

MPLS FAQ For Beginners

Document ID: 4649

Questions

Introduction

What is Multi-Protocol Label Switching (MPLS)?

What is a label? What is the structure of the label?

Where will the label be imposed in a packet?

What is a Forwarding Equivalence Class (FEC)?

What is an upstream label switch router (LSR)? What is a downstream LSR?

Is R3 a "Downstream" LSR to R4 for 10.1.1.0/24?

What do the terms incoming, outgoing, local, and remote mean when referring to labels?

Can an LSR transmit/receive a native IP packet (non-MPLS) on an MPLS interface?

Can an LSR receive/transmit a labeled packet on a non-MPLS interface?

What platforms and Cisco IOSes support MPLS?

Generic Routing Encapsulation (GRE) tunnel has an overhead of 24 bytes. How much overhead does an MPLS LSP tunnel have?

How does the LSR know which is the top label, bottom label, and a middle label of the label stack?

What is the range of label values? What label values are reserved? What do the reserved values signify?

What protocol and port numbers do LDP and TDP use to distribute labels to LDP/TDP peers?

What restrictions exist for MPLS support on the Catalyst 6500 and 7600 Optical Services Router (OSR)?

Where can I find MPLS configuration samples?

What options are available for load balancing MPLS packets?

Can we configure an 802.1Q trunk between two Cisco Catalyst switches on different sites over a MPLS connection?

Does the outgoing MPLS EXP value inherit the DSCP value in incoming IP packets by default or is the incoming DSCP trusted without any additional configuration on an MPLS enabled interface?

Does DHCP relay function work in the MPLS VPN network?

Related Information

Introduction

This document answers the most frequently asked questions related to Multiprotocol Label Switching (MPLS) from a beginner level.

Q. What is Multi-Protocol Label Switching (MPLS)?

A. MPLS is a packet-forwarding technology which uses labels to make data forwarding decisions. With MPLS, the Layer 3 header analysis is done just once (when the packet enters the MPLS domain). Label inspection drives subsequent packet forwarding. MPLS provides these beneficial applications:


- ◆ Virtual Private Networking (VPN)
- ◆ Traffic Engineering (TE)
- ◆ Quality of Service (QoS)
- ◆ Any Transport over MPLS (AToM)

Additionally, it decreases the forwarding overhead on the core routers. MPLS technologies

are applicable to any network layer protocol.

Q. What is a label? What is the structure of the label?


A. A label is a short, four-byte, fixed-length, locally-significant identifier which is used to identify a Forwarding Equivalence Class (FEC). The label which is put on a particular packet represents the FEC to which that packet is assigned.



- ◆ **Label** Label Value (Unstructured), 20 bits
- ◆ **Exp** Experimental Use, 3 bits; currently used as a Class of Service (CoS) field.
- ◆ **S** Bottom of Stack, 1 bit
- ◆ **TTL** Time to Live, 8 bits

Q. Where will the label be imposed in a packet?

A. The label is imposed between the data link layer (Layer 2) header and network layer (Layer 3) header. The top of the label stack appears first in the packet, and the bottom appears last. The network layer packet immediately follows the last label in the label stack.




Q. What is a Forwarding Equivalence Class (FEC)?

A. FEC is a group of IP packets which are forwarded in the same manner, over the same path, and with the same forwarding treatment. An FEC might correspond to a destination IP subnet, but it also might correspond to any traffic class that the Edge-LSR considers significant. For example, all traffic with a certain value of IP precedence might constitute a FEC.


Q. What is an upstream label switch router (LSR)? What is a downstream LSR?

A. Upstream and downstream are relative terms in the MPLS world. They always refer to a prefix (more appropriately, an FEC). These examples further explain this.




For FEC 10.1.1.0/24, R1 is the "Downstream" LSR to R2.

For FEC 10.1.1.0/24, R2 is the "Upstream" LSR to R1.




For FEC 10.1.1.0/24, R1 is the "Downstream" LSR to R2. And, R2 is the "Downstream" LSR to R3.



For FEC 10.1.1.0/24, R1 is the "Downstream" LSR to R2. For FEC 10.2.2.0/24, R2 is the "Downstream" LSR to R1.

Data flows from upstream to downstream to reach that network (prefix).




The R4 routing table has R1 and R2 as the "next-hops" to reach 10.1.1.0/24.

Q. Is R3 a "Downstream" LSR to R4 for 10.1.1.0/24?

A. No, data flows from upstream to downstream.

Q. What do the terms incoming, outgoing, local, and remote mean when referring to labels?

A. Consider R2 and R3 in this topology. R2 distributes a label L for FEC F to R3. R3 uses label L when it forwards data to FEC-F (because R2 is his downstream LSR for FEC-F). In this scenario:



- ◆ L is the incoming label for F on R2.
- ◆ L is the outgoing label for FEC-F on R3.
- ◆ L is the local binding for FEC F on R2.
- ◆ L is the remote binding for FEC-F on R3.

Q. Can an LSR transmit/receive a native IP packet (non-MPLS) on an MPLS interface?

A. Yes, if the IP is enabled on the interface. Native packets are received/transmitted as usual. IP is just another protocol. MPLS packets have a different Layer 2 encoding. The receiving LSR is aware of the MPLS packet, based on the Layer 2 encoding.

Q. Can an LSR receive/transmit a labeled packet on a non-MPLS interface?

A. No. Packets are never transmitted on an interface which is not enabled for that protocol. MPLS has a certain Ether type code associated with it (just as IP, IPX, and Appletalk have unique Ether types). When a Cisco router receives a packet with an Ether type which is not enabled on the interface, it drops the packet. For example, if a router receives an Appletalk packet on an interface which does not have Appletalk enabled, it drops the packet. Likewise, if an MPLS packet is received on an interface which does not have MPLS enabled, the packet is dropped.

Q. What platforms and Cisco IOSes support MPLS?

A. The Cisco Series 2691, 3640, 3660, 3725, 3745, 6400-NRP-1, 6400-NRP-2SV, 6400-NSP, Catalyst 5000 with Route Switch Module (RSM), 7200, 7301, 7400, 7500, Catalyst 6500/Cisco 7600 Series with WS-SUP720-3B and WS-SUP720-3BXL, Gigabit Switch Router (GSR), Route Processor Module (RPM), Universal Broadband Router (UBR) 7200, AS5350, and IGX8400-URM all support MPLS.

These platforms support the Cisco Tag Distribution Protocol (TDP) as the label distribution protocol.

Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), and Border Gateway Protocol (BGP) information can be found using the Software Advisor (registered customers only) tool. Software Advisor provides a complete list of feature sets supported in the different Cisco IOS versions and on different platforms.

Q. Generic Routing Encapsulation (GRE) tunnel has an overhead of 24 bytes. How much overhead does an MPLS LSP tunnel have?

A. An MPLS LSP tunnel has one label (four bytes) or two labels (for example, when using Link Protection Fast reroute) of overhead. Unlike GRE tunnel, MPLS does not change the IP header. Instead, the label stack is imposed on to the packet that takes the tunnel path.

Q. How does the LSR know which is the top label, bottom label, and a middle label of the label stack?

A. The label immediately after the Layer 2 header is the top label, and the label with the S bit set to 1 is the bottom label. No application requires LSR to read/identify the middle labels. However, a label will be a middle label if it is not at the top of the stack and the S bit is set to 0.

Q. What is the range of label values? What label values are reserved? What do the reserved values signify?

A. These values can also be found in RFC3032 – MPLS Label Stack Encoding .

Theoretically, the range is 0 through ($2^{20}-1$). Label values 0–15 are reserved, and values 4–15 are reserved for future use. Values 0–3 are defined as:

- ◆ A value of 0 represents the "IPv4 Explicit NULL Label". This label indicates that the label stack must be popped, and the packet forwarding must be based on the IPv4 header. This helps to keep Exp bits safe until the egress router. It is used in MPLS based QoS.
- ◆ A value of 1 represents the "Router Alert Label". When a received packet contains this label value at the top of the label stack, it is delivered to a local software module for processing. The actual packet forwarding is determined by the label beneath it in the stack. However, if the packet is forwarded further, the Router Alert Label should be pushed back onto the label stack before forwarding. The use of this label is analogous to the use of the "Router Alert Option" in IP packets (for example, ping with record route option)
- ◆ A value of 2 represents the "IPv6 Explicit NULL Label". It indicates that the label stack must be popped, and the packet forwarding must be based on the IPv6 header.
- ◆ A value of 3 represents the "Implicit NULL Label". This is a label that an LSR can assign and distribute. However, it never actually appears in the encapsulation. It indicates that the LSR pops the top label from the stack and forwards the rest of the packet (labeled or unlabeled) through the outgoing interface (as per the entry in Lfib). Although this value might never appear in the encapsulation, it needs to be specified in the Label Distribution Protocol, so a value is reserved.

Q. What protocol and port numbers do LDP and TDP use to distribute labels to LDP/TDP peers?

A. LDP uses TCP port 646, and TDP uses TCP port 711. These ports are opened on the router interface only when **mpls ip** is configured on the interface. The use of TCP as a transport protocol results in reliable delivery of LDP/TDP information with robust flow control and congestion handling mechanisms.

Q. What restrictions exist for MPLS support on the Catalyst 6500 and 7600 Optical Services Router (OSR)?

A. The interface connected to the MPLS domain must use one of the Optical Services Modules (OSM) (for example, any module that utilizes Parallel Express Forwarding (PXF) complex) or an interface in the FlexWAN module. The same restriction exists for MPLS Layer 3 VPN. That is, the IP frame must enter on a WAN interface which is either an OSM or an interface in a FlexWAN module. These restrictions do not exist on a Supervisor 720.

Q. Where can I find MPLS configuration samples?

A. There are many MPLS configuration documents located at Implementation and Configuration: MPLS.

Q. What options are available for load balancing MPLS packets?

A. MPLS packets can be load balanced with the MPLS label information and/or the source and destination address of the essential IP header.

Q. Can we configure an 802.1Q trunk between two Cisco Catalyst switches on different sites over a MPLS connection?

A. When you connect to a remote site through MPLS, it is a layer 3 connection, and the 802.1Q trunk is a layer 2 protocol, so you cannot have an 802.1Q trunk across a MPLS connection. You need to have a Metro Ethernet connection or 802.1Q tunneling to expand your VLAN, which is provided by the ISP. In the MPLS cloud, the ISP communicates through VRF.

Refer to [Configuring IEEE 802.1Q Tunneling](#) for more information.

Q. Does the outgoing MPLS EXP value inherit the DSCP value in incoming IP packets by default or is the incoming DSCP trusted without any additional configuration on an MPLS enabled interface?

A. Yes, no additional configuration is needed.

Q. Does DHCP relay function work in the MPLS VPN network?

A. Yes, the DHCP request is forwarded within the VRF across the MPLS VPN network and the egress Provider Edge sends it in the same VRF to the DHCP server.

Related Information

- [MPLS Support Page](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2009 – 2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Sep 18, 2008

Document ID: 4649
