



Implementing MPLS Forwarding

This module describes how to implement MPLS Forwarding on Cisco ASR 9000 Series Aggregation Services Routers.

All Multiprotocol Label Switching (MPLS) features require a core set of MPLS label management and forwarding services; the MPLS Forwarding Infrastructure (MFI) supplies these services.

Feature History for Implementing MPLS-TE

Release	Modification
Release 3.7.2	This feature was introduced.
Release 3.9.0	No modification.
Release 6.0	The Label Security for BGP Inter-AS Option-B feature was modified.

- [Prerequisites for Implementing Cisco MPLS Forwarding, on page 1](#)
- [Restrictions for Implementing Cisco MPLS Forwarding, on page 2](#)
- [Information About Implementing MPLS Forwarding, on page 2](#)
- [How to Implement MPLS Forwarding, on page 4](#)
- [Enable Per-Prefix Statistics Collection, on page 5](#)
- [Additional References, on page 7](#)

Prerequisites for Implementing Cisco MPLS Forwarding

These prerequisites are required to implement MPLS Forwarding:

- You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.
- Router that runs Cisco IOS XR software.
- Installed composite mini-image and the MPLS package, or a full composite image.

Restrictions for Implementing Cisco MPLS Forwarding

- Label switching on a Cisco router requires that Cisco Express Forwarding (CEF) be enabled.
- CEF is mandatory for Cisco IOS XR software and it does not need to be enabled explicitly.

Information About Implementing MPLS Forwarding

To implement MPLS Forwarding, you should understand these concepts:

MPLS Forwarding Overview

MPLS combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. MPLS enables service providers to meet the challenges of growth in network utilization while providing the opportunity to differentiate services without sacrificing the existing network infrastructure. The MPLS architecture is flexible and can be employed in any combination of Layer 2 technologies. MPLS support is offered for all Layer 3 protocols, and scaling is possible well beyond that typically offered in today's networks.

Based on routing information that is stored in the VRF IP routing table and VRF CEF table, packets are forwarded to their destination using MPLS.

A PE router binds a label to each customer prefix learned from a CE router and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet received from a CE router across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet it pops the label and uses it to direct the packet to the correct CE router. Label forwarding across the provider backbone, is based on either dynamic label switching or traffic engineered paths. A customer data packet carries two levels of labels when traversing the backbone:

- Top label directs the packet to the correct PE router
- Second label indicates how that PE router should forward the packet to the CE router

Label Switching Functions

In conventional Layer 3 forwarding mechanisms, as a packet traverses the network, each router extracts all the information relevant to forwarding the packet from the Layer 3 header. This information is then used as an index for a routing table lookup to determine the next hop for the packet.

In the most common case, the only relevant field in the header is the destination address field, but in some cases, other header fields might also be relevant. As a result, the header analysis must be done independently at each router through which the packet passes. In addition, a complicated table lookup must also be done at each router.

In label switching, the analysis of the Layer 3 header is done only once. The Layer 3 header is then mapped into a fixed-length, unstructured value called a *label*.

Many different headers can map to the same label, as long as those headers always result in the same choice of next hop. In effect, a label represents a forwarding equivalence class—that is, a set of packets which, however different they may be, are indistinguishable by the forwarding function.

The initial choice of a label need not be based exclusively on the contents of the Layer 3 packet header; for example, forwarding decisions at subsequent hops can also be based on routing policy.

Once a label is assigned, a short label header is added at the front of the Layer 3 packet. This header is carried across the network as part of the packet. At subsequent hops through each MPLS router in the network, labels are swapped and forwarding decisions are made by means of MPLS forwarding table lookup for the label carried in the packet header. Hence, the packet header does not need to be reevaluated during packet transit through the network. Because the label is of fixed length and unstructured, the MPLS forwarding table lookup process is both straightforward and fast.

Distribution of Label Bindings

Each label switching router (LSR) in the network makes an independent, local decision as to which label value to use to represent a forwarding equivalence class. This association is known as a label binding.



Note The distribution of label bindings cannot be done statically for the Layer 2 VPN pseudowire.

Each LSR informs its neighbors of the label bindings it has made. This awareness of label bindings by neighboring routers is facilitated by these protocols:

Label Distribution Protocol (LDP)

Supports MPLS forwarding along normally routed paths.

Resource Reservation Protocol (RSVP)

Supports MPLS traffic engineering.

Border Gateway Protocol (BGP)

Supports MPLS virtual private networks (VPNs).

When a labeled packet is sent from LSR A to the neighboring LSR B, the label value carried by the IP packet is the label value that LSR B assigned to represent the forwarding equivalence class of the packet. Thus, the label value changes as the IP packet traverses the network.

MFI Control-Plane Services

The MFI control-plane provides services to MPLS applications, such as Label Distribution Protocol (LDP) and Traffic Engineering (TE), that include enabling and disabling MPLS on an interface, local label allocation, MPLS rewrite setup (including backup links), management of MPLS label tables, and the interaction with other forwarding paths (IP Version 4 [IPv4] for example) to set up imposition and disposition.

MFI Data-Plane Services

The MFI data-plane provides a software implementation of MPLS forwarding in all of these forms:

- Imposition

- Disposition
- Label swapping

MPLS Maximum Transmission Unit

MPLS maximum transmission unit (MTU) indicates that the maximum size of the IP packet can still be sent on a data link, without fragmenting the packet. In addition, data links in MPLS networks have a specific MTU, but for labeled packets. All IPv4 packets have one or more labels. This does imply that the labeled packets are slightly bigger than the IP packets, because for every label, four bytes are added to the packet. So, if n is the number of labels, $n * 4$ bytes are added to the size of the packet when the packet is labeled. The MPLS MTU parameter pertains to labeled packets.

Label Security for BGP Inter-AS Option-B

Option-B is a method to exchange VPNv4/VPNv6 routes between Autonomous Systems (AS), as described in RFC-4364. When a router configured with Option-B, peers with a router from another confederation, or an autonomous system, and receives a labeled packet from such an external peer, the router ensures the following:

- the top label is advertised to the source of traffic
- label stack on the packet received from the external peer contains at least one label (explicit null label is not included)

How to Implement MPLS Forwarding

These topics explain how to configure a router for MPLS forwarding.

Configuring MPLS Label Security

Perform this task to configure the MPLS label security on the interface.

SUMMARY STEPS

1. **configure**
2. **interface** *type interface-path-id*
3. **mpls label-security rpf**
4. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	
Step 2	interface <i>type interface-path-id</i> Example:	Enters the interface configuration mode.

	Command or Action	Purpose
	RP/0/RSP0/CPU0:router (config) # interface tunnel-te 1	
Step 3	mpls label-security rpf Example: RP/0/RSP0/CPU0:router (config-if) # mpls label-security rpf	Configures the MPLS label security on the specified interface and checks for RPF label on incoming packets.
Step 4	commit	

Enable Per-Prefix Statistics Collection

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Enable Per-Prefix Statistics Collection	Release 7.10.1	<p>You can now enable the router to collect only the per-prefix statistics, that is, the aggregate statistics for all paths available for a destination prefix. This helps conserve the hardware resources of the router.</p> <p>In earlier releases, there was no option to disable per-path statistics collection. The router collected both per-prefix statistics and per-path statistics for each destination prefix. When there were multiple paths for a destination prefix, the router required much more hardware resources to store the per-path statistics.</p> <p>This feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • hw-module 13 feature per-prefix-egress-stats enable <p>YANG Data Model:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-um-hw-module-13-cfg.yang</code> (see GitHub, YANG Data Models Navigator)

In large networks, each destination prefix has multiple paths. The router requires a lot of hardware resources to collect and store statistics for each of the various paths for every destination prefix. Hence per-path statistics collection is resource-intensive.

Now you can enable the router to collect only the aggregate statistics for destination prefixes, that is, the per-prefix statistics collection. Hence the router doesn't have to allocate memory for statistics collection for each path of a destination prefix. For every destination prefix, the router collects only one set of statistics. This frees up the platform hardware resources and allows you to scale up the number of destination prefixes in a router. You can enable this feature with the **hw-module 13 feature per-prefix-egress-stats enable** command.

Configure Per-Prefix Statistics Collection

```
Router(config)#hw-module 13 feature per-prefix-egress-stats enable
Router(config)#commit
Router(config)#exit
Router#reload location all
```



Note Reload the router after you enable or disable the per-prefix egress statistics collection. By default, per-prefix egress statistics collection is disabled.

Verification of Per-Prefix Statistics Collection

When you enable the `hw-module 13 feature per-prefix-egress-stats enable` command, the following changes apply for the command outputs for `show mpls forwarding labels`:

- The field **Bytes Switched** is 0 for every path of a destination prefix in the output of the command `show mpls forwarding labels label detail`
- The router displays **Total Packets/Bytes Switched** at the bottom of the command output for `show mpls forwarding labels label detail`.
- The command `show mpls forwarding labels label` displays a value of 0 for the **Bytes Switched** field for every path of a destination prefix.

```
Router#show mpls forwarding labels 24001 detail
Local  Outgoing  Prefix          Outgoing      Next Hop      Bytes
Label  Label     or ID           Interface     -----      Switched
-----
24001  Pop       30.30.30.1/32   Hu0/0/0/2    11.1.1.2     0 <-- per-path stats is
0
    Updated: Oct  7 03:29:30.877
    Version: 129, Priority: 3
    Label Stack (Top -> Bottom): { Imp-Null }
    NHID: 0xa, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
    MAC/Encaps: 4/4, MTU: 1500
    Outgoing Interface: HundredGigE0/0/0/2 (ifhandle 0x000000c0)
    Packets Switched: 0

    Pop       30.30.30.1/32   Hu0/0/0/1    12.1.1.2     0 <-- per-path stats
is 0
    Updated: Oct  7 03:29:30.877
    Version: 129, Priority: 3
    Label Stack (Top -> Bottom): { Imp-Null }
    NHID: 0x9, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
    MAC/Encaps: 4/4, MTU: 1500
    Outgoing Interface: HundredGigE0/0/0/1 (ifhandle 0x00000140)
    Packets Switched: 0

    Total Packets/Bytes Switched: 30/3000          <---- Per-Prefix Stats (aggregate
packets/bytes switched from all paths)
```

```
Router#show mpls forwarding labels 24000
Local  Outgoing  Prefix          Outgoing      Next Hop      Bytes
Label  Label     or ID           Interface     -----      Switched
-----
24000  24000    10.10.10.1/32   Hu0/0/0/2    11.1.1.1     0 <-- per-path stats
is 0
```

```

                24000      10.10.10.1/32      Hu0/0/0/1      12.1.1.1      0 <-- per-path stats
is 0

```

When you disable the per-prefix egress statistics collection mode with the **no hw-module l3 feature per-prefix-egress-stats enable** command, the router populates the **Bytes Switched** field for each path of a destination prefix.

```

Router#show mpls forwarding labels 24001 detail
Local  Outgoing  Prefix          Outgoing      Next Hop      Bytes
Label  Label      or ID          Interface     Interface     Switched
-----
24001  Pop        10.3.1.0/24    Te0/0/0/1/0  10.1.1.1      76636        <--per-path
stats is displayed
  Updated: May  5 03:22:33.070
  Version: 17, Priority: 3
  Label Stack (Top -> Bottom): { Imp-Null }
  NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
  MAC/Encaps: 4/4, MTU: 1500
  Outgoing Interface: TenGigE0/0/0/1/0 (ifhandle 0x000004c0)
  Packets Switched: 1666

      Pop        10.3.1.0/24    Te0/0/0/1/1  10.2.1.1      153318        <--per-path
stats is displayed
  Updated: May  5 03:22:33.070
  Version: 17, Priority: 3
  Label Stack (Top -> Bottom): { Imp-Null }
  NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
  MAC/Encaps: 4/4, MTU: 1500
  Outgoing Interface: TenGigE0/0/0/1/1 (ifhandle 0x00000500)
  Packets Switched: 3333

      Pop        10.3.1.0/24    Te0/0/0/1/5  10.4.1.1      230046        <--per-path
stats is displayed
  Updated: May  5 03:22:40.033
  Version: 17, Priority: 3
  Label Stack (Top -> Bottom): { Imp-Null }
  NHID: 0x0, Encap-ID: N/A, Path idx: 2, Backup path idx: 0, Weight: 0
  MAC/Encaps: 4/4, MTU: 1500
  Outgoing Interface: TenGigE0/0/0/1/5 (ifhandle 0x00000600)
  Packets Switched: 5001

Total Packets/Bytes Switched: 10000/460000
<--Per-Prefix Stats (aggregate packets/bytes switched from all paths)

```

```

Router#show mpls forwarding labels 24001
Local  Outgoing  Prefix          Outgoing      Next Hop      Bytes
Label  Label      or ID          Interface     Interface     Switched
-----
24001  Pop        10.3.1.0/24    Te0/0/0/1/0  10.1.1.1      76636        <--per-path
stats is displayed
      Pop        10.3.1.0/24    Te0/0/0/1/1  10.2.1.1      153318        <--per-path
stats is displayed
      Pop        10.3.1.0/24    Te0/0/0/1/5  10.4.1.1      230046        <--per-path
stats is displayed

```

Additional References

For additional information related to implementing MPLS Forwarding, refer to the following references:

Related Documents

Related Topic	Document Title
MPLS Forwarding commands	<i>MPLS Forwarding Commands on Cisco ASR 9000 Series Router</i> module in <i>Cisco ASR 9000 Series Aggregation Services Routers MPLS Command Reference</i>
Getting started material	<i>Cisco ASR 9000 Series Aggregation Services Routers Getting Started Guide</i>

Related Topic	Document Title
MPLS Forwarding commands	<i>MPLS Forwarding Commands on Cisco IOS XR Software</i> module in <i>Cisco IOS XR MPLS Command Reference for the Cisco ASR 9000 Series Router</i>
Getting started material	<i>Cisco IOS XR Getting Started Guide for the Cisco ASR 9000 Series Router</i>
Information about user groups and task IDs	<i>Configuring AAA Services on Cisco IOS XR Software</i> module of <i>Cisco IOS XR System Security Configuration Guide for the Cisco ASR 9000 Series Router</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
RFC 3031	<i>Multiprotocol Label Switching Architecture</i>
RFC 3443	<i>Time to Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks</i>
RFC 4105	<i>Requirements for Inter-Area MPLS Traffic Engineering</i>