JUNIPER NETWORKS

Technology Overview

Frequently Asked Questions: MPLS in Juniper Networks Switches

Published: 2014-03-10

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Technology Overview Frequently Asked Questions: MPLS in Juniper Networks Switches
NCE0115

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

Introduction 1
MPLS in Juniper Networks Switches FAQ Overview 1
MPLS in Juniper Networks Switches FAQ 1

Introduction

This document answers the most frequently asked questions about MPLS support on Juniper Networks® EX Series Ethernet Switches using the Junos® operating system (Junos OS).

MPLS in Juniper Networks Switches FAQ Overview

MPLS provides a mechanism to engineer network traffic patterns that are independent of routing tables. MPLS assigns short labels to network packets that describe how to forward them through the network. MPLS is independent of any routing protocol and can be used for unicast packets.

In the traditional Layer 3 forwarding paradigm, as a packet travels from one router to the next, an independent forwarding decision is made at each hop. The IP network layer 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 performed just once, when a packet enters the MPLS cloud. The packet is then assigned to a stream, 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, the label forwarding table stores forwarding information. You can associate additional information with a label—such as class-of-service (CoS) values that can be used to prioritize packet forwarding.

To use MPLS in your network, you must define an MPLS domain. When an IP packet enters the MPLS domain, the first switch in the domain (known as the ingress provider edge (PE) switch) analyzes the Layer 3 header of the IP packet and inserts an MPLS label in the packet header. This analysis is performed only once, by the ingress PE switch. The label transforms the network layer 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 MPLS packet is then forwarded to the next provider switch in the label-switched path (LSP). This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, the MPLS switches use the label to look up information in their label forwarding table, replacing the old MPLS label stack with a new MPLS label stack before forwarding the labeled packet to the next switch in the path. When the packet reaches the egress PE switch, the MPLS label stack is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

This document presents the most frequently asked questions about MPLS support on EX Series switches using Junos OS.

Related Documentation

- MPLS in Juniper Networks Switches FAQ on page 1

MPLS in Juniper Networks Switches FAQ

This section presents frequently asked questions and answers related to MPLS in Juniper Networks EX Series switches.

What is MPLS?

MPLS provides a mechanism to engineer network traffic patterns that are independent of routing tables. MPLS assigns short labels to network packets that describe how to forward them through the network. MPLS is independent of any routing protocol and can be used for unicast packets.

Why should I use MPLS?

MPLS enables service providers and data centers to utilize far more bandwidth, resulting in a considerable cost savings, by engineering network traffic patterns. MPLS can be used to create state-of-the-art any-to-any networks. With the use of class of service (CoS), MPLS can efficiently accommodate multiple applications such as VoIP, video conferencing, or data traffic over a single link.

What is a label?

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 stack (32 bits). Labels are used as lookup indexes for the label forwarding table. No additional parsing or lookup is performed on the encapsulated packet, enabling MPLS to support the transmission of any other protocols within the packet payload.

Figure 1 on page 2 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 1: Label Encoding

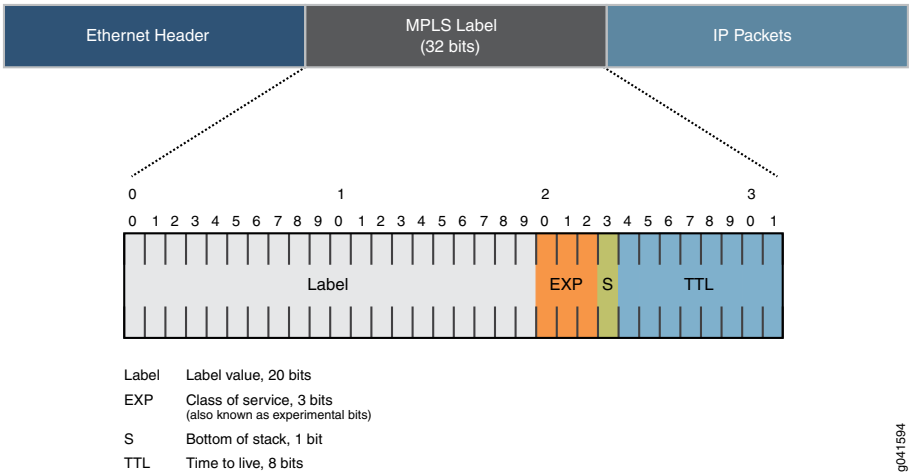

What is a label-switched path?

A label-switched path (LSP) is a unidirectional path through the MPLS network. You can set up an LSP using any of the signaling protocols such as LDP, RSVP, or BGP.

The path begins at an ingress provider edge (PE) switch, which makes a decision on which label to prefix to a packet based on the appropriate forwarding equivalence class (FEC). FECs are a set of packets that have identical characteristics (they use the same

next hop, interface) and are to be forwarded in a similar way. The PE switch then forwards the packet along to the next switch (a provider switch) in the path, which swaps or pops the packet outer label for another label, and forwards it to the next switch. The action—pop or swap—that a provider switch performs is determined by the position of the switch in the LSP. The penultimate provider switch or the last switch (the egress provider edge switch) in the path removes (pops) the label from the packet and forwards the packet based on the header of its next layer (for example IPv4). Due to the forwarding of packets through an LSP being opaque to higher network layers, an LSP is also sometimes referred to as an MPLS tunnel.

What are the possible applications of MPLS with EX Series switches?

Following are the MPLS applications that you can configure using EX Series switches:

- **IP over MPLS**—Uses the traffic engineering capabilities of MPLS to efficiently utilize the existing network (service provider or data center) to manage traffic and to achieve network resiliency.

  In IP over MPLS:

  1. When an IP packet enters the MPLS domain, the first switch in the MPLS domain, also known as the ingress provider edge (PE) switch, analyzes the Layer 3 header of the IP packet and inserts an MPLS label in the packet header. The Layer 3 analysis is done only once, by the ingress PE switch. 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.
  2. The packet is 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.
  3. The provider switches replace the old label with a new label and forward the packet to the next switch in the path.
  4. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

- **Layer 2 over MPLS (CCC)**—Enables you to create an MPLS circuit cross-connect (CCC) between interfaces, connecting two distant interface circuits of the same type by creating MPLS tunnels. MPLS CCC uses the LSPs as the conduit for MPLS packets. Once you have set up a CCC, you must specify an encapsulation for the circuit. The type of CCC-encapsulations that are supported on EX Series switches are Ethernet and VLAN.

- **Layer 2 VPN**—With MPLS Layer 2 VPNs, routing occurs on customer switches, typically on the customer edge (CE) switch.

  In MPLS Layer 2 VPN:

  1. The CE switch connected to a service provider on a Layer 2 VPN selects the appropriate circuit on which to send traffic. The PE switches do not store or process the customer routes; the CE switches must be configured to send Layer 2 data, such

as Ethernet, Frame Relay, asynchronous transfer mode (ATM), or Point-to-Point Protocol (PPP), to the appropriate IP/MPLS tunnel. This gives customers complete control over their own routing.

2. The PE switch receiving the traffic sends it across the service provider's network to the PE switch connected to the receiving site. The service provider must only detect how much traffic the Layer 2 VPN will need to carry. The service provider's switches carry traffic between the customer's sites using Layer 2 VPN interfaces. The VPN topology is determined by policies configured on the PE switches.

- **Layer 3 VPN**—MPLS Layer 3 VPNs enable service providers to 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. Layer 3 VPNs enable customers to leverage the service provider's technical expertise to ensure efficient site-to-site routing. The customer CE switch uses a routing protocol such as BGP or OSPF to communicate with the provider PE switch and to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use IP over MPLS. Other protocol packets are not supported.
- **VPLS**—Virtual private LAN service (VPLS) is a virtual private network (VPN) technology that provides Ethernet-based multipoint to multipoint communication over IP or MPLS networks. VPLS enables geographically dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires. VPLS is protocol independent, and supports IP, IPX, and other legacy protocols. VPLS also offers additional security benefit to sensitive businesses by enabling enterprises to maintain control over their routing tables and eliminating the need to share the routing tables with the service provider. Because of the limitations in the number of sites supported by VPLS and the availability of Ethernet coverage, many enterprises deploy a combination of MPLS and VPLS services, instead of a pure VPLS network. For instance, VPLS can be used for connecting data centers, whereas MPLS can be used for branches.

What MPLS features do the EX Series switches support?

Table 1 on page 5 lists the MPLS software features, the Junos OS release in which they were introduced, and the first Junos OS release for each switch.

NOTE: A separate software license is required for MPLS. See Understanding Software Licenses for EX Series Switches.

Table 1: MPLS Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Aggregated Ethernet interfaces (LAGs) on circuit cross-connects (CCCs)	–	–	–	–	12.2R1	12.2R1	–	11.1R1	See Table 2 on page 7 for a list of EX9200 MPLS features.
BFD for an LDP-based LSP	–	–	–	–	–	–	–	12.2R1	
BFD for an RSVP-based LSP	–	–	–	–	–	–	–	12.2R1	
CCC between 2 interfaces in the same switch	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
Interior gateway protocol (IGP) IS-IS and OSPF shortcuts	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
IP over MPLS	–	10.1R1	–	–	12.2R1 See Note at end of table.	12.2R1 See Note at end of table.	–	11.1R1	
IPv6 over MPLS label-switched paths (LSPs)	–	–	–	–	12.2R1 See Note at end of table.	12.2R1 See Note at end of table.	–	12.1R1	
LDP-based MPLS	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
LDP tunneling (LDP over RSVP)	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS-based circuit cross-connects (CCCs)	–	9.5R1	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS label-switched router (LSR) support	–	–	–	–	12.2R1	12.2R1	–	11.1R1	

Table 1: MPLS Features on Switches by Junos OS Release *(continued)*

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
MPLS Layer 2 CCC on Ethernet-encapsulated interfaces (RFC 6624)	–	9.5R1	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS Layer 2 CCC on VLAN-encapsulated interfaces (RFC 4905)	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS Layer 2 VLAN CCC on Ethernet-encapsulated interfaces (RFC 6624)	–	9.5R1	–	–	12.2R1	12.2R1	–	11.3R1	
MPLS Layer 2 VLAN CCC on VLAN-encapsulated interfaces (RFC 4905)	–	–	–	–	12.2R1	12.2R1	–	11.3R1	
MPLS Layer 2 VPN over CCC	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS Layer 2 VPN over VLAN CCC	–	–	–	–	12.2R1	12.2R1	–	11.3R1	
MPLS OAM-LSP ping	–	–	–	–	–	–	–	11.1R1	
MPLS over untagged Layer 3 interfaces	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS with class of service (CoS)	–	9.5R1	–	–	12.2R5	12.2R5	–	12.1R1	
MPLS Layer 3 VPNs	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
MPLS with RSVP-based label-switched paths (LSPs)	–	9.5R1	–	–	12.2R1	12.2R1	–	11.1R1	

Table 1: MPLS Features on Switches by Junos OS Release *(continued)*

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Layer 3 subinterfaces as MPLS core interfaces	–	–	–	–	12.2R1 See Note at end of table..	12.2R1 See Note at end of table.	–	12.1R1	
Routed VLAN interfaces (RVIs) as MPLS core interfaces	–	–	–	–	–	–	–	12.1R1	
Path maximum transmission unit (MTU) and unicast reverse-path forwarding (RPF) checks for VPNs	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
RSVP-traffic engineering (RSVP-TE)	–	–	–	–	12.2R1	12.2R1	–	11.1R1	
Standby secondary path protection	–	12.1R1	–	–	12.2R1	12.2R1	–	11.1R1	
Static LSPs	–	–	–	–	12.2R1	12.2R1	–	12.1R1	

NOTE: For EX4500 and EX4550 switches to support Layer 3 subinterfaces as MPLS core interfaces, the peer switch that the Layer 3 subinterfaces connect to must be an EX8200 switch.

NOTE: The EX4500 and EX4550 switches do not support IP over MPLS (single MPLS label in the packet) when the switch is positioned as a non-penultimate-hop popping (non-PHP) switch. However, these switches support CCC, Layer 2 VPN, and Layer 3 VPN.

Table 2: MPLS Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Bypass static LSPs	12.3R2
LDP LSP action based on a BFD failure event	12.3R2

Table 2: MPLS Features on EX9200 Switches by Junos OS Release *(continued)*

Feature	Junos OS Release
LDP downstream on demand	12.3R2
LDP, BGP, and VPLS interworking	12.3R2
P2MP LSP traceroute	12.3R2
Static LSP: • Revert timer • Statistics • Traceoptions • At the ingress switch • At the transit switch	12.3R2
Statistics for P2MP LSPs	12.3R2

For a complete list of software features supported on EX Series standalone switches and virtual chassis, see EX Series Switch Software Features Overview or EX Series Virtual Chassis Software Features Overview.

Can I implement class of service with MPLS?

Yes. You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. Juniper Networks EX Series Ethernet Switches support Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

Can I use firewall filters (or ACLs) with MPLS?

Yes. You can configure an MPLS firewall filter to count packets based on the 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: **You cannot apply MPLS firewall filters to Ethernet (fxp0) or loopback (lo0) interfaces.**

What is penultimate-hop popping?

With penultimate-hop popping (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.

Penultimate-hop popping reduces the processing load on the egress PE switch, because it is not responsible for both popping the MPLS label and performing the IP route lookup prior to forwarding the traffic. EX Series switches enable PHP by default with IP over MPLS configurations. You can however change this behavior if required.

Where can I find MPLS configuration examples for EX Series switches?

MPLS configuration examples for EX Series switches can be accessed from the MPLS for EX Series page.

Related Documentation

- MPLS in Juniper Networks Switches FAQ Overview on page 1

