostname and address of the node that the link is going to. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**brief**
hostname	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**detail**

Table 53: show ted link Output Fields *(continued)*

Field Name	Field Description	Level of Output
hostname	Hostname and address of the node that the link is going to. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**detail**
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	**detail extensive**
Remote	Address of the interface on the remote node.	**detail extensive**
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	**detail**
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	**detail**

Sample Output

show ted link brief

```
user@host> show ted link brief
ID                          ->ID                         LocalPath LocalBW
Router-B.02                     Router-A.00                          0 0bps
Router-B.02                     Router-B.00                          0 0bps
Router-C.02                     Router-B.00                          0 0bps
Router-C.02                     Router-C.00                          0 0bps
Router-D.02                     Router-F.00                          0 0bps
Router-D.02                     Router-D.00                          0 0bps
Router-D.03                     Router-D.00                          0 0bps
Router-D.03                     Router-C.00                          0 0bps
Router-E.02                     Router-A.00                          0 0bps
Router-E.02                     Router-E.00                          0 0bps
Router-F.02                     Router-E.00                          0 0bps
Router-F.02                     Router-F.00                          0 0bps
```

show ted link detail

```
user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
  localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
  localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
```

```
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
  Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    LocalPath: 0, Metric: 0, AvailBW: 0bps
    localBW [0] 0bps  [1] 0bps  [2] 0bps  [3] 0bps
    localBW [4] 0bps  [5] 0bps  [6] 0bps  [7] 0bps
```


show ted protocol

List of Syntax

Syntax on page 431
Syntax (EX Series Switches) on page 431

Syntax

show ted protocol
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show ted protocol
<brief | detail>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.

Options

none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.

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

List of Sample Output

show ted protocol on page 432

Output Fields

Table 54 on page 431 describes the output fields for the **show ted protocol** command. Output fields are listed in the approximate order in which they appear.

Table 54: show ted protocol Output Fields

Field Name	Field Description
Protocol name	Protocol that reported the node information: • **IS-IS(1)**—IS-IS Level 1. • **IS-IS(2)**—IS-IS Level 2. • **OSPF** (***area-number***)—OSPF from the specified area.
Credibility	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
Self node	Address the protocol uses as the local address.

Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name        Credibility  Self node
IS-IS(2)             2  (highest) corriedale.00(123.456.1.11)
IS-IS(1)             1            corriedale.00(123.456.1.11)
```


CHAPTER 7

RSVP

Using RSVP

Understanding MPLS Components

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

IP Over MPLS for Customer Edge Interfaces

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “Configuring MPLS on Provider Edge Switches” on page 179.

BGP Layer 3 VPN Configuration

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

Routing Instances for Layer 3 VPN

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- Interior Gateway Protocol on page 435
- MPLS Protocol on page 435
- RSVP on page 435
- Family mpls on page 436

Interior Gateway Protocol

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.

NOTE: The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to **family mpls**, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

RSVP

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

Family mpls

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.

NOTE: **You can enable family mpls on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.**

Related Documentation

- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 160
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156
- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183
- Configuring Rewrite Rules for MPLS EXP Classifiers on page 187
- Configuring a Global MPLS EXP Classifier on page 175
- Configuring Ethernet over MPLS (L2 Circuit) on page 170
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

RSVP Overview

The RSVP protocol is used by routers to deliver quality-of-service (QoS) requests to all nodes along data flow path(s) and to establish and maintain state for the requested service. RSVP requests generally result in resource reservations in each node along the data path. RSVP has the following attributes:

- Makes resource reservations for unidirectional data flows.
- Allows the receiver of a data flow to initiate and maintain the resource reservation used for that flow, as shown in Figure 16 on page 437.
- Maintains a soft state in routers and hosts, providing graceful support for dynamic membership changes and automatic adaptation to routing changes.
- Depends upon present and future routing protocols, but is not a routing protocol itself.

- Provides several reservation models or styles to fit a variety of applications.
- Supports both IPv4 and IPv6. Note, you can configure the Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. For more information, see the *Junos OS MPLS Applications Library for Routing Devices*.

Figure 16: RSVP Reservation Request and Data Flow

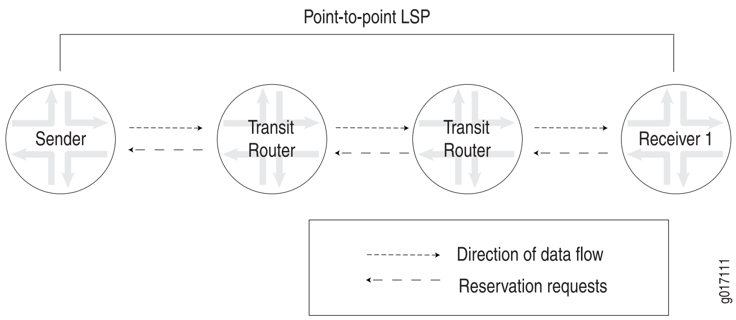

MTU Signaling in RSVP

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value.

For information about how to configure this feature, see *Configuring MTU Signaling in RSVP*.

The following sections describe how MTU signaling in RSVP works:

- *How the Correct MTU Is Signaled in RSVP*
- *Determining an Outgoing MTU Value*
- *MTU Signaling in RSVP Limitations*

Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- Tunneling LDP LSPs in RSVP LSPs Overview on page 5
- Label Operations on page 5

Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.

Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. Configuring the Ingress PE Switch on page 439
2. Configuring the Egress PE Switch on page 440

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

   [edit interfaces]
   user@switch# **set lo0 unit 0 family inet address 192.168.10.1/32**
   user@switch# **set xe-0/0/5 unit 0 family inet address 10.1.5.1/24**
   user@switch# **set xe-0/0/6 unit 0 family inet address 10.1.6.1/24**

   NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:

   NOTE: You can use the switch address as an alternative to the loopback interface.

   [edit protocols ospf]
   user@switch# **set area 0.0.0.0 interface lo0.0**
   user@switch# **set area 0.0.0.0 interface xe-0/0/5.0**
   user@switch# **set area 0.0.0.0 interface xe-0/0/6.0**
3. Configure OSPF traffic engineering:

   [edit protocols ospf]
   user@switch# **set traffic-engineering**
4. Configure RSVP on the loopback interface and the core interfaces:

   [edit protocols rsvp]
   user@switch# **set interface lo0.0**
   user@switch# **set interface xe-0/0/5.0**
   user@switch# **set interface xe-0/0/6.0**
5. Configure MPLS traffic engineering.

   [edit protocols mpls]
   user@switch# **set traffic-engineering**
6. Configure MPLS on the core interfaces:

   [edit protocols mpls]
   user@switch# **set interface xe-0/0/5.0**
   user@switch# **set interface xe-0/0/6.0**
7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

   [edit interfaces]
   user@switch# **set xe-0/0/5 unit 0 family mpls**
   user@switch# **set xe-0/0/6 unit 0 family mpls**
8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

   [edit interfaces]
   user@switch# **set xe-0/0/3 unit 0 family inet address 121.100.10.1/16**
9. Configure this Layer 3 customer edge interface for the routing protocol:

   [edit]

```
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0
```

10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1
```

11. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf
```

12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve
```

Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24
```

NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:

NOTE: You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0
user@switch# set rsvp interface xe-0/0/6.0
```

4. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
```



```
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch#  set static route 121.121.121.0/24 next-hop 192.168.12.1
user@switch# set static route 121.121.121.0/24 resolve
```

Related Documentation

- MPLS Configuration Guidelines on page 217
- Configuring MPLS on Provider Switches on page 183
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding MPLS Components on page 144
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156

Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:

   NOTE: You can use the switch address as an alternative to the loopback interface.

   ```
   [edit protocols ospf]
   user@switch# set area 0.0.0.0 interface lo0.0
   user@switch# set area 0.0.0.0 interface xe-0/0/5.0
   user@switch# set area 0.0.0.0 interface xe-0/0/6.0
   user@switch# set area 0.0.0.0 interface ae0
   ```

   NOTE: You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

   ```
   [edit protocols mpls]
   user@switch# set interface xe-0/0/5.0
   user@switch# set interface xe-0/0/6.0
   user@switch# set interface ae0
   ```

3. Configure RSVP on the loopback interface and the core interfaces:

   ```
   [edit protocols rsvp]
   user@switch# set interface lo0.0
   user@switch# set interface xe-0/0/5.0
   user@switch# set interface xe-0/0/6.0
   user@switch# set interface ae0
   ```

4. Configure an IP address for the loopback interface and the core interfaces:

   ```
   [edit interfaces]
   user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
   user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
   user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
   user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
   ```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

   ```
   [edit interfaces]
   user@switch# set xe-0/0/5 unit 0 family mpls
   user@switch# set xe-0/0/6 unit 0 family mpls
   user@switch# set ae0 unit 0 family mpls
   ```

NOTE: You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- MPLS Configuration Guidelines on page 217
- MPLS Feature Support on QFX Series and EX4600 Switches on page 138
- Understanding MPLS Components on page 144
- Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 156

Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. Verifying the Physical Layer on the Switches on page 443
2. Verifying the Routing Protocol on page 444
3. Verifying the Core Interfaces Being Used for the MPLS Traffic on page 444
4. Verifying RSVP on page 444

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action

```
user@switch> show interfaces xe-* terse

Interface               Admin Link Proto    Local                 Remote
xe-0/0/0                up    up
xe-0/0/0.0              up    up
xe-0/0/1.0              up    up
xe-0/0/2.0              up    up
xe-0/0/3.0              up    up   inet       2.2.2.1/16
xe-0/0/4.0              up    up
xe-0/0/5.0              up    up   inet      10.1.5.1/24
                                   mpls
xe-0/0/6.0              up    up   inet      10.1.6.1/24
                                   mpls
```

Meaning The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

Action `user@switch>` **show ospf neighbor**

```
Address          Interface              State     ID               Pri  Dead
127.1.1.1        xe-0/0/5               Full      10.10.10.10      128    39
```

Meaning The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the *Junos OS Routing Protocols and Policies Command Reference* .

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

Action `user@switch>` **show mpls interface**

```
Interface        State       Administrative groups
ge-0/0/5            Up          <none>
ge-0/0/6            Up          <none>
```

Meaning The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

Verifying RSVP

Purpose Verify the state of the RSVP session. You should perform this verification task on each of the switches.

Action

```
user@switch> show rsvp session

Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.3       127.1.1.1         Up       0  1 FF       -    300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.1       127.1.1.3         Up       0  1 FF   299968        - lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are up.

Related Documentation

- Configuring MPLS on Provider Edge Switches on page 179
- Configuring MPLS on Provider Switches on page 183

CHAPTER 8

Configuration Statements for RSVP

admin-group

Syntax

```
admin-group {
  exclude [ group-names ];
  include-all [ group-names ];
  include-any [ group-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.

Options

exclude *group-names*—Specify the administrative groups to exclude for a bypass LSP.

include-all *group-names*—Specify the administrative groups whose links the bypass LSP must traverse.

include-any *group-names*—Specify the administrative groups whose links the bypass LSP can traverse.

Required Privilege Level

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

Related Documentation

- *Configuring Link Protection on Interfaces Used by LSPs*

authentication-key (Protocols RSVP)

Syntax authentication-key *key*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp peer-interface *peer-interface-name*],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp peer-interface *peer-interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.

RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

All routers that are connected to the same IP subnet must use the same authentication scheme and password.

Options *key*—Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").

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

Related Documentation
- *Configuring RSVP Interfaces*

aggregate (Protocols RSVP)

Syntax (aggregate | no-aggregate);

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp peer-interface *peer-interface-name*],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp peer-interface *peer-interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Control the use of RSVP aggregate messages on an interface or peer interface:

- **aggregate**—Use RSVP aggregate messages.
- **no-aggregate**—Do not use RSVP aggregate messages.

  Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.

  Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.

To have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.

Default Aggregation is disabled.

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

Related Documentation

- *Configuring RSVP Interfaces*
- reliable on page 479

bandwidth (Protocols RSVP)

Syntax
bandwidth *bps*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description
For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.

This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.

Default
The hardware raw bandwidth is used.

Options
bps—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion [also called a thousand million]).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*
- *Configuring Link Protection on Interfaces Used by LSPs*
- *Configuring Link Protection on Interfaces Used by LSPs*

bypass (Signaled LSP)

Syntax

bypass *bypass-name* {
  bandwidth *bps*;
  description *text*;
  hop-limit *number*;
  no-cspf;
  path *address* <strict | loose>;
  priority *setup-priority reservation-priority*;
  to *address*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection]

Release Information

Statement introduced before Junos OS Release 7.4.
The **description** option was added in Junos OS Release 10.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the **[edit protocols rsvp interface *interface-name* link-protection]** hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

Options

bypass-name—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.

description—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show mpls lsp bypass detail** command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.

to *address*—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*

class-of-service (Protocols RSVP)

Syntax class-of-service *cos-value*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options *cos-value*—CoS value. A higher value typically corresponds to a higher level of service.

Range: 0 through 7

Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*

disable (Protocols RSVP)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit logical-systems *logical-system-name* protocols rsvp graceful-restart],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp peer-interface *peer-interface-name*],
[edit protocols rsvp],
[edit protocols rsvp graceful-restart],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp peer-interface *peer-interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.

Default RSVP is enabled on interfaces and peer interfaces configured with the RSVP **interface** statement. RSVP graceful restart is enabled on the router. Link protection is disabled.

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

Related Documentation

- *Minimum RSVP Configuration*
- *Configuring RSVP Graceful Restart*
- *Configuring Link Protection on Interfaces Used by LSPs*

fast-reroute (Protocols RSVP)

Syntax fast-reroute optimize-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement added in Junos OS Release 7.5.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15

Description Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

Options ***seconds***—Specify the number of seconds between fast reroute detour LSP optimizations.

Range: 0 through 65,535 seconds

Default: 0 (disabled)

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

Related Documentation
- *Configuring the Optimization Interval for Fast Reroute Paths*

graceful-deletion-timeout

Syntax graceful-deletion-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Specify the time, in seconds, before completing graceful deletion of signaling.

Options ***seconds***—Time before completing graceful deletion of signaling.

Range: 1 through 300 seconds

Default: 30 seconds

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

Related Documentation
- *Gracefully Tearing Down GMPLS LSPs*

graceful-restart (Enabling Globally)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-helper-recovery-time seconds;
  maximum-helper-restart-time seconds;
  notify-duration seconds;
  recovery-time seconds;
  restart-duration seconds;
  stale-routes-time seconds;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. 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.

Required Privilege Level

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

Related Documentation

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

hello-acknowledgements

Syntax

hello-acknowledgements;

Hierarchy Level

[edit logical-systems *logical-systems-name* protocols rsvp],
[edit protocols rsvp]

Release Information

Statement introduced in Junos OS Release 10.2.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.

Default

Hello acknowledgments are disabled.

Required Privilege Level

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

Related Documentation

- *Configuring Hello Acknowledgments for Nonsession RSVP Neighbors*

hello-interval (Protocols RSVP)

Syntax hello-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp peer-interface *peer-interface-name*],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp peer-interface *peer-interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Enable the sending of hello packets on the interface.

Options *seconds*—Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface.

Range: 1 through 60 seconds

Default: 9 seconds

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

Related Documentation

- *Configuring RSVP Interfaces*

helper-disable (Multiple Protocols)

Syntax
helper-disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ldp | ospf | ospf3) graceful-restart],
[edit protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],
[edit routing-instances *routing-instance-name* protocols (ldp | ospf | ospf3) graceful-restart]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description
Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.

Default
Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.

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

Related Documentation
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*

hop-limit

Syntax hop-limit *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols mpls],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls label-switched-path *lsp-name* fast-reroute],
[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:

- LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.
- Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
- Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.

Options *number*—Maximum number of hops.

Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)

Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)

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

Related Documentation
- *Configuring Fast Reroute*
- *Limiting the Number of Hops in LSPs*

- *Configuring Link Protection on Interfaces Used by LSPs*

interface (Protocols RSVP)

Syntax

```
interface interface-name {
  disable;
  (aggregate | no-aggregate);
  authentication-key key;
  bandwidth bps;
  hello-interval seconds;
  link-protection {
    disable;
    admin-group {
      exclude [ group-names ];
      include-all [ group-names ];
      include-any [ group-names ];
    }
    bandwidth bps;
    bypass bypass-name {
      bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
      }
      description text;
      class-of-service cos-value;
      hop-limit number;
      no-cspf;
      path address <strict | loose>;
      priority setup-priority reservation-priority;
      to address;
    }
    class-of-service cos-value;
    hop-limit number;
    max-bypasses number;
    no-cspf;
    no-node-protection;
    optimize-timer seconds;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    subscription percentage;
  }
  (reliable | no-reliable);
  subscription percentage {
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
  }
  update-threshold threshold;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Enable RSVP on one or more router interfaces.

Default RSVP is disabled on all interfaces.

Options ***interface-name***—Name of an interface. To configure all interfaces, specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

The remaining statements are explained separately.

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

Related Documentation
- *Minimum RSVP Configuration*

keep-multiplier

Syntax keep-multiplier *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Set the keep multiplier value.

Options ***number***—Multiplier value.

Range: 1 through 255

Default: 3

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

Related Documentation
- *Configuring Timers for RSVP Refresh Messages*

link-protection (RSVP)

Syntax

```
link-protection {
  disable;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth bps;
  bypass bypass-name {
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
    description text;
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    to address;
  }
  class-of-service cos-value;
  hop-limit number;
  max-bypasses number;
  no-cspf;
  no-node-protection;
  optimize-timer seconds;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  subscription percentage;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the **link-protection** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level. You can configure single or multiple bypasses for protected interface.

Default

Link protection is disabled.

Options

no-node-protection—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Link Protection on Interfaces Used by LSPs*
- *link-protection (Dynamic LSPs)*

load-balance (Protocols RSVP)

Syntax

```
load-balance {
  bandwidth;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Load-balance traffic between RSVP LSPs.

Options

bandwidth—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.

Required Privilege Level

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

Related Documentation

- *Configuring Load Balancing Across RSVP LSPs*

max-bypasses

Syntax max-bypasses *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Range modified in Junos OS Release 9.3.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure **max-bypasses**, you must also configure the bandwidth statement.

Options ***number***—Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured.

Range: 0 through 99

Default: 1

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*

maximum-helper-recovery-time

Syntax maximum-helper-recovery-time *seconds*;

Hierarchy Level [edit protocols rsvp graceful-restart],
[edit logical-systems *logical-system-name* protocols rsvp graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

Options *seconds*—Legnth of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

Range: 1 through 3600

Default: 180

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

Related Documentation

- *Configuring Graceful Restart Options for RSVP, CCC, and TCC*
- maximum-helper-restart-time (RSVP) on page 469

maximum-helper-restart-time (RSVP)

Syntax maximum-helper-restart-time *seconds*;

Hierarchy Level [edit protocols rsvp graceful-restart],
[edit logical-systems *logical-system-name* protocols rsvp graceful-restart]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.

Options *seconds*—The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down.

Range: 1 through 1800

Default: 60

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

Related Documentation

- *Configuring Graceful Restart Options for RSVP, CCC, and TCC*
- maximum-helper-recovery-time on page 468

no-cspf (Protocols RSVP)

Syntax no-cspf;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information Statement introduced in Junos OS Release 7.5.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.

Default CSPF is enabled.

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*

no-interface-hello

Syntax no-interface-hello;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced in JUNOS Release 10.0.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Allows you to explicitly disable RSVP interface hellos globally on the router. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the hello-interval (Protocols RSVP) statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node ID hellos and other devices support RSVP interface hellos.

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

Related Documentation

- *Configuring RSVP Node ID Hellos*
- hello-interval (Protocols RSVP) on page 459

no-local-reversion

Syntax no-local-reversion;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Disables RSVP local revertive mode as specified in RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. RSVP local revertive mode is supported on all Juniper Networks routers running the Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.

The following information can also be found in RFC 4090. Refer to the full RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:

- Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resignaled along the restored path is switched back.
- Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.

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

no-node-id-subobject

Syntax no-node-id-subobject;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Disable the record route object (RRO) node ID subobject for compatibility with earlier versions of the Junos OS. To interoperate with other vendors' equipment, the Junos OS supports the RRO node ID subobject for use in inter-AS link and node protection configurations.

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

Related Documentation
- *Configuring Inter-AS Node and Link Protection*

no-p2mp-sublsp

Syntax no-p2mp-sublsp;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the **no-p2mp-sublsp** statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.

Default Resv messages that include the S2L_SUB_LSP object are accepted.

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

Related Documentation
- *Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases*

node-hello

Syntax
node-hello;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information
Statement introduced in JUNOS Release 10.0.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description
Enables node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, the JUNOS Software uses interface-based RSVP hellos and node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, the JUNOS software does not accept any node-ID hello packets.

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

Related Documentation
- *Configuring RSVP Node ID Hellos*

optimize-timer (Protocols RSVP)

Syntax optimize-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.

Options ***seconds***—Specify the number of seconds between optimizations.

Range: 0 through 65,535 seconds

Default: 0 (disabled)

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

Related Documentation

- *Configuring Link Protection on Interfaces Used by LSPs*

path (Protocols RSVP)

Syntax path *address* <strict | loose>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.

Default No path is configured. CSPF automatically calculates the path the bypass LSP takes.

Options *address*—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is **loose**. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.

loose—(Optional) The next address in the **path** statement is loose. The LSP can traverse other routers before reaching this router.

Default: **strict**

strict—(Optional) The LSP must go to the next address specified in the **path** statement without traversing other nodes. This is the default.

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

Related Documentation

- *Configuring Link Protection on Interfaces Used by LSPs*

preemption

Syntax

```
preemption {
  (aggressive | disabled | normal);
  soft-preemption {
    cleanup-timer seconds;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Control RSVP session preemption.

Default

normal

Options

aggressive—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.

disabled—Do not preempt RSVP sessions.

normal—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Preempting RSVP Sessions*

priority (Protocols RSVP)

Syntax priority *setup-priority reservation-priority*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.

Options ***reservation-priority***—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 0 (Once the session is set up, no other session can preempt it.)

setup-priority—Setup priority.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 7 (The session cannot preempt any existing sessions.)

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

Related Documentation
- *Configuring Link Protection on Interfaces Used by LSPs*
- *Configuring Priority and Preemption for LSPs*

refresh-time

Syntax refresh-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Set the refresh time.

Options *seconds*—Refresh time.

Range: 1 through 65,535

Default: 30 seconds

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

Related Documentation

- *Configuring Timers for RSVP Refresh Messages*

reliable

Syntax (reliable | no-reliable);

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp peer-interface *peer-interface-name*],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp peer-interface *peer-interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Enable reliable message delivery on the interface.

In order to have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.

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

Related Documentation

- *Configuring RSVP Interfaces*
- aggregate on page 451

setup-protection

Syntax setup-protection;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Description The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the **setup-protection** statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.

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

Related Documentation

- *Configuring RSVP Setup Protection*

subscription

Syntax

```
subscription percentage {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name* link-protection]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). **subscription** is the percentage of the link bandwidth that can be used for the RSVP reservation process.

Options

ct*number percentage*—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.

Range: 0 through 65,000

Default: 100 percent

percentage—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.

Range: 0 through 65,000

Default: 100 percent

Required Privilege Level

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

Related Documentation

- *Configuring the Bandwidth Subscription Percentage for LSPs*
- *Configuring Link Protection on Interfaces Used by LSPs*

soft-preemption (Protocols RSVP)

Syntax

```
soft-preemption {
  cleanup-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp preemption],
[edit protocols rsvp preemption]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.

Options

cleanup-timer—A value of 0 disables soft preemption.

Range: 0 through 180 seconds

Default: 30 seconds

Required Privilege Level

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

Related Documentation

- *Configuring MPLS Soft Preemption*

traceoptions (Protocols RSVP)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable RSVP-level trace options.

Default

The default RSVP-level trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

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**.

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 RSVP tracing output in the file **rsvp-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.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **all**—All tracing operations
- **error**—All detected error conditions
- **event**—RSVP-related events
- **lmp**—RSVP-LMP interactions
- **packets**—All RSVP packets
- **path**—All path messages

- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Enable only certain users to read 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 this 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.

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: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

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

Required Privilege Level

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

Related Documentation

- *Tracing RSVP Protocol Traffic*

tunnel-services (RSVP)

Syntax

```
tunnel-services {
  devices device-names;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced for QFX switches in Junos OS Release13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.

Default

Ultimate-hop popping is disabled.

Options

devices *device-names*—Specify which VT interfaces are used to handle the RSVP traffic.

Range: 0 to 8 devices

Required Privilege Level

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

Related Documentation

- *Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs*

update-threshold

Syntax update-threshold *threshold*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.

Options ***threshold***—Specify the percentage change in bandwidth to trigger an IGP update.

Range: 1 through 20 percent

Default: 10 percent

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

Related Documentation

- *Configuring RSVP Interfaces*

CHAPTER 9

Monitoring Commands for RSVP

clear rsvp session

List of Syntax

Syntax on page 488
Syntax (EX and QFX Series Switches) on page 488

Syntax

clear rsvp session
<connection-destination *address*>
<connection-source *address*>
<gracefully>
<logical-system (all | *logical-system-name*)>
<lsp-id *identifier*>
<name *name*>
<optimize-fast-reroute>
<tunnel-id *identifier*>

Syntax (EX and QFX Series Switches)

clear rsvp session
<connection-destination *address*>
<connection-source *address*>
<gracefully>
<lsp-id *identifier*>
<name *name*>
<optimize-fast-reroute>
<tunnel-id *identifier*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options

none—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.

connection-source *address*—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.

connection-destination *address*—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.

gracefully—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.

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

lsp-id *identifier*—(Optional) LSP identifier (source port) for the RSVP sender template.

name *name*—(Optional) Reset and restart the specified RSVP session.

optimize-fast-reroute—(Optional) Begin fast reroute optimization.

tunnel-id ***identifier***—(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level

clear

Related Documentation

- clear mpls lsp on page 325
- show rsvp session on page 506

List of Sample Output

clear rsvp session on page 489

Output Fields

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

Sample Output

clear rsvp session

`user@host>` **clear rsvp session**

clear rsvp statistics

List of Syntax

Syntax on page 490
Syntax (EX Series Switches) on page 490

Syntax

```
clear rsvp statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
clear rsvp statistics
```

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Clear Resource Reservation Protocol (RSVP) packet and error statistics.

Options

none—Clear RSVP packet and error statistics.

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

Required Privilege Level

clear

Related Documentation

- show rsvp statistics on page 515

List of Sample Output

clear rsvp statistics on page 490

Output Fields

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

Sample Output

clear rsvp statistics

```
user@host> clear rsvp statistics
```


ping mpls rsvp

Syntax

ping mpls rsvp
<*lsp-name*>
<count *count*>
<destination *address*>
<detail>
<dynamic-bypass>
<egress *egress-address*>
<exp *forwarding-class*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<manual-bypass>
<multipoint>
<size *bytes*>
<source *source-address*>
<standby *standby-path-name*>
<sweep>

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.
The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.
The **size** and **sweep** options were introduced in Junos OS Release 9.6.
The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.

Description

Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

NOTE: When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

dynamic-bypass—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress ***egress-address***—(Optional) Only the specified egress router or switch responds to the ping request.

exp ***forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

interface—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

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

lsp-name—Ping an RSVP-signaled LSP using an LSP name.

manual-bypass—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

multipoint—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size ***bytes***—(Optional) Size of the LSP ping request packet (**100** through **65468** bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

source ***source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

standby ***standby-path-name***—(Optional) Name of the standby path.

sweep —(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
 loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
 Local transmit time: 1205310695s 215737us
 Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 ms
 Local transmit time: 1205310615s 347317us
```



```
 Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
 Local transmit time: 1205310615s 347262us
 Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
 Local transmit time: 1205310615s 347167us
 Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

ping mpls rsvp destination detail count size

```
user@host>ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
        Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
        Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp destination detail sweep size

```
user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
        Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
        Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
        Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
        Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
        Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
        Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
        Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
        Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
```



```
        Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
        Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
        Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
        Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
        Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
        Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
        Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
        Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
        Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
        Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

show rsvp interface

List of Syntax

Syntax on page 496
Syntax (EX Series Switches) on page 496

Syntax

```
show rsvp interface
<brief | detail | extensive>
<link-management>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp interface
<brief | detail | extensive>
<link-management>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.

Options

none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.

brief | detail | extensive | link-management—(Optional) Display the specified level of output.

link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).

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

Required Privilege Level

view

List of Sample Output

show rsvp interface brief on page 499
show rsvp interface detail on page 499
show rsvp interface extensive on page 499
show rsvp interface link-management on page 500

Output Fields

Table 55 on page 497 lists the output fields for the **show rsvp interface** command. Output fields are listed in the approximate order in which they appear.

Table 55: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
RSVP interface	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
Interface	Name of the interface.	All levels
Index	Index of the interface.	**detail**
State	State of the interface. • **Disabled**—No traffic engineering information is displayed. • **Down**—Interface is not operational. • **Enabled**—Displays traffic engineering information. • **Up**—Interface is operational.	All levels
NoAuthentication	Interface does not support RSVP authentication.	**detail**
NoAggregate	Interface does not support refresh reduction.	**detail**
NoReliable	Interface does not support refresh reduction message ID extension.	**detail**
NoLinkProtection	Interface does not support link protection.	**detail**
HelloInterval	Frequency at which RSVP hellos are sent on this interface (in seconds).	**detail**
Address	IP address of the local interface.	**detail**
Active control channel	Next-hop link address to transmit messages.	None specified
TElink	Traffic-engineered links that are managed by the peer they are associated with.	None specified
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	All levels
PreemptionCnt	Number of times an RSVP session was preempted on this interface.	**detail**
Update threshold	Percentage change in reserved bandwidth to trigger an IGP update.	**detail**
Subscription	User-configured subscription factor.	All levels
bc *number*	Bandwidth allocated for the specified bandwidth constraint.	**extensive**
ct *number*	Bandwidth allocated for the specified class type.	**extensive**
Static BW	Total interface bandwidth, in bps.	All levels
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels

Table 55: show rsvp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Reserved BW	Currently reserved bandwidth, in bps.	All levels
SoftPreemptionCnt	Number of times a soft preemption occurred on this interface. This number is not included in the **PreemptionCnt** value.	**detail**
Overbooked BW	Currently overbooked bandwidth, in bps, by class type (**ct0** through **ct3**).	**detail**
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bps.	**brief**
PacketType	Type of RSVP packet.	**detail**
Total Sent	Total number of packets sent.	**detail**
Total Received	Total number of packets received since RSVP was enabled.	**detail**
Last 5 seconds Sent	Number of packets sent in the last 5 seconds.	**detail**
Last 5 seconds Received	Number of packets received in the last 5 seconds.	**detail**
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	**detail**
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	**detail**
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	**detail**
Resv	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	**detail**
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	**detail**
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.	**detail**
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	**detail**
Ack	Acknowledge message for refresh reductions.	**detail**
Srefresh	Summary refresh messages.	**detail**
EndtoEnd RSVP	Statistics for the number of end-to-end RSVP messages sent.	**detail**
Queue	CoS transmit queue number and its associated forwarding class designation.	**extensive**

Table 55: show rsvp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
TxRate	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	extensive
Priority	Weight of the queue relative to other configured queues, in percentage.	extensive
queue-priority-value	**Low**, **High**, **None**, or **Exact**. **None** indicates no rate limiting. **Exact** indicates the queue transmits at the configured rate only.	extensive

Sample Output

show rsvp interface brief

```
user@host> show rsvp interface brief
RSVP interface: 1 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface    State resv   iption  BW          BW          BW          mark
de0.0      Up          1     23%     10Mbps      989.992kbps 1.31Mbps    1.31Mbps
```

show rsvp interface detail

```
user@host> show rsvp interface detail
so-0/1/1.0  Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
SoftPreemptionCnt1
OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
PacketType              Total                     Last 5 seconds
                   Sent       Received          Sent       Received
  Path               16              0             1              0
  PathErr             0              0             0              0
  PathTear            1              0             0              0
  Resv                0             11             0              1
  ResvErr             0              0             0              0
  ResvTear            0              0             0              0
  Hello              66             67             1              1
  Ack                 0              0             0              0
  Srefresh            0              0             0              0
  EndtoEnd RSVP       0              0             0              0
...
```

show rsvp interface extensive

```
user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
```

```
bc2 = (ct2+ct3), StaticBW 311.04Mbps
bc3 = ct3, StaticBW 155.52Mbps
ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Queue         TxRate            Priority Exact
  0           155.52Mbps        25%      Low
  1           155.52Mbps        25%      Low
  2           155.52Mbps        25%      Low
  3           155.52Mbps        25%      Low
```

show rsvp interface link-management

```
user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C  State: Up
Active Control Channel: so-0/1/0.0

  TElink: TElnk1, Link ID: 37811
  ActiveResv 0, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

  TElink: TElnk2, Link ID: 37808
  ActiveResv 1, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B  State: Up
Active Control Channel: so-1/0/0.0

  TElink: TElnkAB1, Link ID: 1598
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

  TElink: TElnkAB2, Link ID: 1597
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```


show rsvp neighbor

List of Syntax

Syntax on page 501
Syntax (EX Series Switches) on page 501

Syntax

show rsvp neighbor
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show rsvp neighbor
<brief | detail>

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.

Options

none—Display standard information about RSVP neighbors.

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

List of Sample Output

show rsvp neighbor on page 505
show rsvp neighbor detail on page 505

Output Fields

Table 56 on page 501 lists the output fields for the **show rsvp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 56: show rsvp neighbor Output Fields

Field Name	Field Description	Level of Output
RSVP neighbor	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
via	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	**detail**
Address	Address of a learned neighbor.	All levels
Idle	Length of time the neighbor has been idle, in seconds.	All levels

Table 56: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Up/Dn	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
Up cnt and Down cnt	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	**detail**

status	State of the RSVP neighbor:
	• **Up**—Routing device can detect RSVP Hello messages from the neighbor.
	• **Down**—Routing device has received one of the following indications:
	  ◦ Communication failure from the neighbor.
	  ◦ Communication from IGP that the neighbor is unavailable.
	  ◦ Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.
	• **Restarting**—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.
	• **Restarted**—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.
	• **Dead**—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.
		detail

LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is ***hh:mm:ss***.	All levels
Last changed time	Time elapsed since the neighbor state changed either from up to down or from down to up.	**detail**
HelloInt	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
HelloTx/Rx	Number of hello packets sent to and received from the neighbor.	All levels
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	**detail**
Message received	Number of Path and Resv messages that this routing device has received from the neighbor.	**detail**
Remote Instance	Identification provided by the remote routing device during Hello message exchange.	**detail**
Local Instance	Identification sent to the remote routing device during Hello message exchange.	**detail**

Table 56: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Refresh reduction	Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. **Refresh reduction** can have the following values: • **operational**—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • **incomplete**—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • **no operational**—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.	**detail**
Remote end	Neighboring routing device's status with regard to refresh reduction: • **enabled**—Remote routing device has requested refresh reduction during RSVP message exchanges. • **disabled**—Remote routing device does not require refresh reduction.	**detail**
Ack-extension	An RSVP refresh reduction extension: • **enabled**—Both local and remote routing devices support the ack-extension (RFC 2961). • **disabled**—Remote routing device does not support the ack-extension.	**detail**
Link protection	Status of the MPLS fast reroute mechanism that protects traffic from link failure: • **enabled**—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • **disabled**—No link protection feature has been enabled for this neighbor.	**detail**
LSP name	Name of the bypass LSP.	**detail**
Bypass LSP	Status of the bypass LSP. It can have the following values: • **does not exist**—Bypass LSP is not available. • **connecting**—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • **operational**—Bypass LSP is up and running. • **down**—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.	**detail**
Backup routes	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	**detail**
Backup LSPs	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	**detail**
Bypass explicit route	Explicit route object's (ERO) path that is taken by the bypass LSP.	**detail**

Table 56: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Restart time	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	**detail**
Recovery time	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	**detail**

Sample Output

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address            Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203       0  3/2       13:01        3   366/349
192.168.207.207       0  1/0       22:49        3   448/448
```

show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasyl status: Up
   Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
   Message received: 632
   Hello: sent 673, received 656, interval 3 sec
   Remote instance: 0x6432838a, Local instance: 0x74b72e36
   Refresh reduction:  operational
     Remote end: enabled,  Ack-extension: enabled
   Link protection:  enabled
     LSP name: Bypass_to_192.168.207.203
     Bypass LSP: operational,  Backup routes: 1,  Backup LSPs: 0
     Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

show rsvp session

List of Syntax

Syntax on page 506
Syntax (EX and QFX Series Switches) on page 506

Syntax

show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*lsp-type*>
<name *session-name*>
<p2mp>
<*session-type*>
<statistics>
<te-link *te-link*>

Syntax (EX and QFX Series Switches)

show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<interface *interface-name*>
<*lsp-type*>
<name *session-name*>
<p2mp>
<*session-type*>
<statistics>
<te-link *te-link*>

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
externally-provisioned option added in Junos OS Release 13.3.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Display information about Resource Reservation Protocol (RSVP) sessions.

Options

none—Display standard information about all RSVP sessions.

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

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

bypass—(Optional) Display RSVP sessions for bypass LSPs.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

externally-provisioned—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

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

lsp-type—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name *session-name*—(Optional) Display information about the named session.

p2mp—(Optional) Display point-to-multipoint information.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

statistics—(Optional) Display packet statistics.

te-link *te-link*—(Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level view

Related Documentation

- clear rsvp session on page 488

List of Sample Output

show rsvp session on page 511
show rsvp session statistics on page 511
show rsvp session detail on page 512
show rsvp session detail (Path MTU Output Field) on page 512
show rsvp session detail (GMPLS) on page 512
show rsvp session extensive on page 513
show rsvp session p2mp (Ingress Router) on page 513
show rsvp session p2mp (Transit Router) on page 514

Output Fields Table 57 on page 508 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 57: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	**detail**
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
P2MP name	(Appears only when the **p2mp** option is specified). Name of the point-to-multipoint LSP path.	All levels
P2MP branch count	(Appears only when the **p2mp** option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
From	Source (ingress routing device) of the session.	All levels
State	State of the path: **Up**, **Down**, or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress routing device) of the LSP.	**detail**
From	Source (ingress routing device) of the session.	**detail**
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either **Up**, **Dn** (down), or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	**brief detail**
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the routing table is the primary MPLS table **mpls.0**).	**brief**
Active Route	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (**mpls.0**).	**detail**
LSPname	Name of the LSP.	**brief detail**
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. **LSPpath** can be either **primary** or **secondary** and can be displayed on the ingress, egress, and transit routing devices. **LSPpath** can also indicate when a graceful LSP deletion has been triggered.	**detail**
Bypass	(Egress routing device) Destination address for the bypass LSP.	**detail**

Table 57: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Bidir	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	**detail**
Bidirectional	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	**detail**
Upstream label in	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	**detail**
Upstream label out	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	**detail**
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	**detail**
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	**detail**
Resv style or **Style**	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	**brief detail**
Label in	Incoming label for this LSP.	**brief detail**
Label out	Outgoing label for this LSP.	**brief detail**
Time left	Number of seconds remaining in the lifetime of the reservation.	**brief detail**
Since	Date and time when the RSVP session was initiated.	**detail**
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	**detail**
DiffServ info	Indicates whether the LSP is a multiclass LSP (**multiclass diffServ-TE LSP**) or a Differentiated-Services-aware traffic engineering LSP (**diffServ-TE LSP**).	**detail**
bandwidth	Bandwidth for each class type (**ct0**, **ct1**, **ct2**, or **ct3**).	**detail**
Port number	Protocol ID and sender/receiver port used in this RSVP session.	**detail**
Attrib flags	**Non-PHP** indicates that ultimate hop popping has been requested by the LSP using this RSVP session	extensive

Table 57: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
FastReroute desired	Fast reroute has been requested by the ingress routing device.	**detail**
Soft preemption desired	Soft preemption has been requested by the ingress routing device.	**detail**
FastReroute desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	**detail extensive**
Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	**detail extensive**
Node/Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	**detail extensive**
Type	LSP type: • **Link protected LSP**—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (**extensive**). • **Node/Link protected LSP**—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (**extensive**). • **Protection down**—LSP is not currently protected. • **Bypass LSP**—LSP that is used to protected one or more user LSPs in case of link failure. • **Backup LSP at Point-of-Local-Repair (PLR)**—LSP that has been temporarily established to protected a user LSP at the ingress of a failed link. • **Backup LSP at Merge Point (MP)**—LSP that has been temporarily established to protected a user LSP at the egress of a failed link.	**detail extensive**
New bypass	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	**extensive**
Link protection up, using *bypass-name*	Link protection (the bypass name is also displayed) has been activated for the LSP.	**extensive**
Creating backup LSP, link down	A **link down** event occurred, and traffic is being switched over to the bypass LSP.	**extensive**
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	**extensive**
Path mtu	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the **allow-fragmentation** statement configured at the **[edit protocols mpls path-mtu]** hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	**detail**

Table 57: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	**detail**
Adspec	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	**detail**
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	**detail**
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	**detail**
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	**detail**

Sample Output

show rsvp session

```
user@host> show rsvp session
Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
10.255.245.214  10.255.245.212  AdminDn  0  1 FF       -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

show rsvp session statistics

```
user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To              From            State   Packets      Bytes    LSPname
10.255.245.24   10.255.245.22     Up          0          0     pro3-bd
10.255.245.24   10.255.245.22     Up      44868    2333136   pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To              From            State   Packets      Bytes    LSPname
10.255.245.22   10.255.245.24     Up          0          0     pro3-db
10.255.245.22   10.255.245.24     Up          0          0   pro3-db-2
Total 2 displayed, Up 2, Down 0
```

```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session detail

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session detail (Path MTU Output Field)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding
```

show rsvp session detail (GMPLS)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1–to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: –
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: –
  Resv style: 0 –, Label in: -, Label out: –
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
```



```
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
```

show rsvp session extensive

```
user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299808
  Resv style: 1 FF, Label in: -, Label out: 299808
  Time left:    -, Since: Thu Sep 20 15:54:20 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 61576 protocol 0
  Attrib flags: Non-PHP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
  RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
  Explct route: 10.0.0.18 10.0.0.22
  Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  140, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session p2mp (Ingress Router)

```
user@host> show rsvp session p2mp
```

```
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout LSPname
10.255.10.95    10.255.10.2     Up       0  1 SE    -           3 to-pe1
P2MP name: test2, P2MP branch count: 2
To              From            State   Rt Style Labelin Labelout LSPname
10.255.10.23    10.255.10.2     Up       0  1 SE    -      299776 to-pe3
10.255.10.16    10.255.10.2     Up       0  1 SE    -      299776 to-pe4
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session p2mp (Transit Router)

```
user@host> show rsvp session p2mp
Ingress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin  Labelout LSPname
10.255.10.23    10.255.10.95    Up       0  1 SE    -       299792 to-pe2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout  LSPname
10.255.10.95    10.255.10.2     Up       0  1 SE    3         -     to-pe1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions
P2MP name: test2, P2MP branch count: 2
To              From            State   Rt Style Labelin Labelout  LSPname
10.255.10.23    10.255.10.2     Up       0  1 SE  299776   299808  to-pe3
10.255.10.16    10.255.10.2     Up       0  1 SE  299776   299856  to-pe4
Total 2 displayed, Up 2, Down 0
```


show rsvp statistics

List of Syntax

Syntax on page 515
Syntax (EX Series Switches) on page 515

Syntax

show rsvp statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show rsvp statistics

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Display Resource Reservation Protocol (RSVP) packet and error statistics.

Options

none—Display RSVP packet and error statistics.

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

Required Privilege Level

view

Related Documentation

- **clear rsvp statistics on page 490**

List of Sample Output

show rsvp statistics on page 518

Output Fields

Table 58 on page 515 describes the output fields for the **show rsvp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 58: show rsvp statistics Output Fields

Field Name	Field Description
Packet Type	Statistics about different RSVP messages.
Total Sent	Total number of packets sent since RSVP was enabled.
Total Received	Total number of packets received since RSVP was enabled.
Last 5 seconds Sent	Total number of packets sent in the last 5 seconds.
Last 5 seconds Received	Number of packets received in the last 5 seconds.
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.

Table 58: show rsvp statistics Output Fields *(continued)*

Field Name	Field Description
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, *Resource ReSerVation Protocol (RSVP)*.
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.
No path information	A reservation was received, but no sender is active.

Table 58: show rsvp statistics Output Fields *(continued)*

Field Name	Field Description
Resv style conflict	The same session contains inconsistent reservation styles.
Port conflict	There were inconsistent port numbers for the same session.
Resv no interface	An interface for the receive reservation packets cannot be located.
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
Recv nack	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to recv Hop	Counter of packets discarded because a TE link was not found.
Rcv pkt disabled interface	Number of RSVP packets received on an interface that is not enabled for RSVP.
Transmit buffer full	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
Transmit failure	Number of times the RSVP task failed to send out a packet.
Receive failure	Number of times the RSVP task failed to read an incoming packet.
P2MP RESV discarded by appl	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
Rate limit	Number of RSVP packets dropped due to rate limiting.
Err msg loop detected	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

Sample Output

show rsvp statistics

```
user@host> show rsvp statistics
  PacketType             Total                  Last 5 seconds
                 Sent       Received         Sent       Received
  Path            355            408            0              0
  PathErr           2             13            0              0
  PathTear        101            139            0              0
  Resv FF           0              0            0              0
  Resv WF           0              0            0              0
  Resv SE         419            225            0              0
  ResvErr           0              0            0              0
  ResvTear          0             13            0              0
  ResvConf          0              0            0              0
  Ack             682           1414            0              0
  SRefresh     395198         236030            5              2
  Hello        578809         578221            4              4
  EndtoEnd RSVP     0              0            0              0

  Errors                          Total             Last 5 seconds
  Rcv pkt bad length                  0                          0
  Rcv pkt unknown type                0                          0
  Rcv pkt bad version                 0                          0
  Rcv pkt auth fail                   0                          0
  Rcv pkt bad checksum                0                          0
  Rcv pkt bad format                  0                          0
  Memory allocation fail              0                          0
  No path information                10                          0
  Resv style conflict                 0                          0
  Port conflict                       0                          0
  Resv no interface                   0                          0
  PathErr to client                  38                          0
  ResvErr to client                   0                          0
  Path timeout                        8                          0
  Resv timeout                       57                          0
  Message out-of-order                0                          0
  Unknown ack msg                  2978                          0
  Recv nack                          86                          0
  Recv duplicated msg-id              5                          0
  No TE-link to recv Hop              0                          0
  Rcv pkt disabled interface          0                          0
  Transmit buffer full                0                          0
  Transmit failure                    0                          0
  Receive failure                     0                          0
  P2MP RESV discarded by appl         0                          0
  Rate limit                        306                          0
  Err msg loop detected               0                          0
```

show rsvp version

List of Syntax

Syntax on page 519
Syntax (EX Series Switches) on page 519

Syntax

show rsvp version
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show rsvp version

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced for QFX switches in Junos OS Release 13.2X51-D15
Statement introduced for QFX Virtual Chassis and Virtual Chassis Fabric in Junos OS Release 14.1X53-D30

Description

Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.

Options

none—Display RSVP protocol settings.

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

Required Privilege Level

view

List of Sample Output

show rsvp version on page 521

Output Fields

Table 59 on page 519 describes the output fields for the **show rsvp version** command. Output fields are listed in the approximate order in which they appear.

Table 59: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: **Enabled** or **Disabled**.
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: **Aggressive**, **Disabled**, or **Normal**. The default is **Normal**.
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.

Table 59: show rsvp version Output Fields *(continued)*

Field Name	Field Description
Graceful deleting timeout	Currently configured value for the **graceful-deletion-timeout** statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.
NSR Mode	Status of the nonstop active routing feature for RSVP on the restarting device: **Disabled**, **Enabled/Master**, or **Enabled/Standby**.
NSR State	State of the nonstop active routing feature for RSVP on the restarting device. Possible values are: • **Idle** • **TE-link sync complete** • **Neighbor sync complete** • **Path state sync complete** • **Resv state sync complete** • **Bypass sync complete** • **Init sync complete**
Setup protection	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: **Enabled** or **Disabled**
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: **Enabled** or **Disabled**.
Restart helper mode	Status of the helper mode feature: **Enabled** or **Disabled**. When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
P2p transit LSP nexthop mode	Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are **Chained** or **Unchained**
P2mp transit LSP nexthop mode	Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are **Chained** or **Unchained**

Sample Output

show rsvp version

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
   RSVP protocol:                Enabled
   R(refresh timer):             30 seconds
   K(keep multiplier):           3
   Preemption:                   Normal
   Soft-preemption cleanup:      30 seconds
   Graceful deletion timeout:    30 seconds
   NSR mode:                     Enabled/Master
   NSR state:                    Init sync complete
   Setup protection:             Disabled
   Graceful restart:             Disabled
   Restart helper mode:          Enabled
   Maximum helper restart time:  20000 msec
   Maximum helper recovery time: 180000 msec
   Restart time:                 0 msec
   P2p transit LSP nexthop mode: Unchained
   P2mp transit LSP nexthop mode: Unchained
```

traceroute mpls rsvp

Syntax

traceroute mpls <rsvp> *lsp-name*
<detail>
<egress>
<exp>
<logical-system>
<multipoint>
<no-resolve>
<retries>
<source *source-address*>
<ttl>

Release Information

Command introduced in Junos OS Release 9.2.
egress, **multipoint**, and **ttl** options added in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Trace route to a remote host for an MPLS LSP signaled by RSVP. Use **traceroute mpls rsvp** as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

lsp-name—Specify the name of the LSP to be traced.

detail—(Optional) Display detailed output.

egress—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.

exp—(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.

logical-system—(Optional) Specify the name of the logical system for the traceroute attempt.

multipoint—(Optional) Perform a trace route on a point-to-multipoint LSP.

no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.

retries—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.

source *source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

ttl—(Optional) Specify the number of hops to follow before forcing the trace route to quit.

Required Privilege Level

network

List of Sample Output

traceroute mpls rsvp on page 524
traceroute mpls rsvp detail on page 524
traceroute mpls rsvp multipoint (branch node for sub-LSPs) on page 525
traceroute mpls rsvp multipoint (single-hop sub-LSPs) on page 525

Output Fields

Table 60 on page 523 describes the output fields for the **traceroute mpls rsvp *lsp-name*** and **traceroute mpls rsvp *lsp-name* detail** commands. Output fields are listed in the approximate order in which they appear.

Table 60: traceroute mpls rsvp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the **traceroute mpls rsvp *lsp-name*** command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays **Success** if the trace to a hop is successful or **Egress** if the trace has reached the last router on the path.	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	**detail**
Parent	Address of the previous hop. Parent value for the first hop is null.	**detail**
Return Code	Return code for reporting the result of processing the echo request by the receiver.	**detail**
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	**detail**
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	**detail**
Response time	Time for the echo request to reach the receiver.	**detail**

Table 60: traceroute mpls rsvp Output Fields *(continued)*

Field Name	Field Description	Level of Output
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	**detail**
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	**detail**
Label stack	Label stack used to forward the packet.	**detail**
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

Sample Output

traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta

     Probe options: retries 3, exp 7

     ttl    Label  Protocol    Address          Previous Hop     Probe Status
       1   299792  RSVP-TE        192.168.1.2      (null)           Success
       2   299803  RSVP-TE        192.168.2.3      192.168.1.2      Success
       3        3  RSVP-TE        192.168.3.4      192.168.2.3      Egress

     Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
Probe options: retries 3, exp 7

   Hop 192.168.1.2 Depth 1
     Probe status: Success
     Parent: (null)
     Return code: Label-switched at stack-depth 1
     Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
     Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
     Response time: 26.99 msec
     MTU: Unknown
     Multipath type: IP bitmask
       Address Range 1: 127.0.0.64 ~ 127.0.0.127
     Label Stack:
       Label 1 Value 299792 Protocol RSVP-TE

Hop 192.168.2.3 Depth 2
     Probe status: Success
     Parent: 192.168.1.2
     Return code: Upstream interface index unknown label-switched at stack-depth
 1
     Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
     Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
     Response time: 26.55 msec
     MTU: 1518
     Multipath type: IP bitmask
```

```
    Address Range 1: 127.0.0.64 ~ 127.0.0.127
  Label Stack:
    Label 1 Value 299803 Protocol RSVP-TE
```

traceroute mpls rsvp multipoint (branch node for sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
  Probe options: retries 3, exp 7

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    1   300000  RSVP-TE     81.1.2.2         (null)           Success
    2   299968  RSVP-TE     81.2.3.3         81.1.2.2         Success
    3   299952  RSVP-TE     81.3.4.4         81.2.3.3         Success
    4   299920  RSVP-TE     81.4.6.6         81.3.4.4         Egress

  Path 1 via lt-1/2/0.102 destination 127.0.0.64

  ttl    Label  Protocol    Address          Previous Hop     Probe Status
    4   299920  RSVP-TE     81.4.5.5         81.3.4.4         Egress

  Path 2 via lt-1/2/0.102 destination 127.0.0.64
```

traceroute mpls rsvp multipoint (single-hop sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
  Probe options: retries 3, exp 7

   ttl    Label  Protocol    Address          Previous Hop     Probe Status
     1        0  RSVP-TE     81.1.2.2         (null)           Egress

   Path 1 via lt-1/2/0.102 destination 127.0.0.64

   ttl    Label  Protocol    Address          Previous Hop     Probe Status
     1        0  RSVP-TE     81.1.8.8         (null)           Egress

   Path 2 via lt-1/2/0.108 destination 127.0.0.64

   ttl    Label  Protocol    Address          Previous Hop     Probe Status
     1        0  RSVP-TE     81.1.9.9         (null)           Egress

   Path 3 via lt-1/2/0.109 destination 127.0.0.64
```