

Chapter 8

Configure Miscellaneous MPLS Properties

This chapter discusses the following topics:

Configure MPLS to Pop the Label on the Ultimate-Hop Router on page 99

Configure Traffic Engineering for LSPs on page 100

Configure MPLS to Gather Statistics on page 102

Control MPLS System Log Messages and SNMP Traps on page 103

Trace MPLS and LSP Packets and Operations on page 104

Configure an MPLS Firewall Filter on page 104

Configure MPLS Rewrite Rules on page 107

Ping LSPs on page 108

Configure MPLS to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of a label switched path (LSP). 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. By enabling ultimate-hop popping, label 0 (IPv4 Explicit Null Label) is advertised. Ultimate hop-popping ensures that any packets traversing an MPLS network include a label.

To configure MPLS to pop the label on the ultimate hop, include the explicit-null statement at the [edit protocols ldp] hierarchy level:

```
[edit protocols ldp]
explicit-null;
```



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” on page 20 and “Label Allocation” on page 21.

Configure Traffic Engineering for LSPs

When you configure an LSP, a host route (a 32-bit mask) is installed in the ingress router toward the egress router, the address of the host route is the destination address of the LSP. By default, only BGP can use LSPs in its route calculations (traffic-engineering bgp). By configuring the other traffic-engineering statement options, you can alter this behavior in the master instance. This functionality is not available for specific routing instances. Also, you can enable only one of the traffic-engineering statement options (bgp, bgp-igp, bgp-igp-both-ribs, or mpls-forwarding) at a time.



Note

Enabling or disabling any of the traffic-engineering statement options causes all the MPLS routes to be removed and then reinserted into the routing tables.

You can configure traffic engineering for LSPs as follows:

Use RSVP and LDP Routes for Traffic Forwarding on page 100

Use RSVP and LDP Routes for Traffic Forwarding in VPNs on page 101

Use RSVP and LDP Routes for Traffic Forwarding but not Route Selection on page 101

You can also configure OSPF and traffic engineering to advertise the LSP metric in summary LSAs as described in the following section:

Advertise the LSP Metric in Summary LSAs on page 102

Use RSVP and LDP Routes for Traffic Forwarding

Configure the bgp-igp option of the traffic-engineering statement to cause BGP and the IGPs to use LSPs for forwarding traffic destined for egress routers. The bgp-igp option causes all inet.3 routes to be moved to the inet.0 routing table.

On the ingress router, include the bgp-igp option of the traffic-engineering statement at the [edit protocol mpls] hierarchy level:

```
[edit protocol mpls]
traffic-engineering bgp-igp;
```



Note

The bgp-igp option of the traffic-engineering statement cannot be configured for VPNs. VPN routing instances require that routes be in the inet.3 routing table.

Use RSVP and LDP Routes for Traffic Forwarding in VPNs

VPNs rely on the routes in the inet.3 routing table to function properly. For VPNs, configure the `bgp-igp-both-ribs` option of the `traffic-engineering` statement to cause BGP and the IGP to use LSPs for forwarding traffic destined for egress routers. The `bgp-igp-both-ribs` option installs the ingress routes in both the inet.0 routing table (for IPv4 unicast routes) and the inet.3 routing table (for MPLS path information).

On the ingress router, include the `bgp-igp-both-ribs` option of the `traffic-engineering` statement at the `[edit protocol mpls]` hierarchy level:

```
[edit protocol mpls]
traffic-engineering bgp-igp-both-ribs;
```

Use RSVP and LDP Routes for Traffic Forwarding but not Route Selection

If you configure the `traffic-engineering bgp-igp` statement or the `traffic-engineering bgp-igp-both-ribs` statement, high-priority RSVP and LDP routes can supersede IGP routes in the inet.0 routing table. IGP routes might no longer be redistributed since they are no longer the active routes.

When you configure the `mpls-forwarding` option at the `[edit protocols mpls traffic-engineering]` hierarchy level, RSVP and LDP routes are used for forwarding but are excluded from route selection. These routes are added to both the inet.0 and inet.3 routing tables. The RSVP and LDP routes in the inet.0 routing table are given a low preference when the active route is selected. However, the RSVP and LDP routes in the inet.3 routing table are given a normal preference and are therefore used for selecting forwarding next hops.

When you activate the `mpls-forwarding` option, routes whose state is `ForwardingOnly` are preferred for forwarding even if their preference is lower than that of the currently active route. To examine the state of a route, execute a `show route detail` command.

To configure, include the `mpls-forwarding` option of the `traffic-engineering` statement at the `[edit protocols mpls]` hierarchy level:

```
[edit protocols mpls]
traffic-engineering mpls-forwarding;
```

When you configure the `mpls-forwarding` option, IGP shortcut routes are copied to the inet.0 routing table only.

Advertise the LSP Metric in Summary LSAs

You can configure MPLS and OSPF to treat an LSP as a link. This allows other routers in the network to use this LSP. To accomplish this, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary link-state advertisements (LSAs).

For MPLS, include the traffic-engineering bgp-igp statement at the [edit protocols mpls] hierarchy level:

```
[edit protocols]
mpls {
  traffic-engineering bgp-igp;
  label-switched-path label-switched-path-name {
    to address;
  }
}
```

For OSPF, include the lsp-metric-into-summary statement at the [edit protocols ospf traffic-engineering shortcuts] hierarchy level:

```
[edit protocols]
ospf {
  traffic-engineering {
    shortcuts {
      lsp-metric-into-summary;
    }
  }
}
```

For more information on MPLS traffic engineering, see “Configure Traffic Engineering for LSPs” on page 100. For more information on OSPF traffic engineering, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions. To do this, include the statistics statement at the [edit protocol mpls] hierarchy level:

```
[edit protocol mpls]
statistics {
  auto-bandwidth;
  file filename <size size files number>;
  interval seconds;
}
```

The default interval is 300 seconds.

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.

The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. Note that 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. The following is a sample of the information included in the output file:

```

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

```

Control MPLS System Log Messages and SNMP Traps

Whenever an LSP makes a transition from up to down, or down to up, and whenever an LSP switches from one active path to another, the ingress router generates a system log message and sends an SNMP trap. The following shows a sample system log message:

```

MPLS lsp sheep1 up on primary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on primary(any) Route 192.168.1.1 192.168.1.2
192.168.1.3
MPLS lsp sheep1 down on primary(any)
MPLS lsp sheep1 up on secondary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on secondary(any) to primary(any), Route 192.168.1.1
192.168.1.2 192.168.1.3

```

For information about the MPLS SNMP traps and the proprietary MPLS MIB, see the *JUNOS Internet Software Configuration Guide: Network Management*.

To disable both the generation of system log messages and SNMP traps, include the following log-updown statement at the [edit protocols mpls] hierarchy level:

```

[edit protocols mpls]
log-updown {
  no-syslog;
  no-trap;
}

```

To disable only the generation of system log messages, configure the following:

```

[edit]
user@host# set protocols mpls log-updown no-syslog

```

For scalability reasons, only the ingress router generates SNMP traps. By default, MPLS issues traps for all configured LSPs. If you have many LSPs, the number of traps can become quite large. To disable the generation of SNMP traps, configure the following:

```

[edit]
user@host# set protocols mpls log-updown no-trap

```

Trace MPLS and LSP Packets and Operations

To trace MPLS and LSP packets and operations, include the traceoptions statement at the [edit routing-options], [edit protocol mpls], or [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy levels:

```
[edit]
traceoptions {
  file filename <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

You can specify the following MPLS-specific flags in the MPLS traceoptions statement:

connection—Trace all circuit cross-connect (CCC) activity.

connection-detail—Trace detailed CCC activity.

cspf—Trace CSPF computations.

cspf-link—Trace links visited during CSPF computations.

cspf-node—Trace nodes visited during CSPF computations.

error—Trace MPLS error conditions.

lsping—Trace lsping packets and return codes.

state—Trace all LSP state transitions.

For general information about tracing and global tracing options, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure an MPLS Firewall Filter

You can configure an MPLS firewall filter to count packets based on the experimental (EXP) bits for the top-level MPLS label in a packet.

The following sections provide an overview of MPLS firewall filters and examples of how to configure MPLS firewall filters:

MPLS Firewall Filter Overview on page 105

MPLS Firewall Filter Examples on page 106

MPLS Firewall Filter Overview

You can configure an MPLS firewall filter to count packets based on the experimental (EXP) bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface. 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 that you cannot apply MPLS firewall filters to Ethernet (fxp0) or loopback (lo0) interfaces.

You can configure an MPLS firewall filter on the M-series and the T-series platforms.

You can configure the following match criteria attributes for MPLS filters at the [edit firewall family mpls filter *filter-name* term *term-name* from] hierarchy level:

exp

exp-except

These attributes can accept EXP bits in the range 0 through 7. You can configure:

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];

If you do not specify a match criteria (that is, you do not configure the from statement and use only the then statement with the count action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the [edit firewall family mpls filter *filter-name* term *term-name* then] hierarchy level:

count

accept

discard

next

policer

For more information about how to configure firewall filters, see the *JUNOS Internet Software Configuration Guide: Policy Framework*. For more information about how to configure interfaces, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* and the *JUNOS Internet Software Configuration Guide: Services Interfaces*.

MPLS Firewall Filter Examples

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

This shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

This shows how to apply the MPLS firewall filter to an interface:

```
[edit interfaces]
so-0/0/0 {
  mtu 4474;
  encapsulation ppp;
  sonet-options {
    fcs 32;
  }
  unit 0 {
    point-to-point;
    family mpls {
      filter {
        input expf;
        output expf;
      }
    }
  }
}
```

The MPLS firewall filter is applied to the input and output of an interface (see the input and output statements in the preceding example).

Configure MPLS Rewrite Rules

You can apply a number of different rewrite rules to MPLS packets.

The following sections describe how you can apply rewrite rules to MPLS packets:

Rewrite the EXP Bits of All Three Labels of an Outgoing Packet on page 107

Rewrite MPLS and IPv4 Packet Headers on page 108

For more information on how to configure statements at the [edit class-of-service] hierarchy level, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* manual.

Rewrite the EXP Bits of All Three Labels of an Outgoing Packet

In interprovider, carrier-of-carrier, and complex traffic engineering scenarios, it is sometimes necessary to push three labels on the next hop, using a swap-push-push or triple-push operation.

By default, on M-series routers, the top MPLS EXP label of an outgoing packet is not rewritten when you configure swap-push-push and triple-push operations. On M-series routers, you can rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining the CoS of an incoming MPLS or non-MPLS packet.

To do this on incoming MPLS packets, include the swap-push-push default statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
  exp-swap-push-push default;
```

To do this on incoming non-MPLS packets, include the push-push-push default statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
  exp-push-push-push default;
```

For more information on how to configure statements at the [edit class-of-service] hierarchy level, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* manual.

Rewrite MPLS and IPv4 Packet Headers

You can apply a rewrite rule to MPLS and IPv4 packet headers simultaneously. This allows you to initialize MPLS EXP and IP precedence bits at LSP ingress. You can configure different rewrite rules depending on whether the traffic is VPN or non-VPN.

To rewrite MPLS and IPv4 packet headers, include the protocol statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules exp *rewrite-rule-name*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp
rewrite-rule-name]
protocol types;
```

Use the protocol statement to specify the types of MPLS packets and packet headers to which to apply the rewrite rule. The MPLS packet can be a standard MPLS packet or an MPLS packet with an IPv4 payload. Specify the type of MPLS packet using the following options:

mpls-any—Applies the rewrite rule to MPLS packets and writes the code point value to MPLS headers.

mpls-inet-both—Applies the rewrite rule to VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T-series routers. On M-series routers, the **mpls-inet-both** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.

mpls-inet-both-non-vpn—Applies the rewrite rule to any non-VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T-series routers. On M-series routers, the **mpls-inet-both-non-vpn** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.

For a detailed example on how to configure rewrite rules for MPLS and IPv4 packets and for more information on how to configure statements at the [edit class-of-service] hierarchy level, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

Ping LSPs

You can ping a specific LSP. Echo requests are sent over the LSP as MPLS packets. The payload is a User Datagram Protocol (UDP) packet forwarded to the address 127.0.0.1 and port 3503. The label and interface information for building and sending this information as an MPLS packet is the same as for standard LSP traffic.

When the echo request arrives at the egress node, the receiver checks the contents of the packet and sends a reply containing the correct return value, using UDP. The router sending the echo request waits to receive an echo reply after a timeout of 2 seconds (you cannot configure this value).

You must configure MPLS at the [edit protocols mpls] hierarchy level on the remote router to be able to ping an LSP terminating there. You must configure MPLS even if you only intend to ping LDP forwarding equivalence classes (FECs).

On the egress router (the router receiving the MPLS echo packets), you must configure the address 127.0.0.1/32 on its lo0 interface. If this is not configured, the egress router does not have this forwarding entry and therefore simply drops the incoming MPLS pings and replies with “ICMP host unreachable” messages.

**Note**

You can only ping LDP FECs and RSVP LSPs. You cannot ping VPN LSPs.

The command you use to ping an MPLS LSP is `ping mpls <count count> <ldp <fec>> <rsvp <lsp-name>>`. For a detailed description of this command, see the *JUNOS Internet Software Operational Mode Command Reference: Protocols, Class of Service, Chassis, and Management*.

