



Multiprotocol Label Switching (MPLS) on Cisco Routers

This document describes commands for configuring and monitoring Multiprotocol Label Switching (MPLS) functionality on Cisco routers and switches. This document is a companion to other feature modules describing other MPLS applications.

This document includes the following sections:

- [Feature Overview](#)
- [Supported Standards, MIBs, and RFCs](#)
- [Configuration Tasks](#)
- [Glossary](#)

The following table outlines the release history of this document over time and the changes made thereto in support of new MPLS/LDP functionality.

Release	Modification
11.1CT	The document supporting this release was titled <i>Tag Switching on Cisco 7000 Family</i> .
12.1(3)T	The document supporting this release, titled <i>MPLS Label Switching on Cisco Routers</i> , described a set of generic label switching commands intended for use with Cisco routers and switches. This set of commands marked the beginning of Cisco's transition from tag-switching command syntax and terminology to the new MPLS IETF command syntax and terminology.
12.1(5)T	The document supporting this release described new arguments in the mpls ip propagate-ttl command and added a new mpls ip ttl-expiration pop command.
12.0(14)ST	This document describes a new parameter (vrf vpn-name) for the show mpls forwarding-table command and the show mpls interfaces command that provides support for new MPLS VPN functionality for LDP.



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2007 Cisco Systems, Inc. All rights reserved.

Feature Overview

Multiprotocol label switching (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 explosive 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.

MPLS efficiently enables the delivery of IP services over an ATM switched network. MPLS supports the creation of different routes between a source and a destination on a purely router-based Internet backbone. By incorporating MPLS into their network architecture, service providers can save money, increase revenue and productivity, provide differentiated services, and gain competitive advantages.

Functional Description of MPLS

Label switching is a high-performance packet forwarding technology that integrates the performance and traffic management capabilities of data link layer (Layer 2) switching with the scalability, flexibility, and performance of network layer (Layer 3) routing.

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

- Tag Distribution Protocol (TDP)—Used to support MPLS forwarding along normally routed paths
- Resource Reservation Protocol (RSVP)—Used to support MPLS traffic engineering
- Border Gateway Protocol (BGP)—Used to support MPLS virtual private networks (VPNs)

When a labeled packet is being 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.

Benefits

MPLS provides the following major benefits to service provider networks:

- Scalable support for virtual private networks (VPNs)—MPLS enables VPN services to be supported in service provider networks, thereby greatly accelerating Internet growth.

The use of MPLS for VPNs provides an attractive alternative to the building of VPNs by means of either ATM or Frame Relay permanent virtual circuits (PVCs) or various forms of tunneling to interconnect routers at customer sites.

Unlike the PVC VPN model, the MPLS VPN model is highly scalable and can accommodate increasing numbers of sites and customers. The MPLS VPN model also supports “any-to-any” communication among VPN sites without requiring a full mesh of PVCs or the backhauling (suboptimal routing) of traffic across the service provider network. For each MPLS VPN user, the service provider’s network appears to function as a private IP backbone over which the user can reach other sites within the VPN organization, but not the sites of any other VPN organization.

From a user perspective, the MPLS VPN model enables network routing to be dramatically simplified. For example, rather than having to manage routing over a topologically complex virtual backbone composed of many PVCs, an MPLS VPN user can generally employ the service provider’s backbone as the default route in communicating with all of the other VPN sites.

- Explicit routing capabilities (also called constraint-based routing or traffic engineering)—Explicit routing employs “constraint-based routing,” in which the path for a traffic flow is the shortest path that meets the resource requirements (constraints) of the traffic flow.

In MPLS traffic engineering, factors such as bandwidth requirements, media requirements, and the priority of one traffic flow versus another can be taken into account. These traffic engineering capabilities enable the administrator of a service provider network to

- Control traffic flow in the network
- Reduce congestion in the network
- Make best use of network resources

Thus, the network administrator can specify the amount of traffic expected to flow between various points in the network (thereby establishing a traffic matrix), while relying on the routing system to

- Calculate the best paths for network traffic
- Set up the explicit paths to carry the traffic

- Support for IP routing on ATM switches (also called IP and ATM integration)—MPLS enables an ATM switch to perform virtually all of the functions of an IP router. This capability of an ATM switch stems from the fact that the MPLS forwarding paradigm, namely, label swapping, is exactly the same as the forwarding paradigm provided by ATM switch hardware.

The key difference between a conventional ATM switch and an ATM label switch is the control software used by the latter to establish its virtual channel identifier (VCI) table entries. An ATM label switch uses IP routing protocols and the Tag Distribution Protocol (TDP) to establish VCI table entries.

An ATM label switch can function as a conventional ATM switch. In this dual mode, the ATM switch resources (such as VCI space and bandwidth) are partitioned between the MPLS control plane and the ATM control plane. The MPLS control plane provides IP-based services, while the ATM control plane supports ATM-oriented functions, such as circuit emulation or PVC services.

Restrictions

Label switching on a Cisco router requires that Cisco Express Forwarding (CEF) be enabled on that router (see the “[Configuration Tasks](#)” section below).

Supported Standards, MIBs, and RFCs

The supported standards, MIBs, and RFCs applicable to the MPLS applications appear in the respective feature module for the application.

Configuration Tasks

This section explains how to configure a router for MPLS forwarding by enabling CEF on the router.

Configuration tasks for other MPLS applications are described in the feature module documentation for the application.

Configuring a Router for MPLS Forwarding

MPLS forwarding on Cisco routers requires that CEF be enabled. To enable CEF on a router, issue the following commands:

```
Router# configure terminal
Router(config)# ip cef [distributed]
```



Note

For best performance in network MPLS forwarding activity, use the **distributed** option in the **ip cef** command on routers that support this option.

For more information about the CEF command, refer to the appropriate chapters on CEF in the following documents:

- Cisco IOS Switching Services Command Reference
- Cisco IOS Command Reference

Verifying Configuration of MPLS Forwarding

To verify that CEF has been configured properly on a Cisco router, issue the **show ip cef summary** command, which generates output similar to that shown below:

```
Router# sho ip cef summary
IP CEF with switching (Table Version 49), flags=0x0
  43 routes, 0 resolve, 0 unresolved (0 old, 0 new)
  43 leaves, 49 nodes, 56756 bytes, 45 inserts, 2 invalidations
  2 load sharing elements, 672 bytes, 2 references
  1 CEF resets, 4 revisions of existing leaves
  4 in-place modifications
  refcounts: 7241 leaf, 7218 node

Adjacency Table has 18 adjacencies
Router#
```

Glossary

ATM edge LSR—A router that is connected to the ATM-LSR cloud through LC-ATM interfaces. The ATM edge LSR adds labels to unlabeled packets and strips labels from labeled packets.

ATM-LSR—A label switch router with a number of LC-ATM interfaces. The router forwards the cells among these interfaces using labels carried in the VPI/VCI field of the ATM cell header.

CoS—Class of service. A feature that provides scalable, differentiated types of service across an MPLS network.

IP precedence—A 3-bit value in a ToS byte used for assigning precedence to IP packets.

label—A short fixed-length label that tells switching nodes how to forward data (packets or cells).

label-controlled ATM interface (LC-ATM interface)—An interface on a router or switch that uses label distribution procedures to negotiate label VCs.

label edge router (LER)—A router that performs label imposition.

label imposition—The action of putting the first label on a packet.

label switch—A node that forwards units of data (packets or cells) on the basis of labels.

label-switched path (LSP)—A sequence of hops (Router 0...Router n) in which a packet travels from R0 to Rn by means of label switching mechanisms. A label-switched path can be chosen dynamically, based on normal routing mechanisms, or it can be configured manually.

label-switched path (LSP) tunnel—A configured connection between two routers, in which label switching techniques are used for packet forwarding.

label switching router (LSR)—A Layer 3 router that forwards a packet based on the value of a label encapsulated in the packet.

label VC (LVC)—An ATM virtual circuit that is set up through ATM LSR label distribution procedures.

LFIB—Label Forwarding Information Base. The data structure used by switching functions to switch labeled packets.

LIB—Label information base. A database used by an LSR to store labels learned from other LSRs, as well as labels assigned by the local LSR.

MPLS—Multiprotocol label switching. An emerging industry standard that defines support for MPLS forwarding of packets along normally routed paths (sometimes called MPLS hop-by-hop forwarding).

QoS—Quality of service. A measure of performance for a transmission system that reflects its transmission quality and service availability.

tailend—The downstream, received end of a tunnel.

TDP—Tag Distribution Protocol. The protocol used to distribute label bindings to LSRs.

traffic engineering—The techniques and processes used to cause routed traffic to travel through the network on a path other than the one that could have been chosen if standard routing methods had been applied.

traffic engineering tunnel—A label-switched tunnel that is used for traffic engineering. Such a tunnel is set up through means other than normal Layer 3 routing; it is used to direct traffic over a path different from the one that Layer 3 routing could cause the tunnel to take.

VPN—Virtual private network. Enables IP traffic to use tunneling to travel securely over a public TCP/IP network.



Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2007 Cisco Systems, Inc. All rights reserved.

