## JUNOS MPLS for EX-series Switches Overview

JUNOS MPLS for EX-series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

JUNOS MPLS for EX-series switches supports RSVP-based label switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs).

---

**NOTE:** MPLS configurations on EX-series switches are compatible with configurations on other Juniper Networks devices that support MPLS and CCC.

---



### *Benefits of MPLS*

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports may be assigned different labels.
- A packet arriving at a particular provider edge switch may be assigned a different label than the same packet entering the network at a different provider edge switch. As a result, forwarding decisions that depend on the ingress provider edge 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 being 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 identity of the explicit route need not be carried with the packet.

---

**NOTE:** MPLS configurations on EX-series switches do not support:

- LDP-based MPLS
- Routed VLAN interfaces (RVIs)

---

### *Additional Benefits of MPLS and Traffic Engineering*

MPLS is the packet-forwarding component of the JUNOS traffic engineering architecture. Traffic engineering provides the capabilities to do the following:

- Route primary paths around known bottlenecks or points of congestion in the network.
- Provide precise control over how traffic is rerouted when the primary path is faced with single or multiple failures.
- Provide more efficient use of available aggregate bandwidth and long-haul fiber by ensuring that subsets of the network do not become overutilized while other subsets of the network along potential alternate paths are underutilized.
- Maximize operational efficiency.
- Enhance the traffic-oriented performance characteristics of the network by minimizing packet loss, minimizing prolonged periods of congestion, and maximizing throughput.
- Enhance statistically bound performance characteristics of the network (such as loss ratio, delay variation, and transfer delay) required to support a multiservice Internet.

For additional information on MPLS traffic protection on EX-series switches, see Understanding MPLS and Path Protection on EX-series Switches.

### *MPLS Label Switched Paths and MPLS Labels on EX-series Switches*

When a packet enters the MPLS network, it is assigned to a label switched path (LSP). Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the packet. 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.

---

**NOTE:** MPLS for EX-series switches supports only single-label packets.

---

### *MPLS Label Operations on EX-series Switches*

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 an LSP, the ingress provider edge 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

provider edge switch, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

EX-series switches support the following label operations:

- Push
- Pop
- Swap

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 has an existing 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 1 shows an IP packet without a label arriving on the customer-edge interface (ge-0/0/1) of the provider edge ingress switch. The provider edge ingress switch examines the packet and identifies that packet's destination is the provider edge egress switch. The provider edge ingress 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 to 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 provider edge egress switch. The provider edge egress 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 scope of the tunnel.

**Figure 1: MPLS Label Swapping**

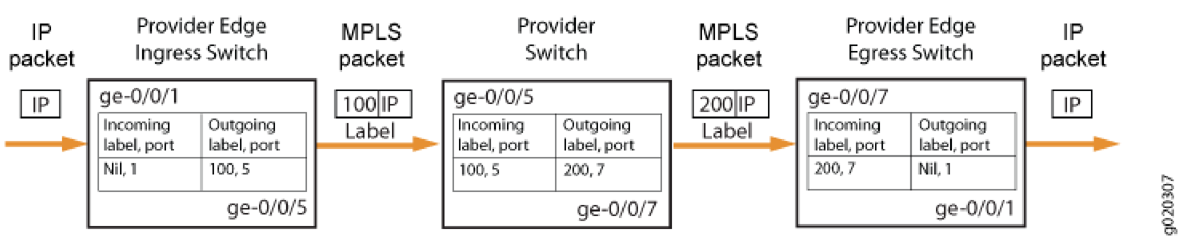


Figure 1 shows the path of a packet as it passes in one direction from the provider edge ingress switch to the provider edge egress switch. However, the MPLS

configuration also allows traffic to travel in the reverse direction. Thus, each provider edge switch operates as both an ingress switch and an egress switch.

**Related Topics**

- Understanding JUNOS MPLS Components for EX-series Switches
- Example: Configuring MPLS on EX-series Switches
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure)
- Configuring Path Protection in an MPLS Network (CLI Procedure)
- *JUNOS Software MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/junos95/index.html
- *JUNOS Software VPNs Configuration Guide* at http://www.juniper.net/techpubs/software/junos/junos95/index.html