MPLS

Implementing Cisco MPLS

Volume 1

Version 2.2

Student Guide

Text Part Number: 97-2389-02

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Students, this letter describes important course evaluation access information!

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco Systems is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. We would appreciate a few minutes of your time to complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web.

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,

Cisco Systems Learning

Table of Contents

Volume 1

Course Introduction

Overview

Service providers (and enterprises acting as service providers) are faced with many challenges in terms of customer demand, including an ongoing need for value-added services. Conventional IP packet forwarding has several limitations, and more and more service providers realize that something else is needed. Not only must service providers be concerned with protecting their existing infrastructure, but they must also find ways to generate new services that are not currently supportable using existing technologies.

Multiprotocol Label Switching (MPLS) is a high-performance method for forwarding packets through a network. MPLS enables routers at the edge of a network to apply simple labels to packets. This practice allows the edge devices—ATM switches or existing routers in the center of the service provider core—to switch packets according to labels, with minimal lookup overhead. MPLS integrates the performance and traffic-management capabilities of data link Layer 2 with the scalability and flexibility of network Layer 3 routing. When used in conjunction with other standard technologies, MPLS allows service providers the ability to support value-added features that are critical for their networks.

Implementing Cisco MPLS (MPLS) v2.2 is recommended training for individuals seeking certification as a Cisco CCIP®. The focus of this course is on MPLS technology issues as those issues apply to service providers and on how to configure new features and functions in an existing routed environment.

Learner Skills and Knowledge

This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should complete to benefit fully from this course.

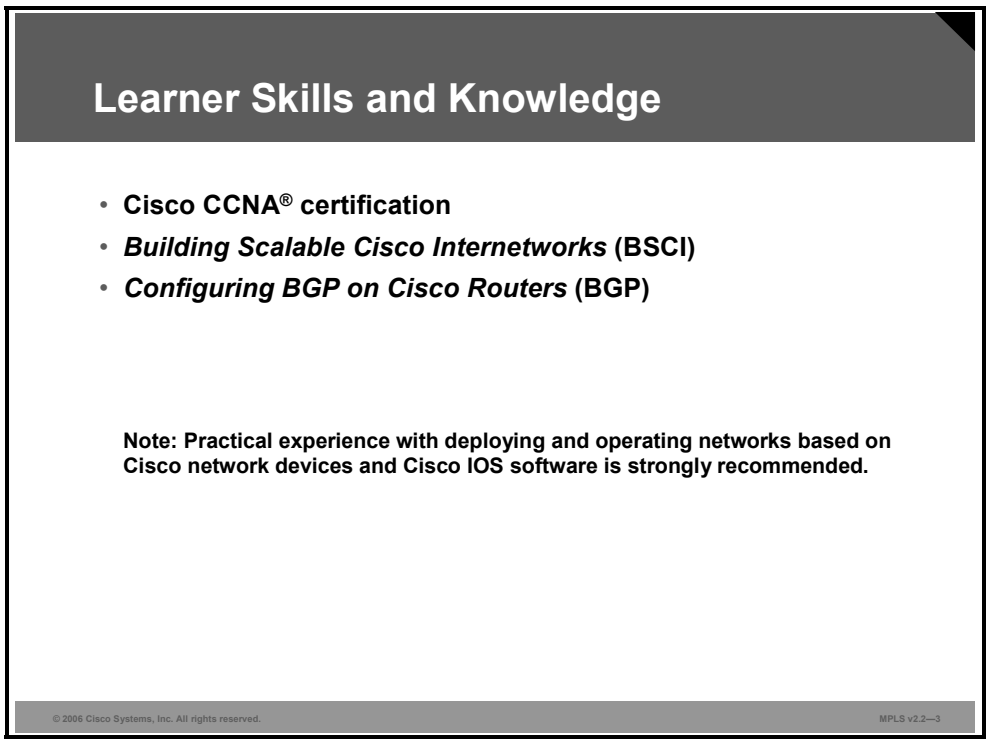

Course Goal and Objectives

This topic describes the course goal and objectives.

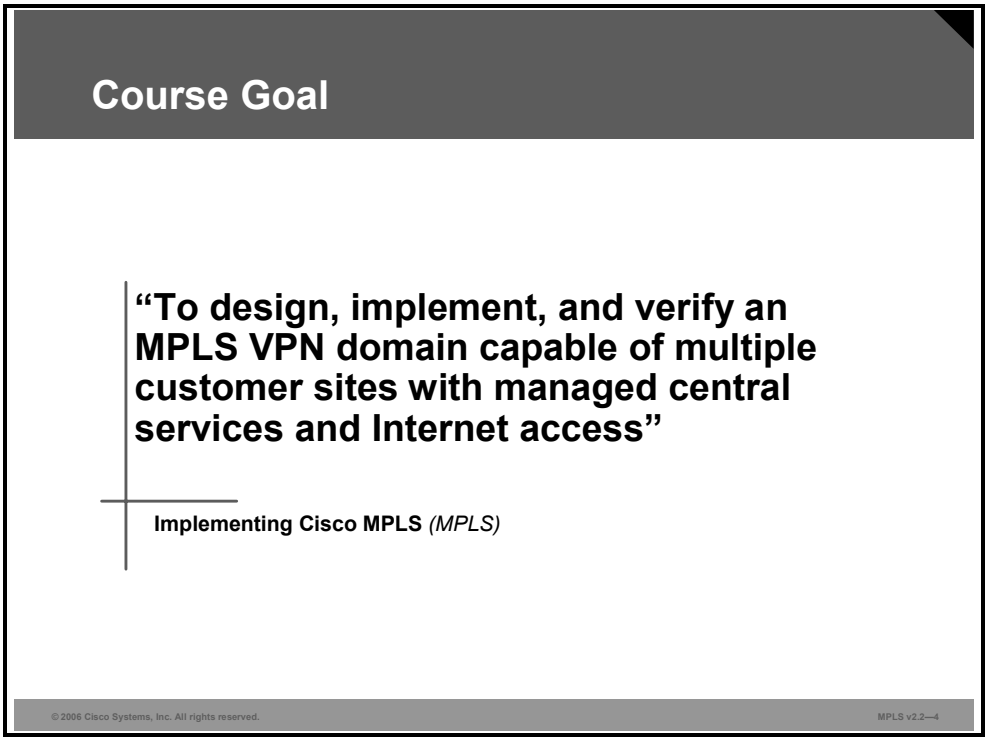

Upon completing this course, you will be able to meet these objectives:

- Describe the features of MPLS
- Describe how MPLS labels are assigned and distributed
- Configure and troubleshoot MPLS on frame-mode Cisco IOS platforms
- Describe the MPLS peer-to-peer architecture and explain the routing and packet-forwarding model in this architecture
- Configure, monitor, and troubleshoot VPN operations
- Describe how the overlapping model can be used to implement managed services and Internet access
- Describe the various Internet access implementations that are available and the benefits and drawbacks of each model; configure, monitor, and troubleshoot basic Internet access
- Configure, monitor, and troubleshoot basic MPLS TE functions

Course Flow

This topic presents the suggested flow of the course materials.

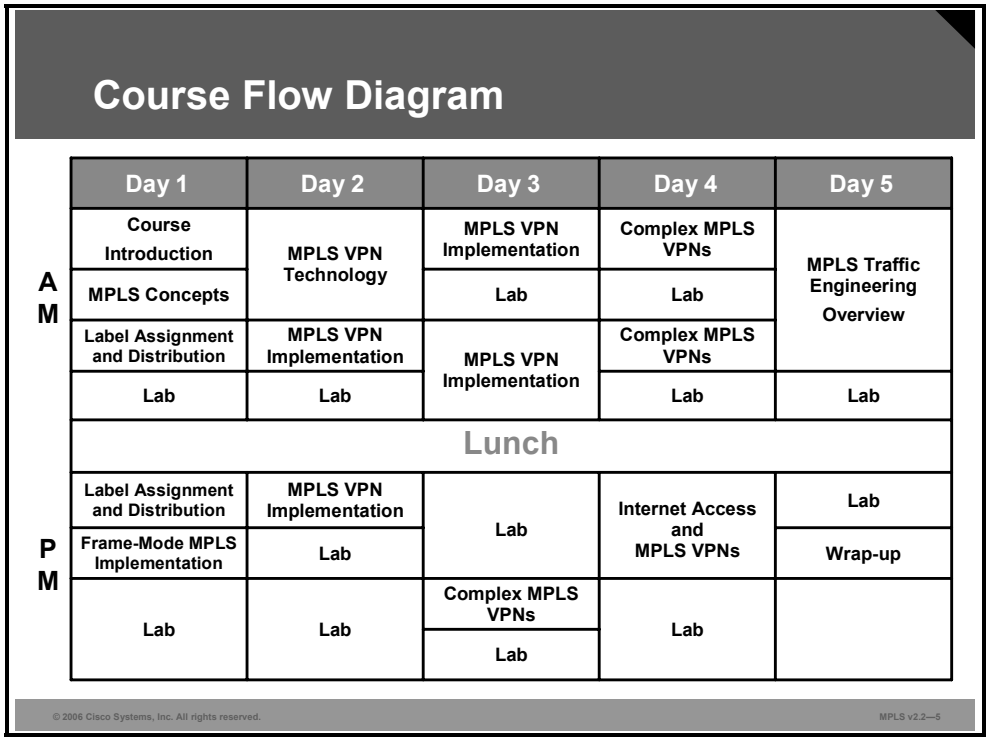

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.

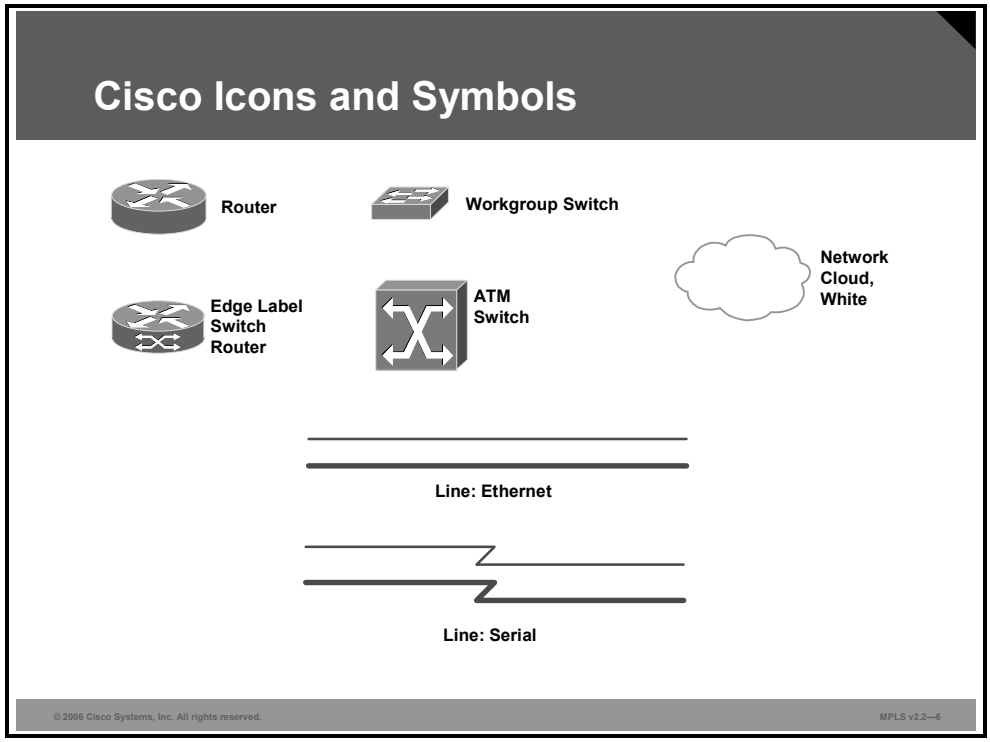

Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/index.htm.

Your Training Curriculum

This topic presents the training curriculum for this course.

You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, or CCSP™). It provides a gathering place for Cisco certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit http://www.cisco.com/en/US/learning/le3/le2/le41/learning_certification_level_home.html.

Module 1

MPLS Concepts

Overview

This module explains the features of Multiprotocol Label Switching (MPLS) compared with those of traditional hop-by-hop IP routing. MPLS concepts and terminology, along with MPLS label format and label switch router (LSR) architecture and operations, are explained in this module.

Module Objectives

Upon completing this module, you will be able to describe the features of MPLS. This ability includes being able to meet these objectives:

- Describe the basic MPLS concepts
- Describe the structure and function of MPLS labels and MPLS label stacks
- Describe the different MPLS applications in which you can use MPLS

Lesson 1

Introducing Basic MPLS Concepts

Overview

This lesson discusses the basic concepts and architecture of Multiprotocol Label Switching (MPLS). The lesson provides information about MPLS components and labels. This lesson lays the foundation for subsequent lessons that cover key areas, such as MPLS implementations and Virtual Private Networks (VPNs).

It is important to have a clear understanding of the role of MPLS and the makeup of the devices and components. This understanding will help you have a clear picture of how to differentiate between the roles of certain devices and to understand how information gets transferred across an MPLS domain.

Objectives

Upon completing this lesson, you will be able to describe the basic MPLS concepts, including some advantages as compared to traditional IP routing. This ability includes being able to meet these objectives:

- Describe the foundations of traditional IP routing
- Describe the basic features of MPLS
- Describe the benefits of MPLS
- Describe the main components of the MPLS architecture
- Describe the function of the different types of LSRs

What Are the Foundations of Traditional IP Routing?

This topic describes the foundations of traditional IP routing.

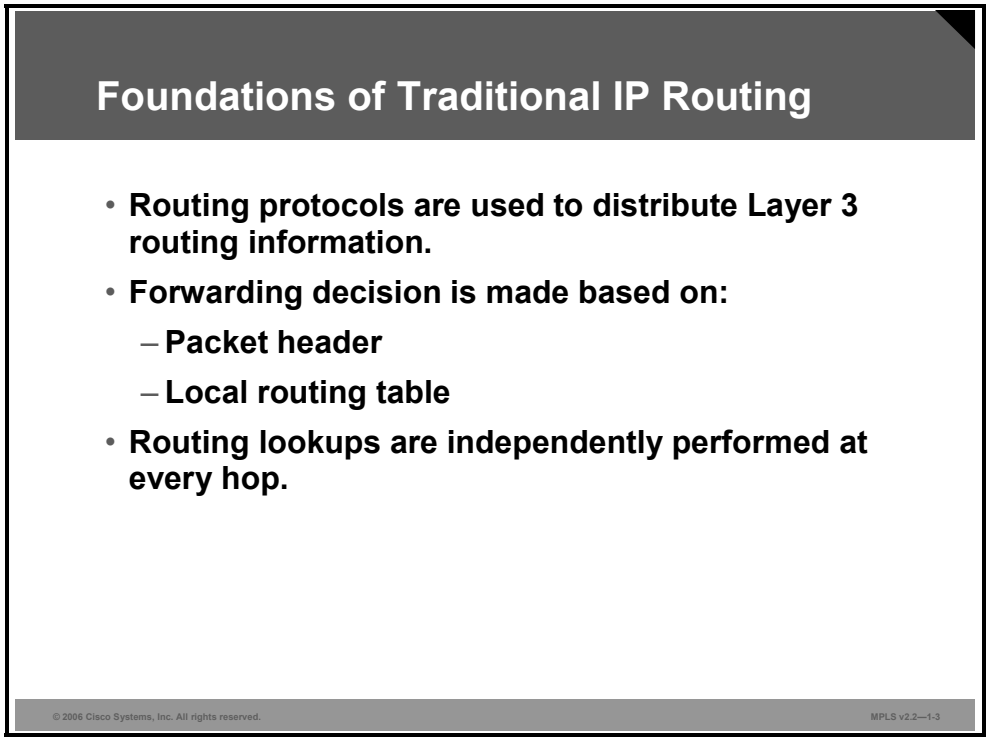

Before basic MPLS functionality is explained, these three foundations of traditional IP routing need to be highlighted:

- Routing protocols are used on all devices to distribute routing information.
- Each router analyzes the Layer 3 header of each packet compared to the local routing table and makes a decision about where to forward the packet. Regardless of the routing protocol, routers forward packets contingent on a destination address-based routing lookup.

Note	The exception to this rule is policy-based routing (PBR), where routers will bypass the destination-based routing lookup.

- The routing lookup is performed independently on every router in the network.

Example: Traditional IP Routing

This diagram shows traditional IP routing.

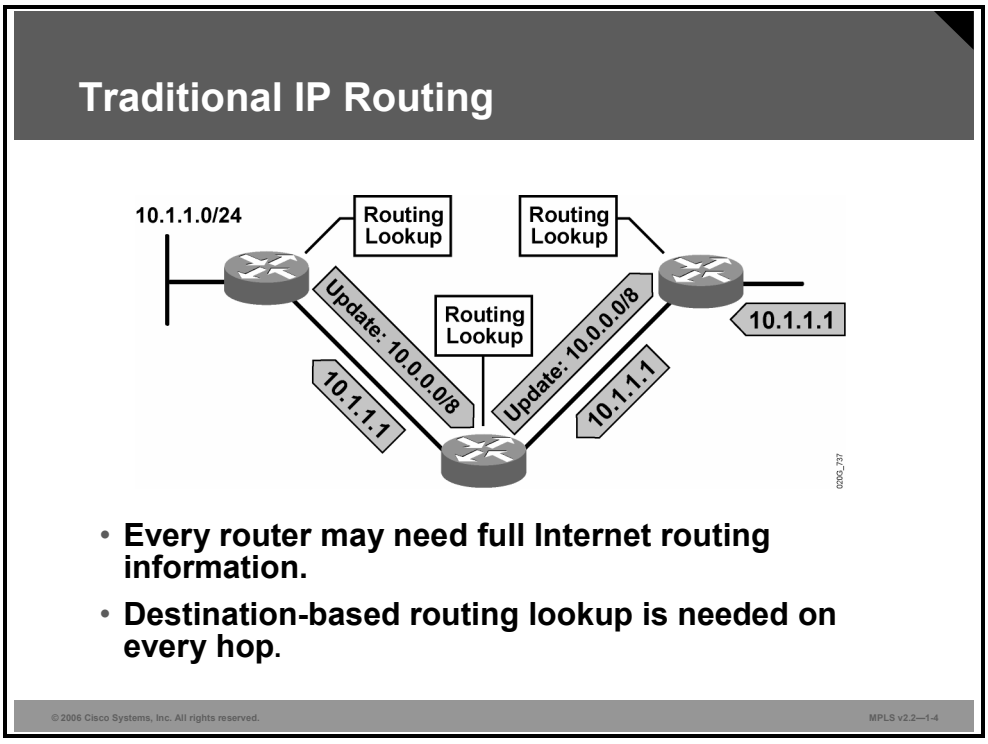

Basic MPLS Features

This topic describes the basic features of MPLS.

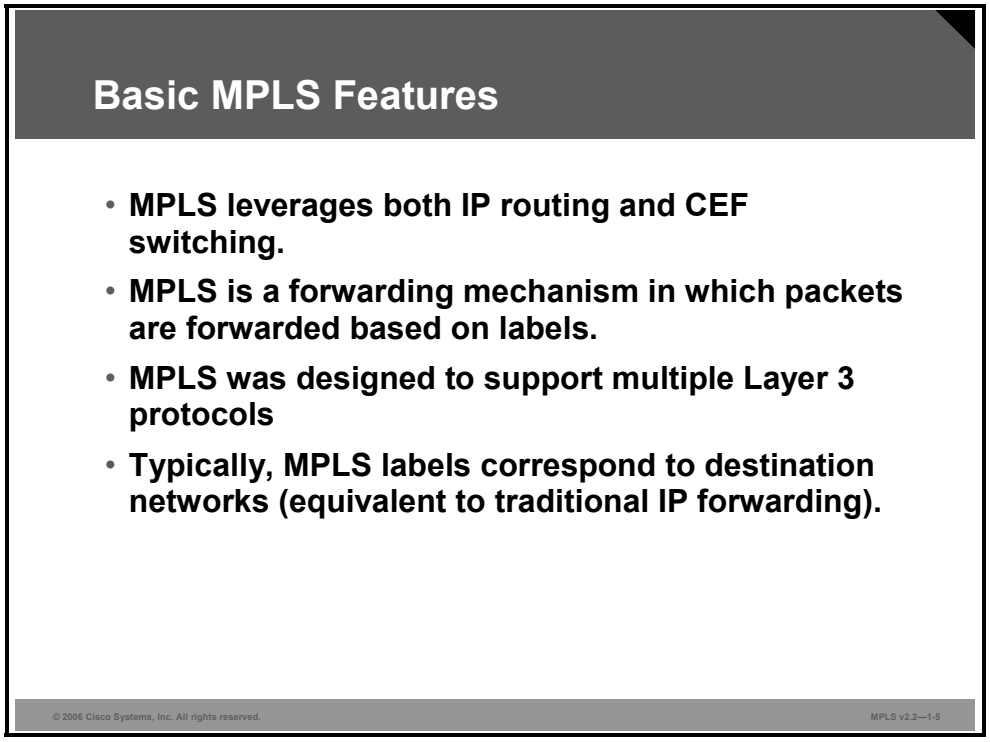

MPLS is designed to leverage the intelligence associated with IP routing and the switching model associated with Cisco Express Forwarding (CEF) switching.

Note	CEF switching will be discussed in detail in the "Frame-Mode MPLS Implementation on Cisco IOS Platforms" module. In summary, CEF uses a complete IP switching table, the Forwarding Information Base (FIB) table, to make forwarding decisions. Because the FIB contains the complete IP switching table, the router can make definitive forwarding decisions based on the information in it.

MPLS is a packet-forwarding technology that uses appended labels to make forwarding decisions for packets.

- Within the MPLS network, the Layer 3 header analysis is done just once (when the packet enters the MPLS domain). Labels are appended to the packet, and then the packet is forwarded into the MPLS domain.
- Simple label inspection integrated with CEF switching drives subsequent packet forwarding.

MPLS was designed to support efficiently forwarding packets across the network core based on a simplified header. Label switching is performed regardless of the Layer 3 routing protocol.

MPLS labels typically correspond to Layer 3 destination addresses (equal to destination-based routing). Labels can also correspond to other parameters, such as quality of service (QoS), source address, or a Layer 2 circuit. An MPLS label is a short, 4-byte, fixed-length, locally significant identifier.

Benefits of MPLS

This topic describes some of the benefits of MPLS.

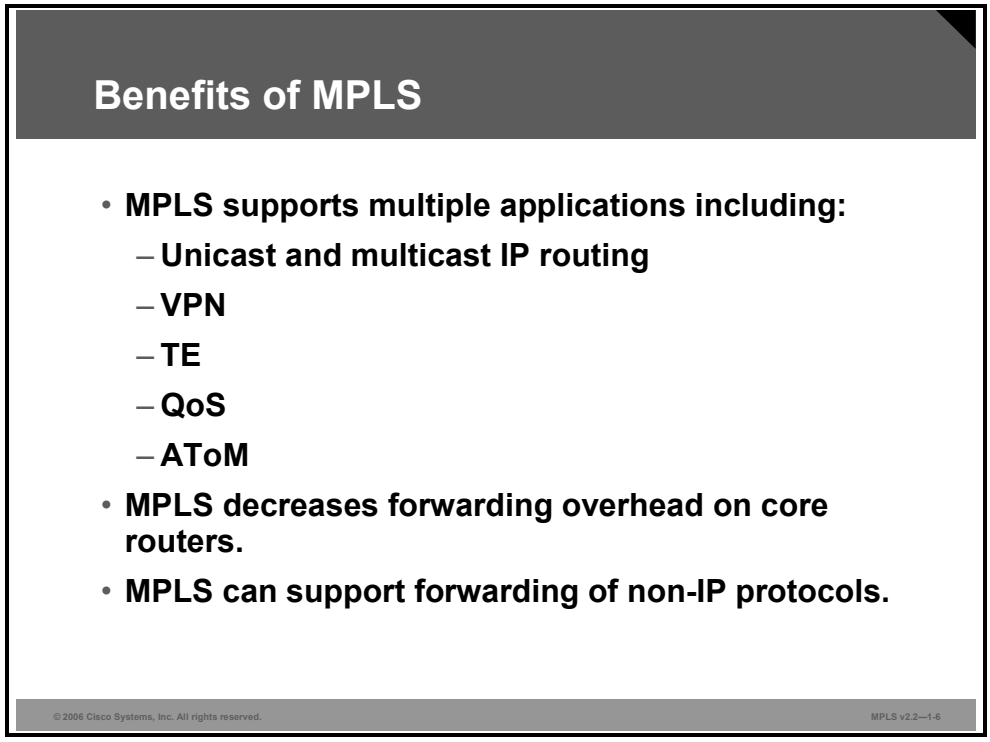

There are several benefits to MPLS:

- MPLS decreases the forwarding overhead on the core routers.
- MPLS supports multiple useful applications such as those listed here:
  - Unicast and multicast IP routing
  - VPN
  - Traffic engineering (TE)
  - QoS
  - Any Transport over MPLS (AToM).

Note An overview of these applications will be provided in the “Identifying MPLS Applications” lesson in this module.

- MPLS supports the forwarding of non-IP protocols, because MPLS technologies are applicable to any network layer protocol.

What Are the MPLS Architecture Components?

MPLS consists of these two major components:

- Control plane
- Data plane

MPLS Control Plane

The control plane takes care of the routing information exchange and the label exchange between adjacent devices.

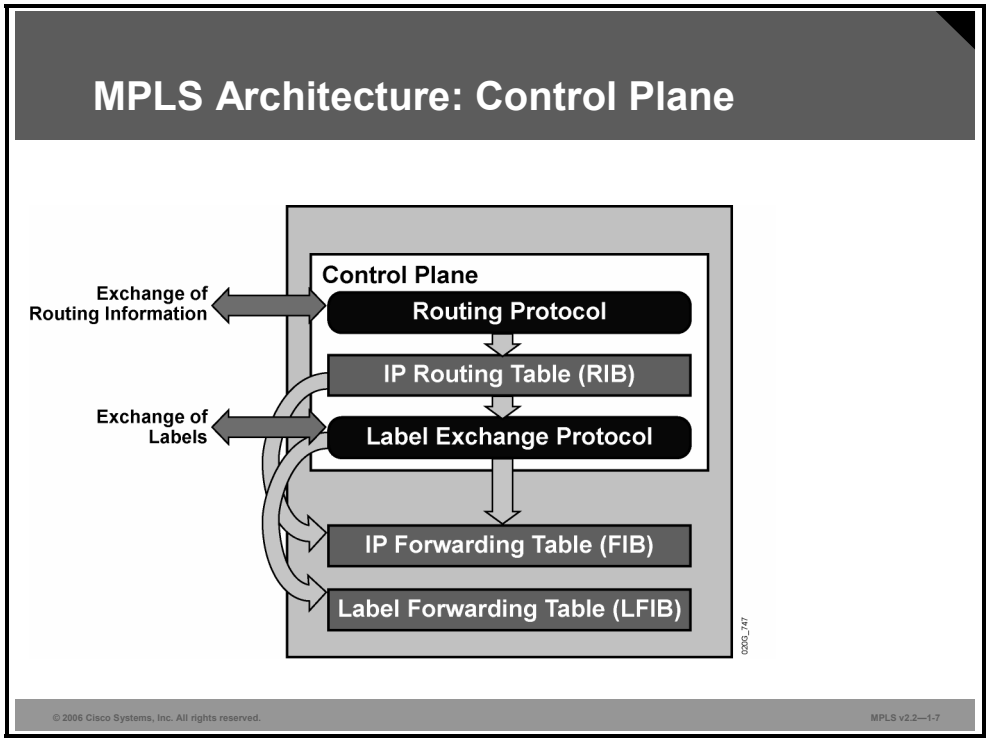

The control plane builds a routing table (Routing Information Base [RIB]) based on the routing protocol. Various routing protocols, such as Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP), can be used in the control plane for managing Layer 3 routing.

The control plane uses a label exchange protocol to create and maintain labels internally, and to exchange these labels with other devices. The label exchange protocol binds labels to networks learned via a routing protocol. Label exchange protocols include MPLS Label Distribution Protocol (LDP), the older Cisco Tag Distribution Protocol (TDP), and BGP (used by MPLS VPN). Resource Reservation Protocol (RSVP) is used by MPLS TE to accomplish label exchange.

The control plane also builds two forwarding tables, a FIB from the information in the RIB, and a label forwarding information base (LFIB) table based on the label exchange protocol and the RIB. The LFIB table includes label values and associations with the outgoing interface for every network prefix.

MPLS Data Plane

The data plane takes care of forwarding based on either destination addresses or labels; the data plane is also known as the forwarding plane.

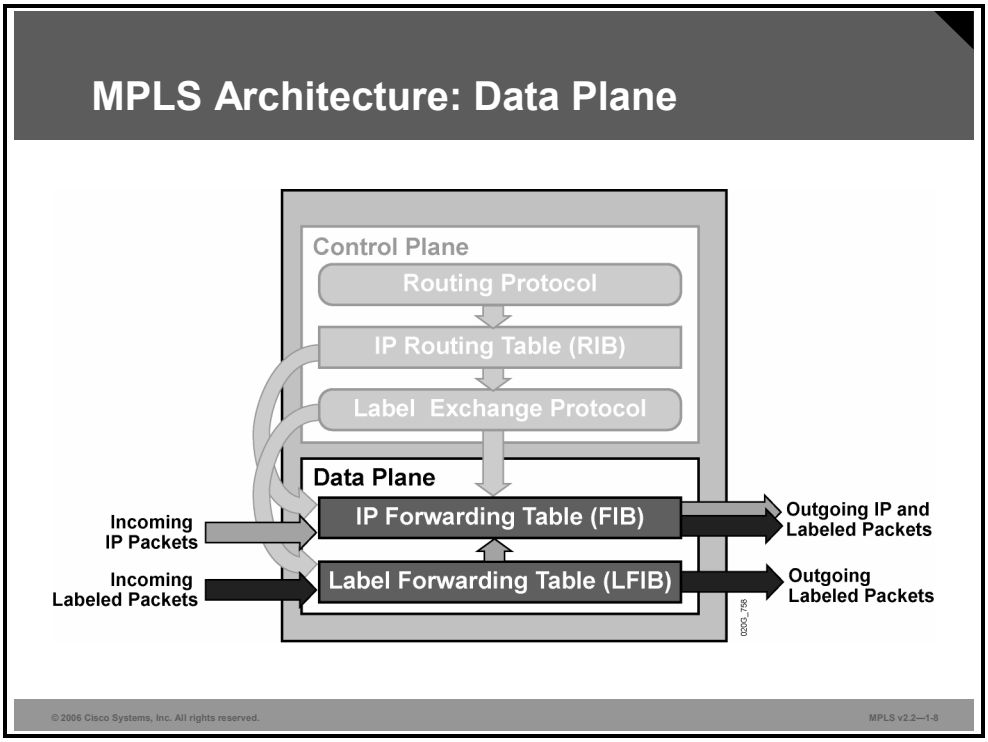

The data plane is a simple forwarding engine that is independent of the type of routing protocol or label exchange protocol being used. The data plane forwards packets to the appropriate interface based on the information in the LFIB or the FIB tables.

MPLS LSRs

This topic describes the function of two types of MPLS devices.

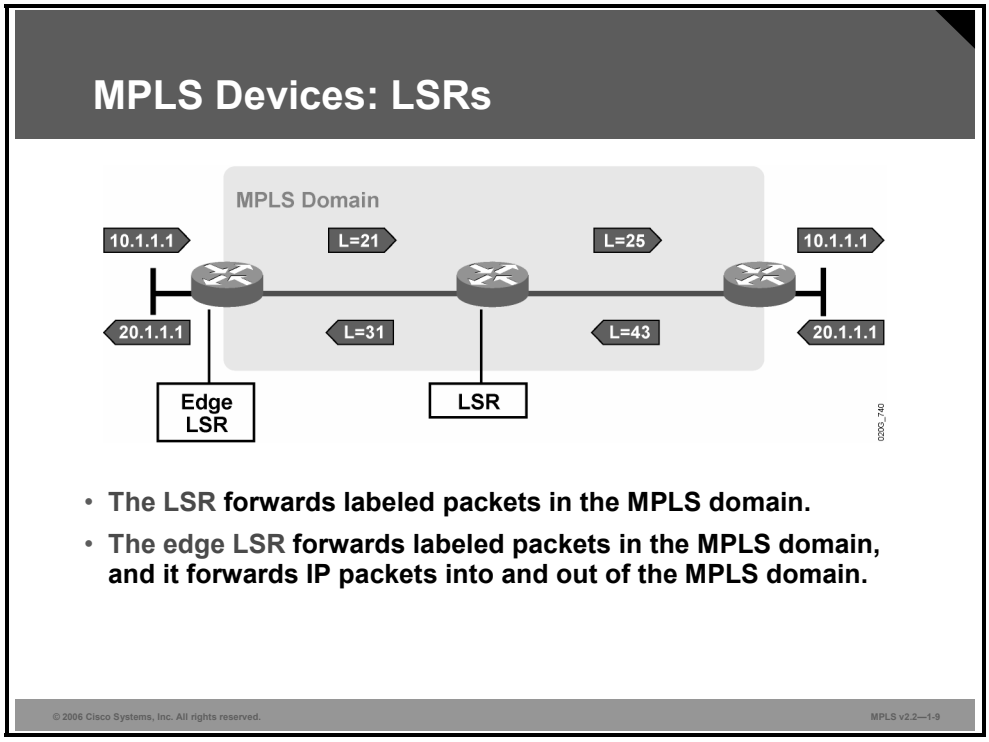

The label switch router (LSR) is the basic MPLS device used for most MPLS applications. Here are two definitions:

- **LSR:** A device that implements label distribution procedures and primarily forwards packets based on labels
- **Edge LSR:** An LSR on the edge of an MPLS domain that implements label distribution procedures, forwards packets based on labels, and primarily inserts labels on packets or removes labels for non-MPLS devices

LSRs and edge LSRs are usually capable of doing both label switching and IP routing. Their names are based on their positions in an MPLS domain. Routers that have all interfaces enabled for MPLS are called LSRs because they mostly forward labeled packets. Routers that have some interfaces that are not enabled for MPLS are usually at the edge of an MPLS domain. These routers also forward packets based on IP destination addresses and label them if the outgoing interface is enabled for MPLS.

Note	In a service provider MPLS environment, an edge LSR is typically known as a provider edge (PE) router, and an LSR is known as a provider (P) router.

Label Switch Routers: Architecture of LSRs

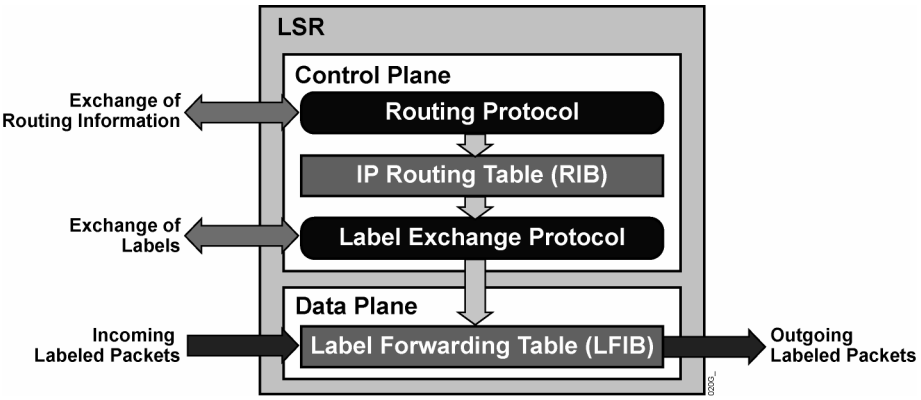

The primary LSR functions are to exchange labels with other LSRs and to forward labeled packets. Therefore, every LSR needs a Layer 3 routing protocol (for example, OSPF, EIGRP, or IS-IS), and a label exchange protocol (for example, LDP or TDP).

LDP populates the LFIB table in the data plane that is used to forward labeled packets.

Note	LSRs may not be able to forward unlabeled packets if they do not have sufficient routing information.

Example: LSR Architecture

The figure illustrates the main components of the control and data planes for MPLS operation on a LSR.

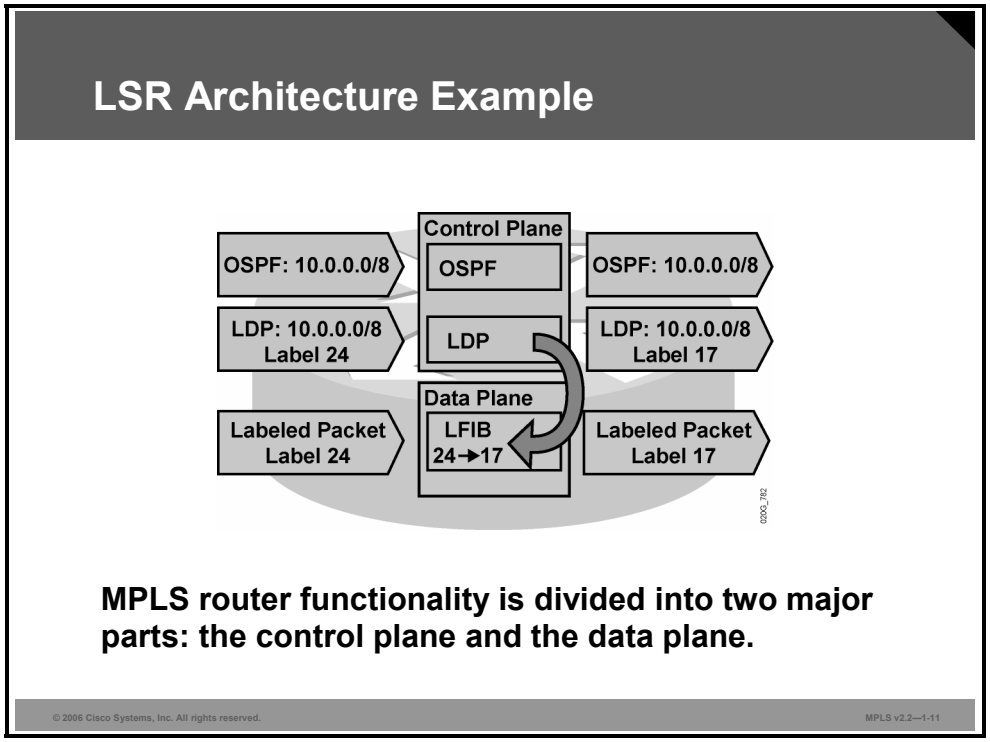

In the example LSR architecture, the control plane uses these protocols:

- A routing protocol (OSPF), which receives and forwards information about IP network 10.0.0.0/8
- A label exchange protocol (LDP), which receives label 24 to be used for packets with destination address 10.0.0.0/8

(A local label 17 is generated and is sent to upstream neighbors so that these neighbors can label packets with the appropriate label.)

The data plane uses an LFIB to forward packets based on labels:

- The LFIB receives an entry from LDP, where label 24 is mapped to label 17. When the data plane receives a packet labeled with a 24, it replaces label 24 with label 17 and forwards the packet through the appropriate interfaces.

Note	In the example, both packet flow and routing and label updates are from left to right.

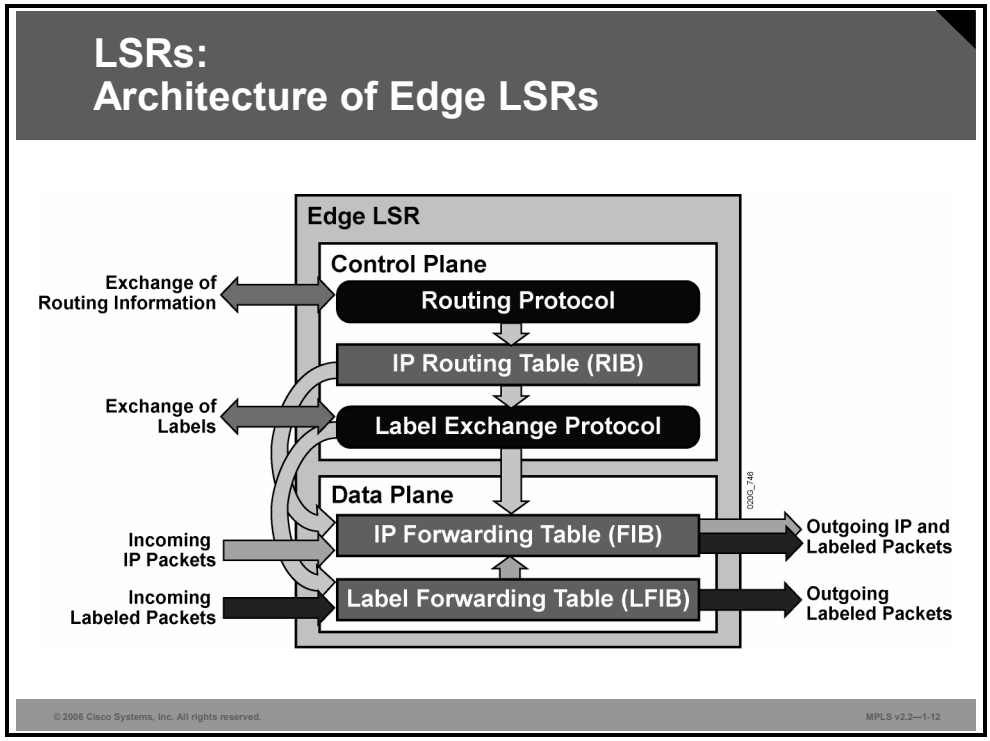

Besides forwarding labeled packets, edge LSRs also forward IP packets into and out of the MPLS domain.

These combinations are possible on an edge LSR:

- A received IP packet is forwarded based on the IP destination address (sent as an IP packet.)
- A received IP packet is labeled based on the IP destination address and is forwarded as a labeled packet.
- A received labeled packet is forwarded after the label is swapped (sent as a labeled packet).
- A received labeled packet is forwarded as an IP packet after the label is removed.

These scenarios are possible if the network is not configured properly:

- A received labeled packet is dropped if the label is not found in the LFIB table, even if the IP destination exists in the IP forwarding table—also called the FIB.
- A received IP packet is dropped if the destination is not found in the FIB table, even if there is an MPLS label-switched path toward the destination.

Example: Basic MPLS

The figure illustrates a situation in which the intermediary router in the MPLS core does not have to perform a time-consuming routing lookup. Instead, this router simply swaps a label with another label (25 is replaced by 23) and forwards the packet based on the swapped label (23).

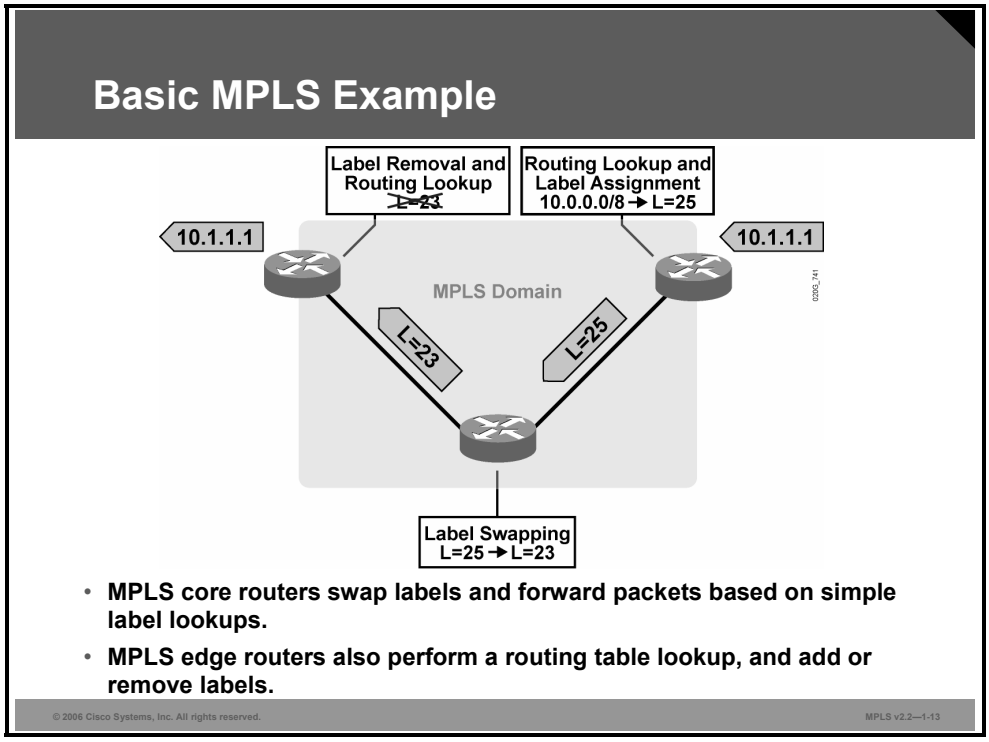

At the egress LSR, the MPLS label is removed, and an IP packet is forwarded out of the MPLS domain.

In larger networks, the result of MPLS labeling is that only the edge routers perform a routing lookup. The core routers simply forward packets based on the labels.

Summary

This topic summarizes the key points that were discussed in this lesson.

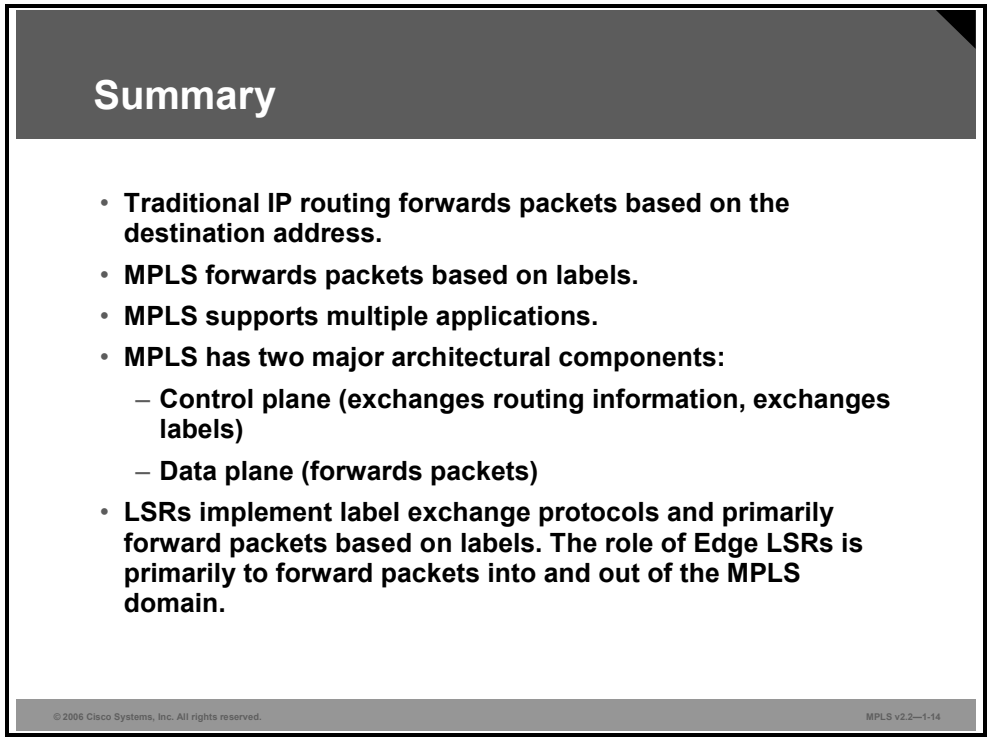

Lesson 2

Introducing MPLS Labels and Label Stacks

Overview

This lesson explains the four fields that make up a Multiprotocol Label Switching (MPLS) label. This lesson also explains how label stacking is used and how labels are forwarded in frame-mode environments.

To fully understand MPLS, it is necessary to have a clear understanding of the format of an MPLS label, and the definition for each field in that label. You also need to know exactly how information is passed from node to node in the network.

Objectives

Upon completing this lesson, you will be able to describe MPLS labels and MPLS label stacks, including the format of the MPLS label and also when and why a label stack is created. This ability includes being able to meet these objectives:

- Describe the features of MPLS labels
- Describe the format and fields of an MPLS label
- Describe where MPLS labels are inserted in an IP packet
- Describe the features of an MPLS label stack
- Describe MPLS label operations

What Are MPLS Labels?

This topic describes the features of MPLS labels.

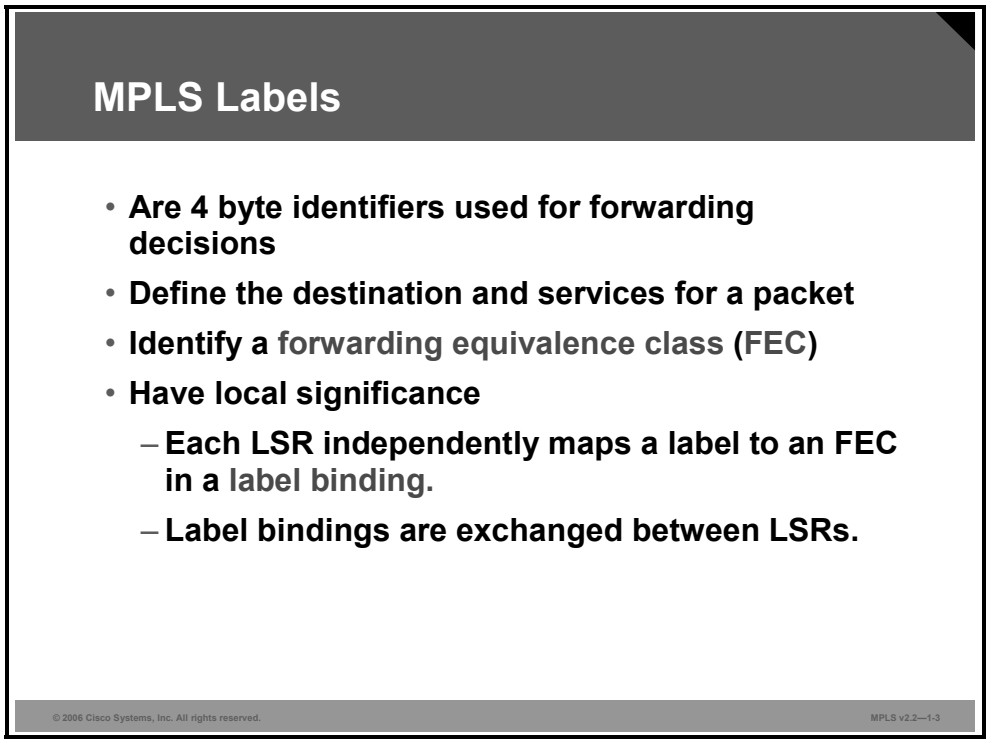

An MPLS label is a 4-byte, fixed-length, locally significant identifier that is used by network core devices to make forwarding decisions for a packet. Labels define the destination and services for each packet, and identify a forwarding equivalence class (FEC). The label put on a particular packet represents the FEC to which the packet is assigned.

Labels have local significance to a label switch router (LSR). Each LSR in the network makes an independent, local decision regarding which label value to use to represent an FEC. This mapping is known as a label binding.

Each LSR informs its neighbors of the label bindings that it has made.

Note	Details on how the label binding are exchanged will be covered in the “Label Assignment and Distribution” module.

FEC and MPLS Forwarding

An FEC is an integral part of MPLS forwarding.

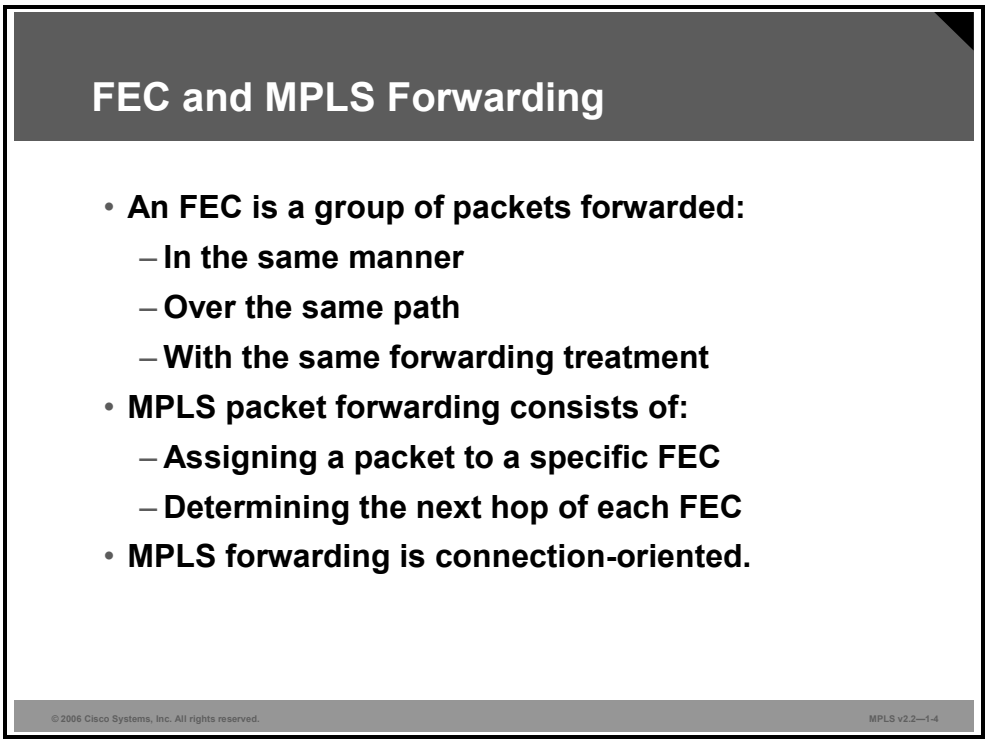

The FEC is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. An FEC might correspond to a destination IP subnetwork, but it also might correspond to any traffic class that the edge LSR considers significant. For example, all traffic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

MPLS packet forwarding consists of these two elements:

- At the ingress to the MPLS network, packets are classified and assigned to a specific FEC using a label. No further packet classification is done in the MPLS network.
- Throughout the MPLS network, all packets in an FEC are forwarded using the next-hop address for the FEC. The label value changes as the IP packet traverses the network. When a labeled packet is sent from one LSR to the next-hop LSR, the label value carried by the packet is the label value that the next-hop LSR assigned to represent the FEC of the packet.

Note	Details on MPLS forwarding will be discussed in the “Frame-Mode MPLS Implementation on Cisco IOS Platforms » module.

MPLS uses FEC-based forwarding to evolve connectionless IP networks to connection-oriented networks.

What Is the MPLS Label Format?

This topic describes the format and fields of an MPLS label.

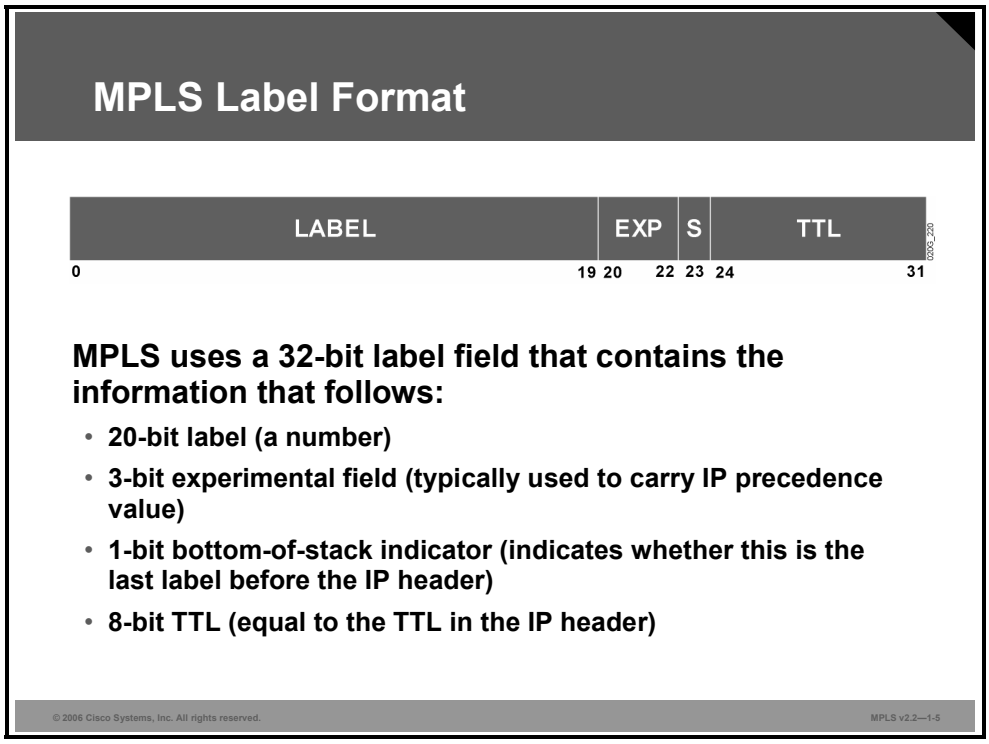

A label contains the fields listed in this table.

Label Fields

Field	Description
`20-bit label`	This is the actual label. The values 0 to 15 are reserved.
`3-bit experimental field`	This field is typically used to define a class of service (CoS) or IP precedence value.
`Bottom-of-stack bit`	MPLS allows multiple labels to be inserted; this bit determines if this label is the last label in the packet. If this bit is set (1), it indicates that this is the last label.
`8-bit TTL field`	This field has the same purpose as the time-to-live (TTL) field in the IP header; it is used to prevent the indefinite looping of packets.

Where Are MPLS Labels Inserted?

This topic describes where MPLS labels are inserted in an IP packet.

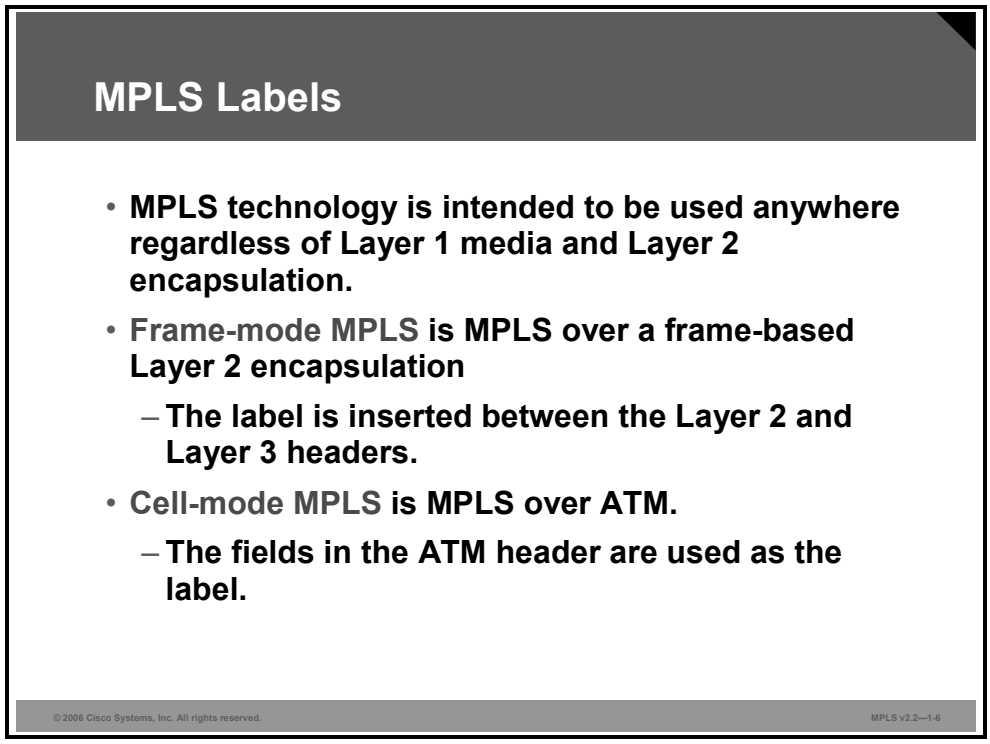

MPLS is designed for use on virtually any media and Layer 2 encapsulation.

Frame-mode MPLS is the typical mode of MPLS, because most Layer 2 encapsulations are frame-based. In frame-mode MPLS, the 32-bit label is inserted between the Layer 2 and Layer 3 headers.

ATM is a special case of Layer 2 encapsulation where fixed-length cells are used.

Note	If you are using ATM as a WAN link and the ATM switches do not act as LSRs, you are still running frame-mode MPLS.

Cell-mode MPLS is MPLS using ATM Layer 2 encapsulation, where the ATM switch is participating as an LSR. In cell-mode MPLS, a label cannot be inserted on every cell; therefore, the virtual path identifier/virtual channel identifier (VPI/VCI) fields in the ATM header are used as a label.

Note	Cell-mode MPLS is briefly mentioned here for reference purposes. This course will focus on frame-mode MPLS, which is more prevalent in the industry.

Example: MPLS Label Insertion—Frame-Mode MPLS

The figure shows an edge router that receives a normal IP packet.

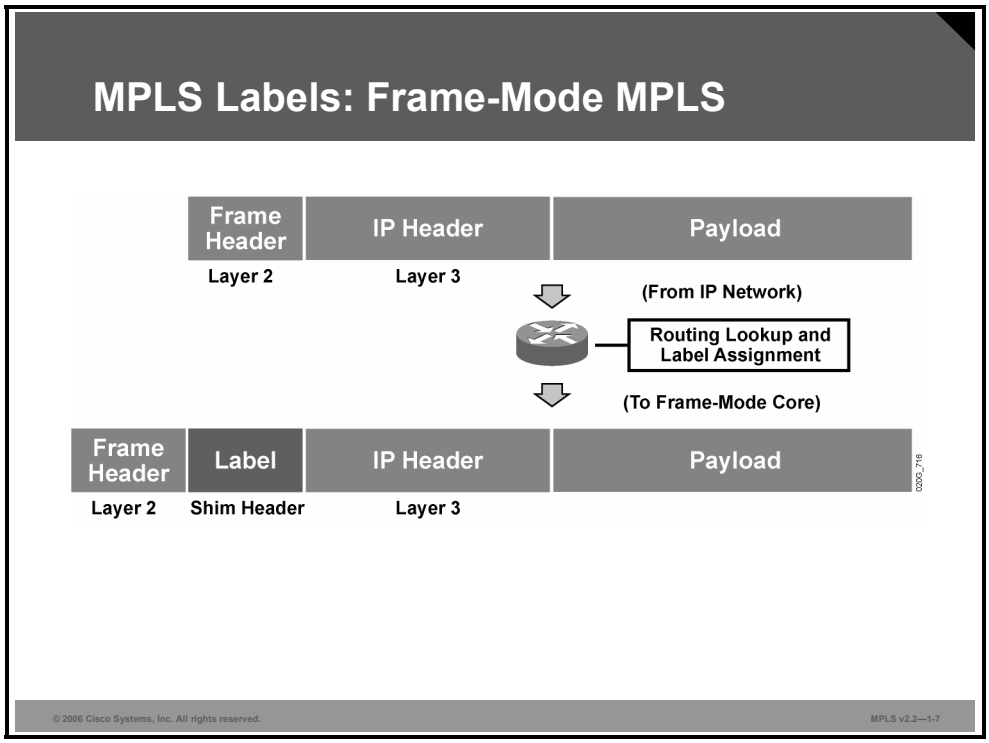

The edge LSR then does these tasks:

- The router performs routing lookup to determine the outgoing interface.
- The router assigns and inserts a label between the Layer 2 frame header and the Layer 3 packet header, if the outgoing interface is enabled for MPLS and if a next-hop label for the destination exists. This inserted label is also called the shim header.
- The router then changes the Layer 2 protocol identifier (PID) or Ethertype value in the Layer 2 frame header to indicate that this is a labeled packet.
- The router sends the labeled packet.

Note	Other routers in the MPLS core simply forward packets based on the received label.

What Is an MPLS Label Stack?

This topic describes the features of an MPLS label stack.

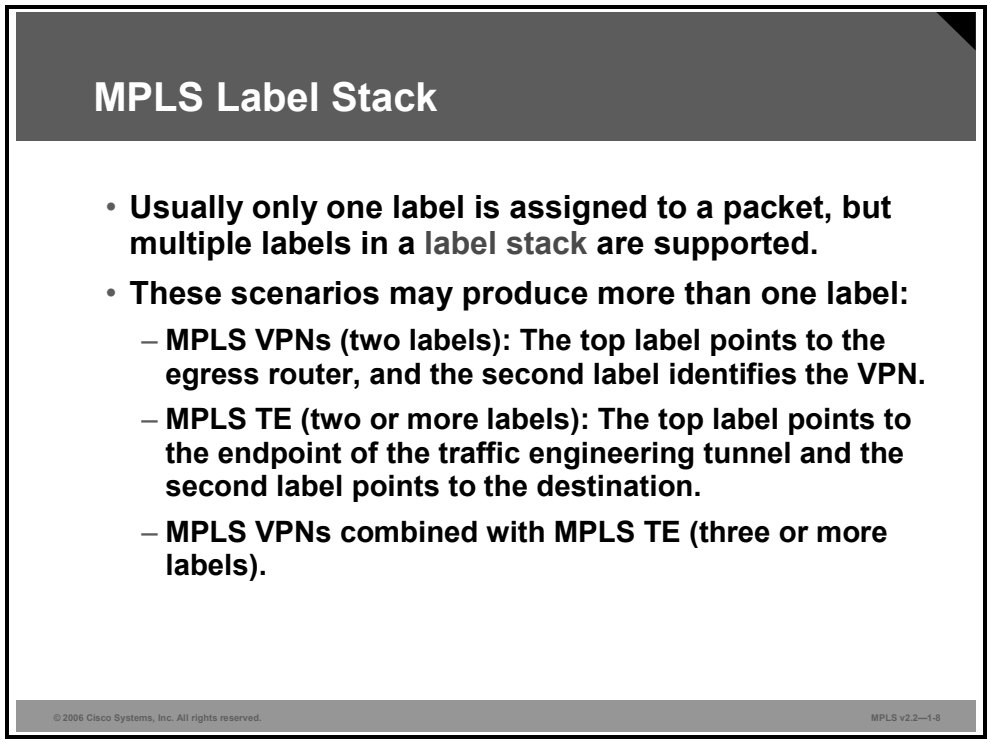

Simple MPLS uses just one label in each packet. However, MPLS does allow multiple labels in a label stack to be inserted in a packet.

These applications may add labels to packets:

- **MPLS Virtual Private Networks (VPNs):** With MPLS VPNs, Multiprotocol Border Gateway Protocol (MP-BGP) is used to propagate a second label that is used in addition to the one propagated by Label Distribution Protocol (LDP) or Tag Distribution Protocol (TDP).
- **Cisco MPLS Traffic Engineering (MPLS TE):** MPLS TE uses Resource Reservation Protocol (RSVP) to establish label-switched path (LSP) tunnels. RSVP also propagates labels that are used in addition to the one propagated by LDP or TDP.

A combination of these mechanisms and other advanced features might result in three or more labels being inserted into one packet.

Example: MPLS Label Stack

This figure shows multiple labels in an MPLS label stack.

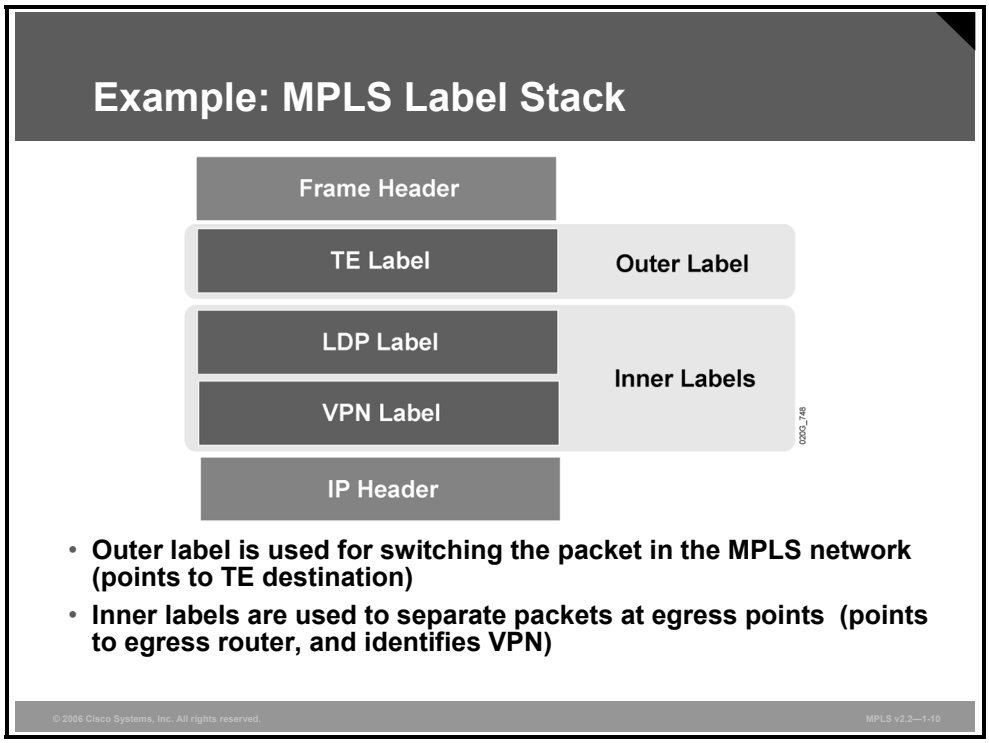

The outer label is used to switch the MPLS packet across the network. In this case, the outer layer is a traffic engineering (TE) label pointing to the endpoint of a TE tunnel.

The inner labels are ignored by the intermediary routers. In this case, the inner labels are used to point to the egress router and to identify the VPN for the packet.

Example: MPLS Label Stack Format

This figure shows the format of multiple labels in a frame-mode MPLS label stack.

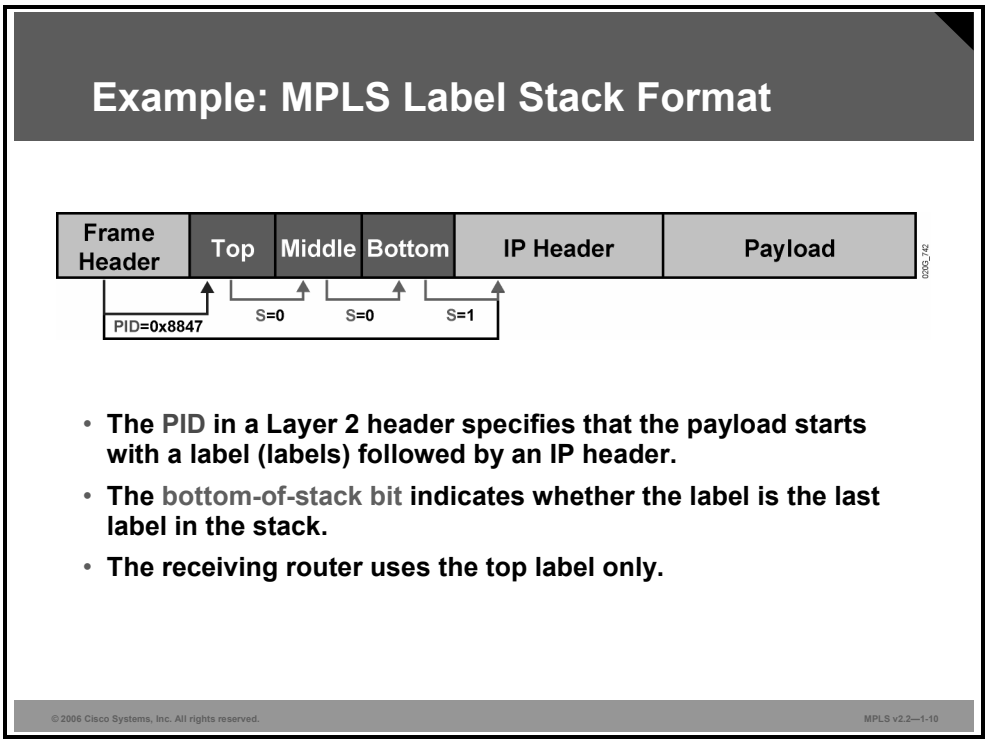

An MPLS PID in the Layer 2 header is used to identify every MPLS packet type. These Ethertype values are used to identify Layer 3 protocols with most Layer 2 encapsulations:

- **Unlabeled IP unicast:** PID = 0x0800 identifies that the frame payload is an IP packet.
- **Labeled IP unicast:** PID = 0x8847 identifies that the frame payload is a unicast IP packet with at least one label preceding the IP header. The bottom-of-stack bit indicates when the IP header actually starts.
- **Labeled IP multicast:** PID = 0x8848 identifies that the frame payload is a multicast IP packet with at least one label preceding the IP header. The bottom-of-stack bit indicates when the IP header actually starts.

The top label of the label stack appears first in the packet, and the bottom label appears last. The bottom-of-stack bit in each MPLS label defines if the label is the last label in the packet. If this bit is set to 1, it indicates that this is the last label.

What Are MPLS Label Operations?

This topic describes MPLS label operations used when forwarding packets.

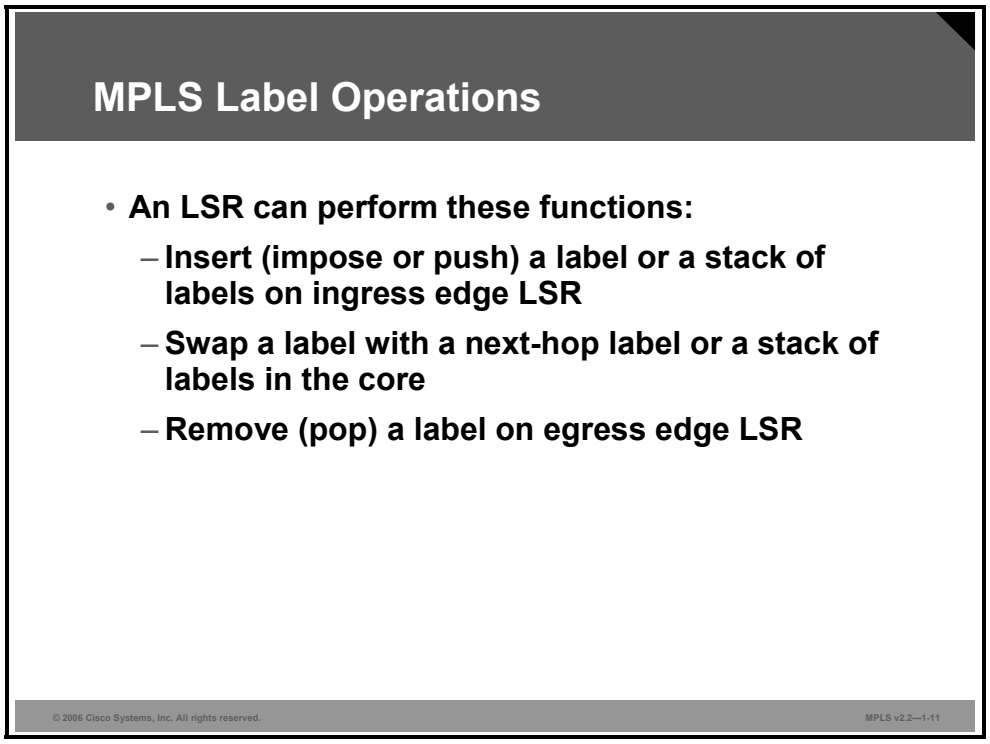

An IP packet going through an MPLS domain experiences the following:

- On an ingress edge LSR, a label or a stack of labels is inserted (imposed). The label corresponds to the assigned FEC for the packet.
- On a core or interior LSR, the top label is swapped with a next-hop label or a stack of labels.
- At the egress edge LSR, the label is removed.

Example: MPLS Label Operations—Frame-Mode MPLS

This figure shows label operations in an MPLS network using frame-mode MPLS.

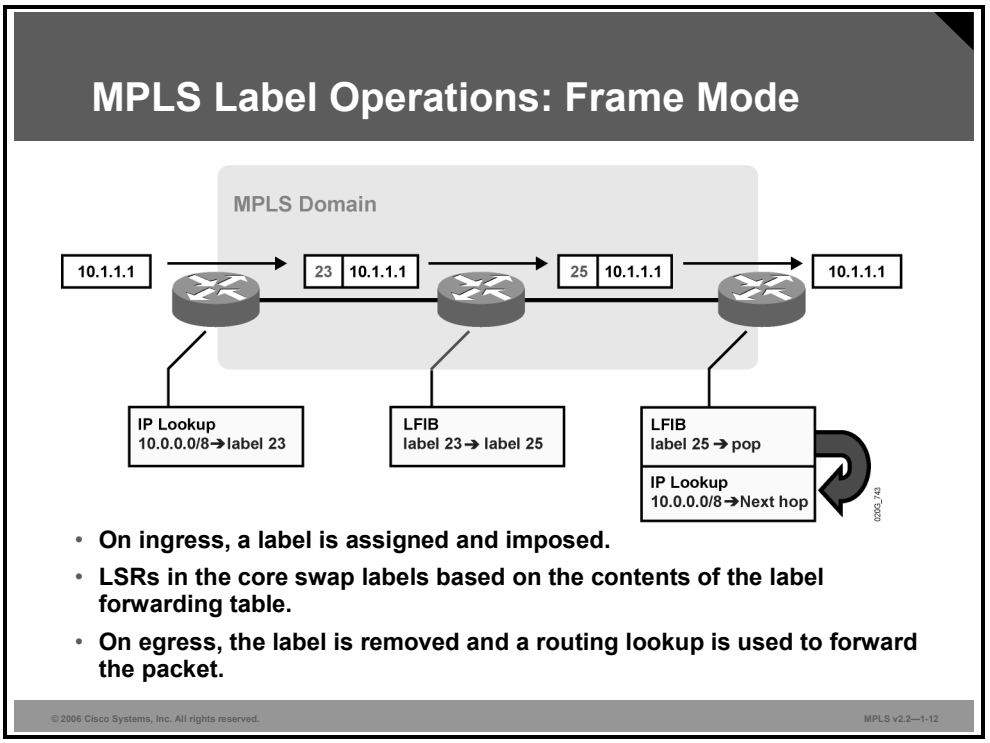

All LSRs are capable of forwarding IP packets or labeled packets.

In this example, the ingress edge LSR performs a routing lookup and assigns a label.

The middle router simply swaps the label.

The egress edge LSR removes the label and optionally performs a routing lookup.

Summary

This topic summarizes the key points that were discussed in this lesson.

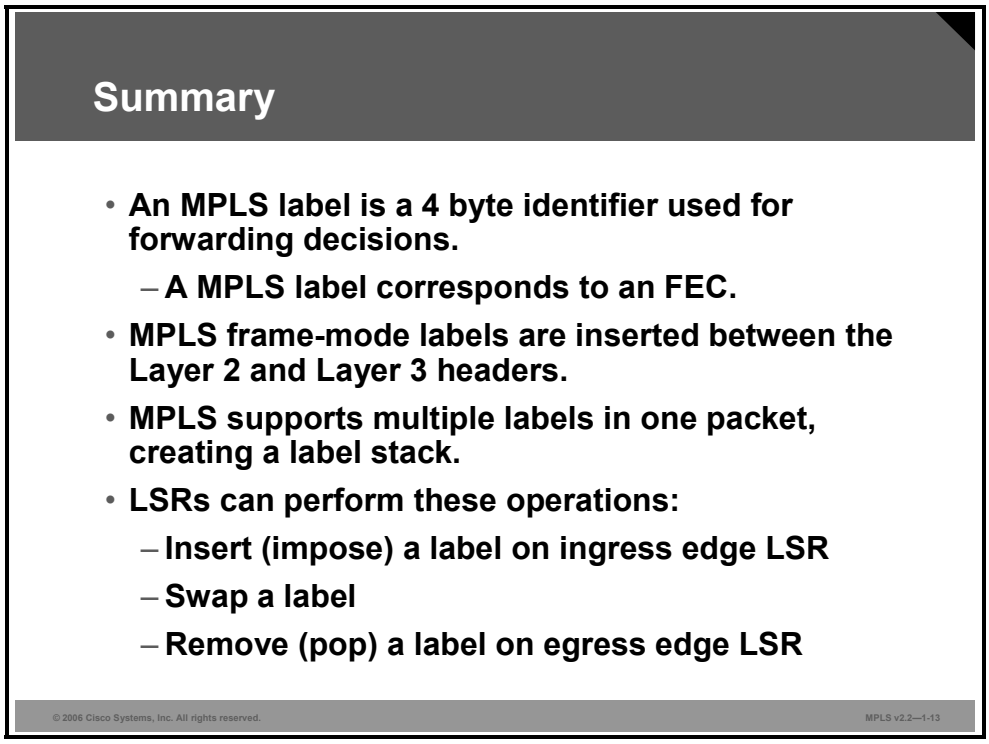

Lesson 3

Identifying MPLS Applications

Overview

This lesson describes some of the different types of applications with which you can use Multiprotocol Label Switching (MPLS). These applications are discussed at a high level. Interaction among multiple applications is also discussed because there are various methods for exchanging labels. Regardless of the differences in the control plane, all of the applications use a single label-forwarding engine in the data plane.

Objectives

Upon completing this lesson, you will be able to describe the different MPLS applications with which you can use MPLS. This ability includes being able to meet these objectives:

- Describe the various applications that are used with MPLS
- Describe the features of MPLS unicast IP routing
- Describe the features of MPLS multicast IP routing
- Describe MPLS use in VPNs
- Describe MPLS use in TE environments
- Describe MPLS use in QoS environments
- Describe AToM
- Identify the interactions that occur between various MPLS applications

Which Applications Are Used with MPLS?

This topic describes various applications that are used with MPLS.

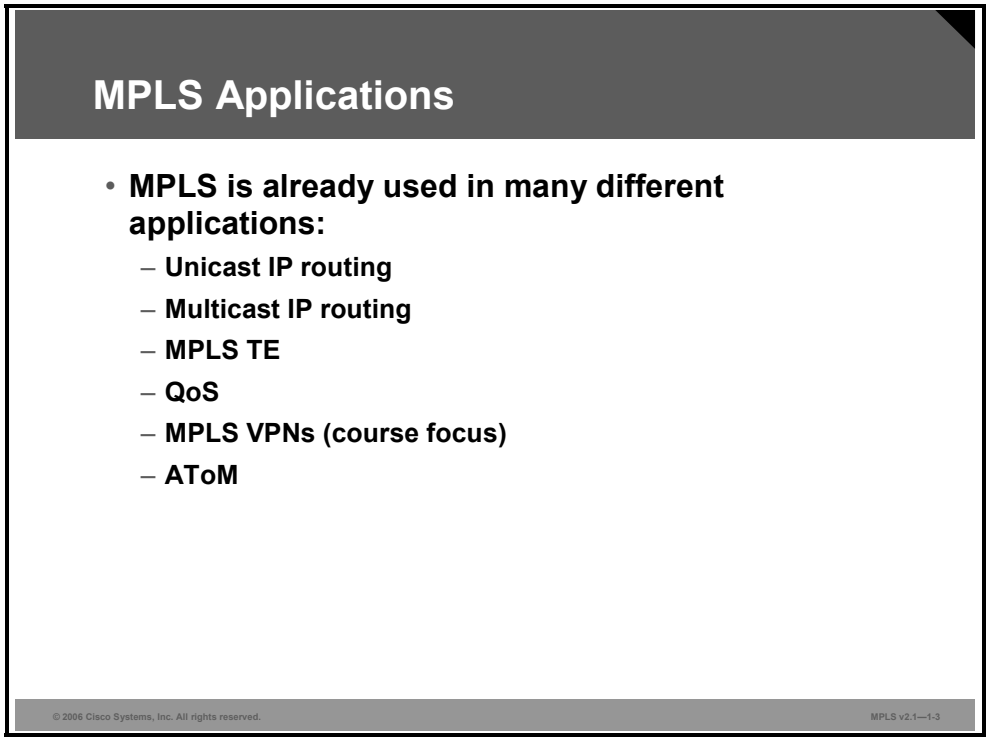

MPLS is a technology used for the delivery of IP services. MPLS can be used in different applications, as outlined here:

- Unicast IP routing is the most common application for MPLS.
- Multicast IP routing is treated separately because of different forwarding requirements.
- MPLS TE is an add-on to MPLS that provides better and more intelligent link use.
- Differentiated QoS can also be provided with MPLS.
- MPLS VPNs are implemented using labels to allow overlapping address space between VPNs. MPLS VPN is the focus of this course.
- AToM is a solution for transporting Layer 2 packets over an IP or MPLS backbone.

What Is MPLS Unicast IP Routing?

This topic describes the features of MPLS unicast IP routing.

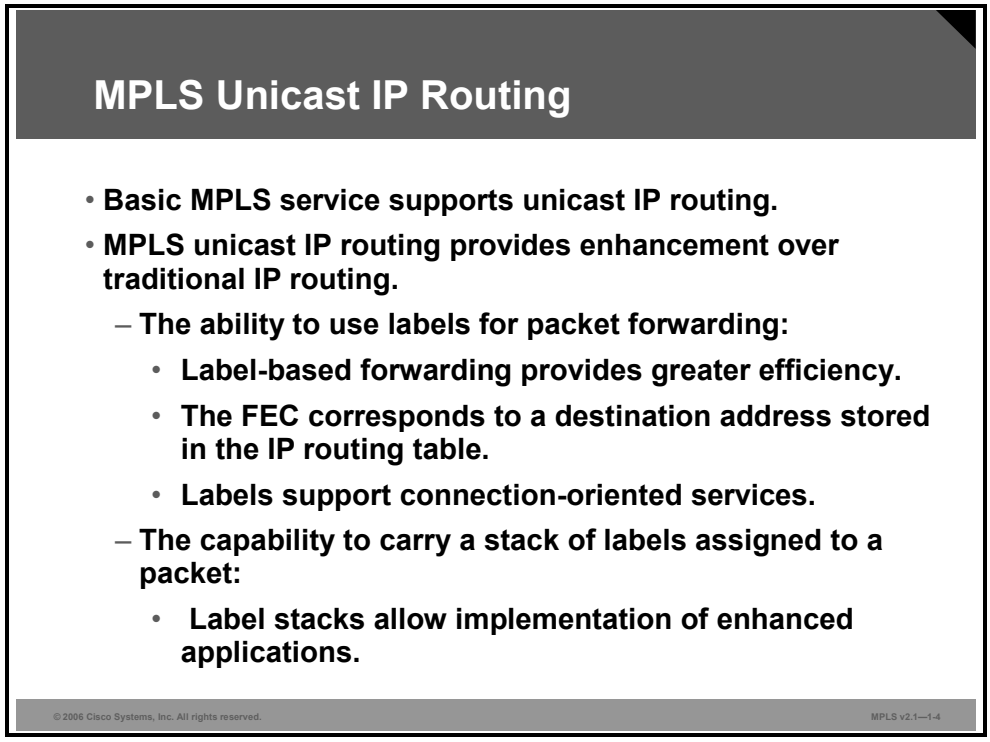

Basic MPLS supports unicast IP routing.

There are two significant enhancements that MPLS unicast IP routing provides over traditional IP routing:

- The ability to use labels for packet forwarding
- The capability to carry a stack of labels assigned to a packet

As discussed in the "Introducing Basic MPLS Concepts" lesson, using labels for packet forwarding increases efficiency in network core devices because the label swapping operation is less CPU intensive than a routing lookup. MPLS can also provide connection-oriented services to IP traffic due to forwarding equivalence class (FEC)-based forwarding.

Note	The MPLS unicast IP traffic FEC corresponds to a destination network stored in the IP routing table.

MPLS support for a label stack allows implementation of enhanced applications, such as Virtual Private Networks (VPNs), traffic engineering (TE), and enhanced quality of service (QoS).

What Is MPLS Multicast IP Routing?

This topic describes the features of MPLS multicast IP routing.

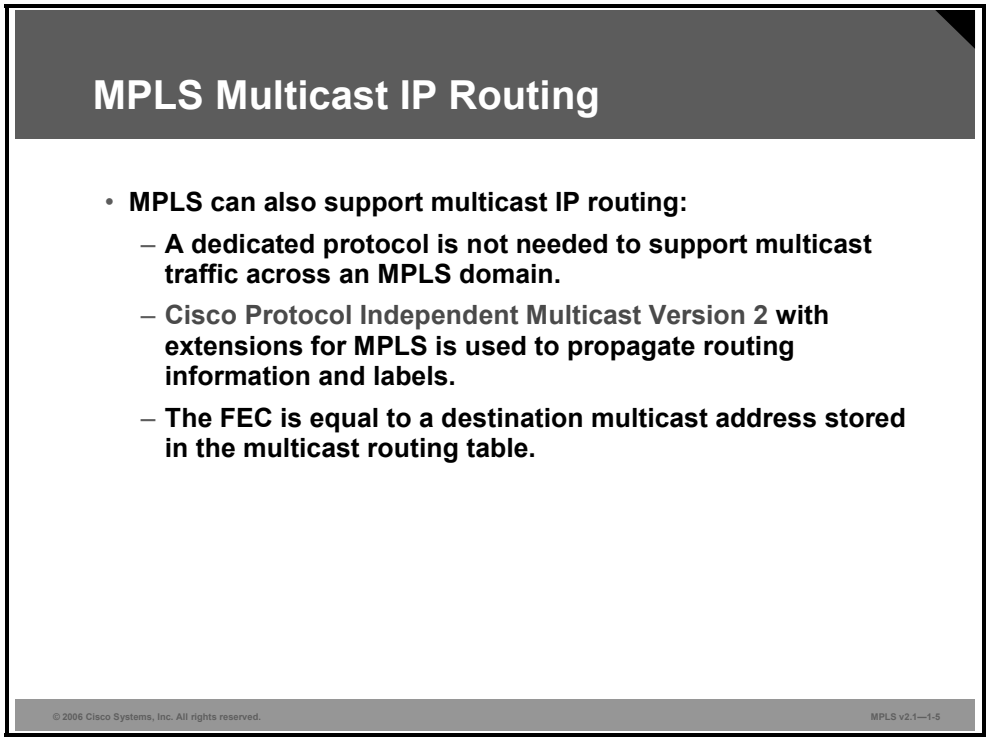

Multicast IP routing can also use MPLS. Cisco Protocol Independent Multicast (PIM) Version 2 with extensions for MPLS is used to propagate routing information and labels.

The FEC is equal to a destination multicast address.

What Are MPLS VPNs?

This topic describes MPLS use in VPNs.

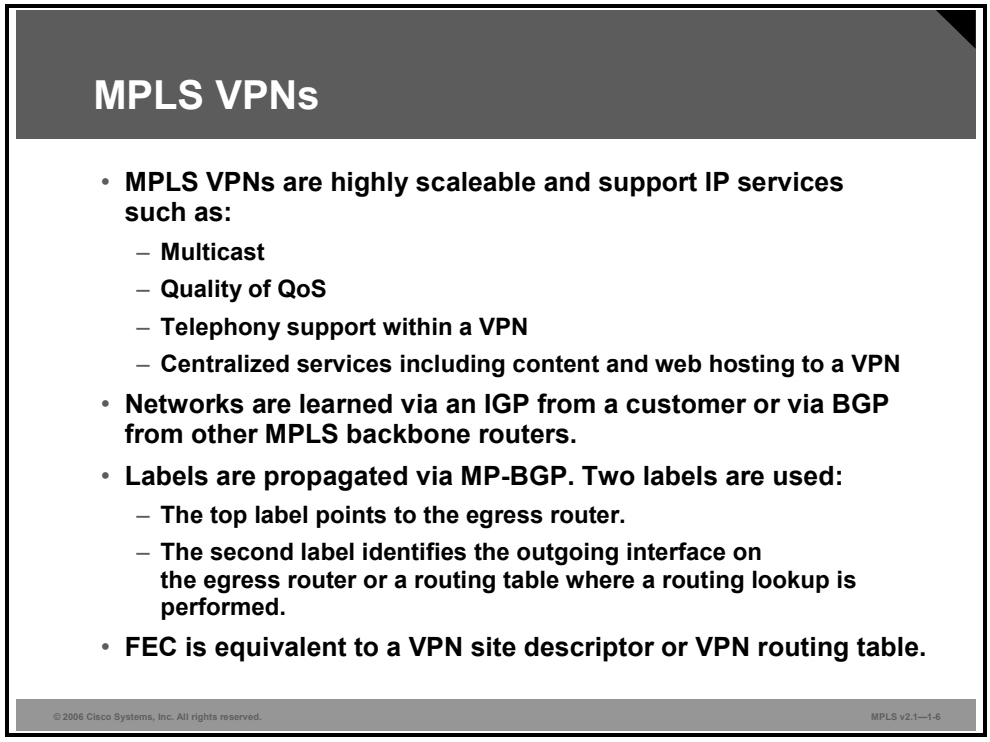

MPLS enables highly scaleable VPN services to be supported. For each MPLS VPN user, the network appears to function as a private IP backbone over which the user can reach other sites within the VPN organization, but not the sites of any other VPN organization. MPLS VPNs are a common application for service providers. Building VPNs in Layer 3 allows delivery of targeted services to a group of users represented by a VPN.

MPLS VPNs are seen as private intranets, and support IP services such as those listed here:

- Multicast
- QoS
- Telephony support within a VPN
- Centralized services including content and web hosting to a VPN

Customer networks are learned via an Interior Gateway Protocol (IGP) (Open Shortest Path First [OSPF], External Border Gateway Protocol [EBGP], Enhanced Interior Gateway Routing Protocol [EIGRP], Routing Information Protocol version 2 [RIPv2], or static) from a customer, or via Border Gateway Protocol (BGP) from other MPLS backbone routers.

MPLS VPNs use two labels:

- The top label points to the egress router.
- The second label identifies the outgoing interface on the egress router or a routing table where a routing lookup is performed.

Label Distribution Protocol (LDP) is needed in the top label to link edge label switch routers (LSRs) with a single label-switched path (LSP) tunnel. Multiprotocol BGP (MP-BGP) is used in the second label to propagate VPN routing information and labels across the MPLS domain. The MPLS VPN FEC is equivalent to a VPN site descriptor or VPN routing table.

What Is MPLS TE?

This topic describes MPLS use in TE environments.

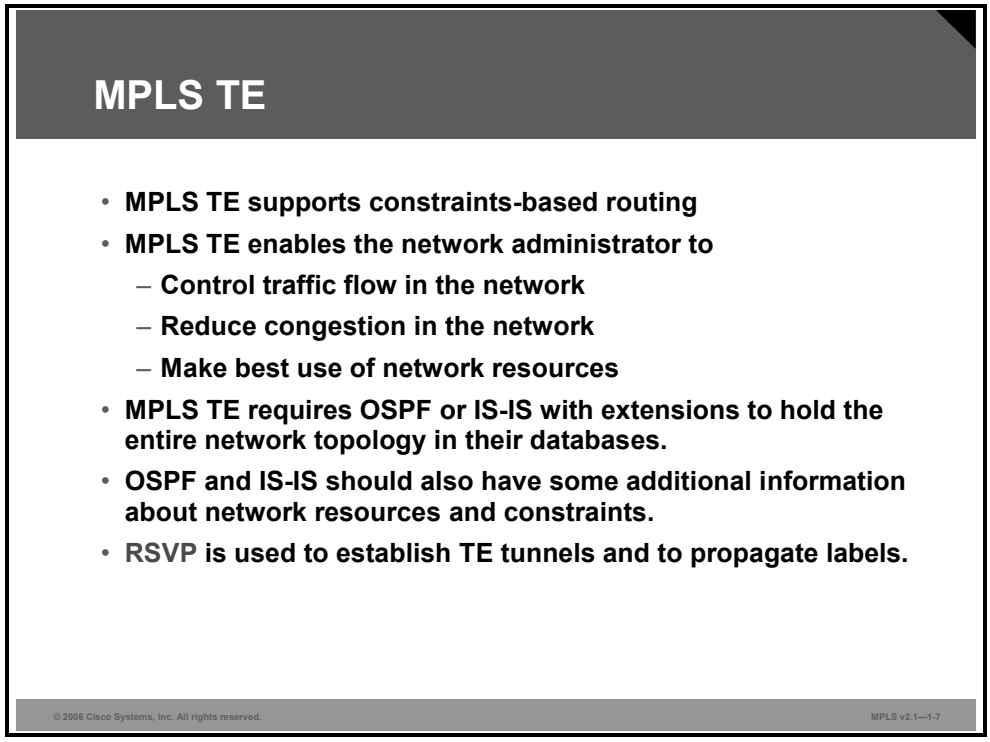

Another application of MPLS is TE. MPLS TE enables an MPLS backbone to replicate and expand upon the TE capabilities of Layer 2 ATM and Frame Relay networks. MPLS TE supports constraint-based routing in which the path for a traffic flow is the shortest path that meets the resource requirements (constraints) of the traffic flow. Factors such as bandwidth requirements, media requirements, and the priority of one traffic flow versus another can be taken into account. TE capabilities enable the administrator of a network to accomplish these goals:

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

MPLS TE has these special requirements:

- Every LSR must see the entire topology of the network (only OSPF and Intermediate System-to-Intermediate System [IS-IS] hold the entire topology).
- Every LSR needs additional information about links in the network. This information includes available resources and constraints. OSPF and IS-IS have extensions to propagate this additional information.
- Resource Reservation Protocol (RSVP) is used to establish TE tunnels and to propagate the labels.

Every edge LSR must be able to create an LSP tunnel on demand. RSVP is used to create an LSP tunnel and to propagate labels for TE tunnels.

What Is MPLS QoS?

This topic describes MPLS use in QoS environments.

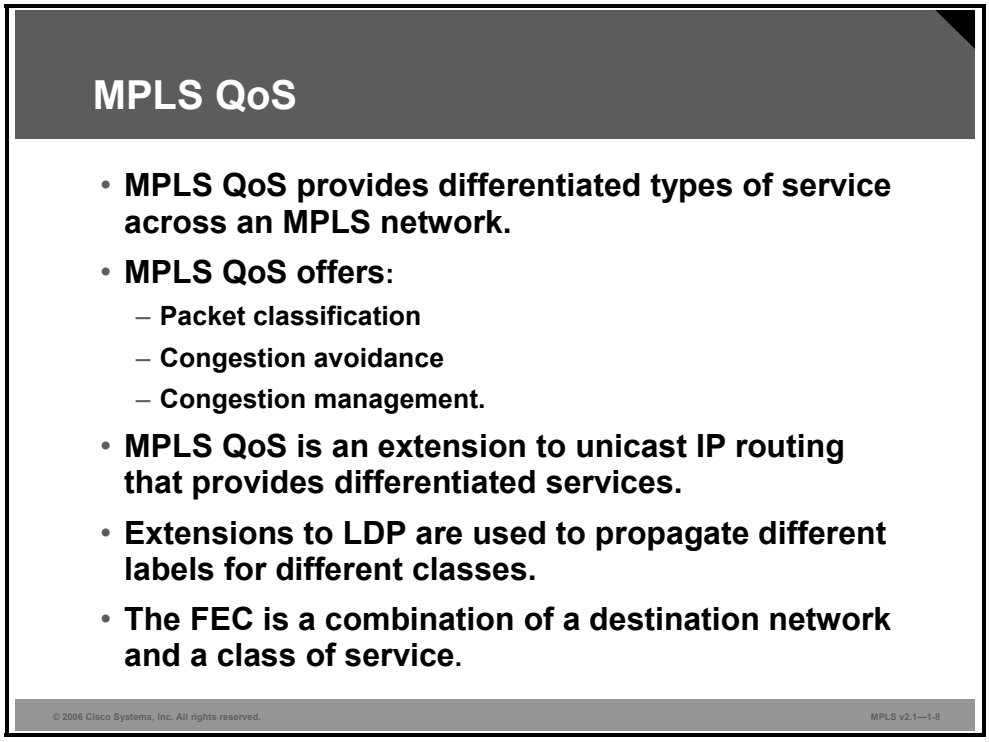

MPLS QoS enables network administrators to provide differentiated types of service across an MPLS network. MPLS QoS offers packet classification, congestion avoidance, and congestion management.

Note	MPLS QoS functions map nearly one-for-one to IP QoS functions on all interface types.

Differentiated QoS is achieved by using MPLS experimental bits or by creating separate LSP tunnels for different classes. Extensions to LDP are used to create multiple LSP tunnels for the same destination (one for each class).

The FEC for MPLS QoS is equal to a combination of a destination network and a class of service (CoS).

What Is AToM?

This topic describes Any Transport over MPLS (AToM).

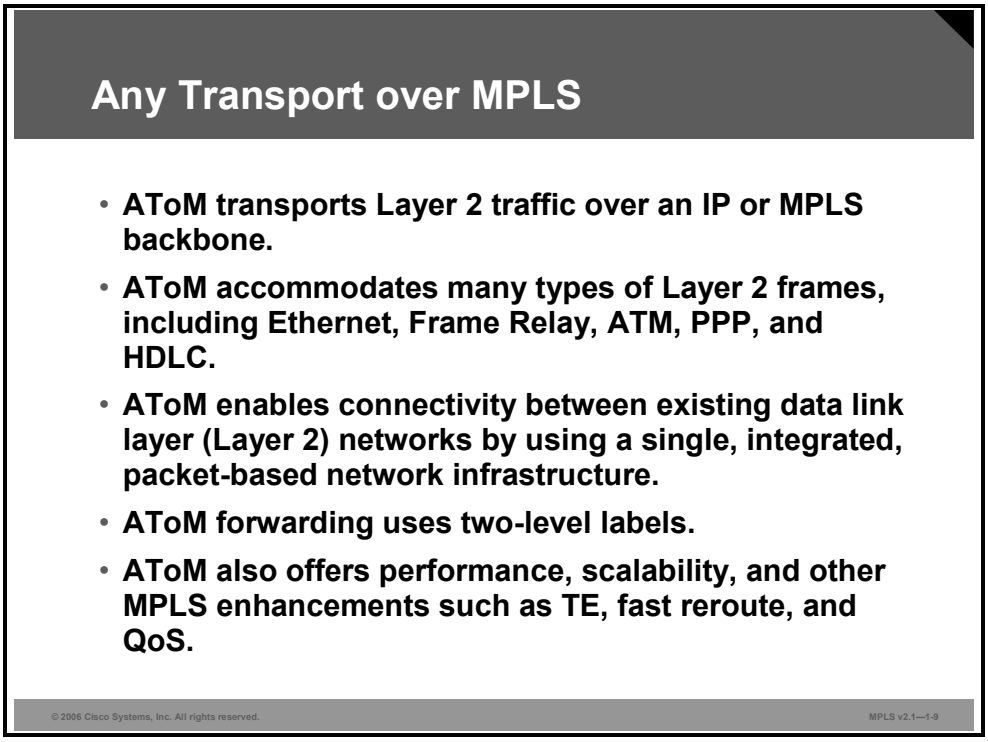

AToM is the Cisco solution for transporting Layer 2 traffic over an IP or MPLS backbone. AToM extends the usability of an IP or MPLS backbone by enabling it to offer both Layer 2 and Layer 3 services. The AToM product set accommodates many types of Layer 2 frames, including Ethernet, Frame Relay, ATM, PPP, and High-Level Data Link Control (HDLC), across various Cisco router platforms including Cisco 7200, 7400, 7500, 7600, 10700, and 12000 Series Routers.

AToM enables service providers to supply connectivity between customer sites with existing data link layer (Layer 2) networks by using a single, integrated, packet-based network infrastructure.

AToM uses a directed LDP session between edge LSRs (or provider edge (PE) routers) for setting up and maintaining connections. Directed LDP is unicast based, and establishes a TCP session across potentially multiple hops. Forwarding occurs through the use of two-level labels, switching between the PE routers. The external label (tunnel label), routes the packet over the MPLS backbone to the egress PE from the ingress PE. The virtual circuit (VC) label determines the egress interface, and it binds the Layer 2 egress interface to the tunnel label.

AToM also offers performance, scalability, and new value-added services using other MPLS enhancements such as TE, fast reroute, and QoS.

Note A detailed discussion of AToM is beyond the scope of this course. A general overview of AToM with links to more detailed information may be found at http://www.cisco.com/en/US/products/ps6646/products_ios_protocol_option_home.html

A technical overview of AToM may be found at http://www.cisco.com/en/US/products/ps6603/products_white_paper09186a00804fbda5.shtml

AToM Examples

This topic provides an overview of some AToM technologies.

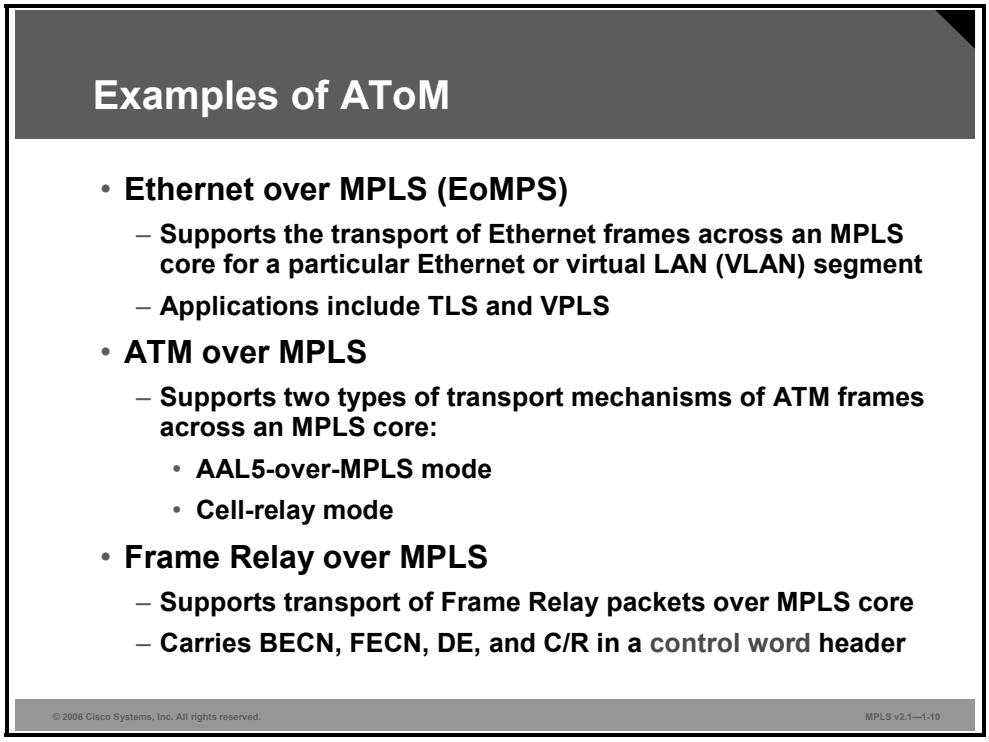

Ethernet over MPLS (EoMPLS) is the transport of Ethernet frames across an MPLS core. It transports all frames received on a particular Ethernet or VLAN segment, regardless of the destination MAC information. It does not perform MAC learning or MAC lookup for forwarding packets from the Ethernet interface. Some applications include Transparent LAN Services (TLS) between facilities, and Virtual Private LAN Services (VPLS), which is a class of VPN that supports the connection of multiple sites in a single bridged domain over a managed IP or MPLS network.

ATM over MPLS is another supported technology. There are two types of transport mechanisms for ATM over MPLS:

- **ATM adaptation layer 5 (AAL5)-over-MPLS mode:** ATM interface assembles the AAL5 protocol data unit (PDU) with either AAL5 Subnetwork Access Protocol (AAL5 SNAP) or AAL5 multiplexer (AAL5 MUX) encapsulation at the boundary and transports it across the network as a single MPLS packet.
- **Cell-relay mode:** The ATM interface receives cells and transports them across the MPLS core. Cell relay with cell packing is used to send multiple cells in one MPLS frame, improving the efficiency of cell transport.

Frame Relay over MPLS (FRoMPLS) is also supported. In this application, traffic is encapsulated in MPLS packets and forwarded across the MPLS network. When encapsulating FRoMPLS, the Frame Relay header and the frame check sequence (FCS) are stripped from the packet. The bits for backward explicit congestion notification (BECN), forward explicit congestion notification (FECN), discard eligibility (DE), and command/response (C/R) are carried across the MPLS network in the control word header.

What Are the Interactions Between MPLS Applications?

This topic identifies the interactions that occur between MPLS applications.

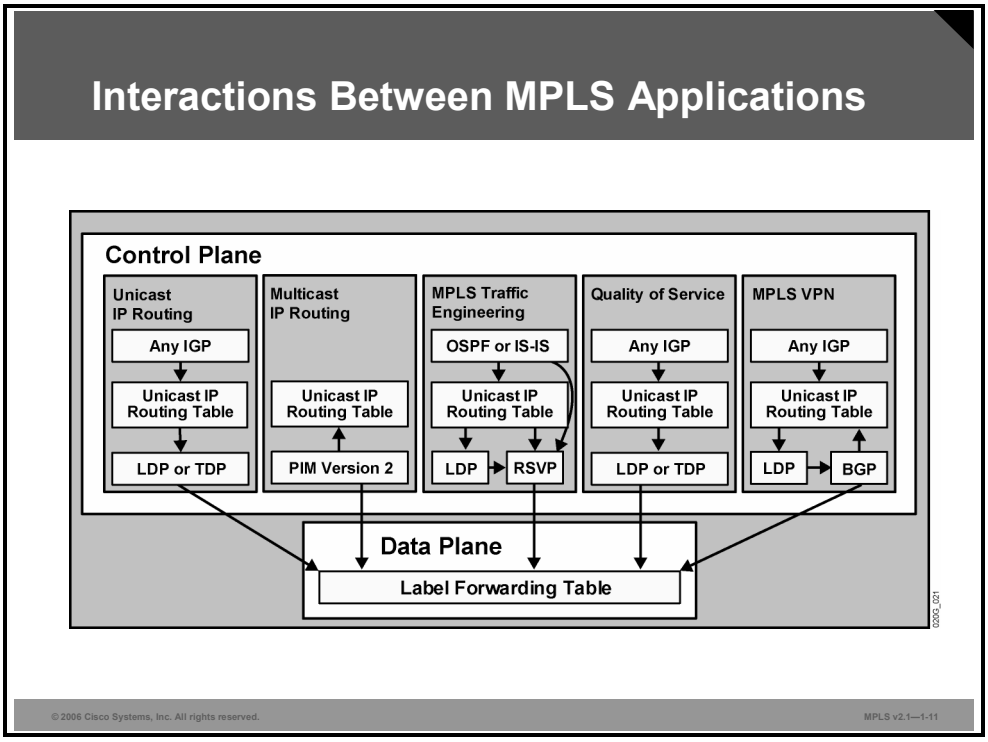

The figure shows the overall architecture when multiple applications are used.

Regardless of the application, the functionality is always split into the control plane and the data (forwarding) plane, as discussed here:

- The applications may use a different routing protocol and a different label exchange protocol in the control plane.
- The applications all use a common label-switching data (forwarding) plane.
- Edge LSR Layer 3 data planes may differ to support label imposition and disposition.
- In general, a label is assigned to an FEC.

Summary

This topic summarizes the key points that were discussed in this lesson.

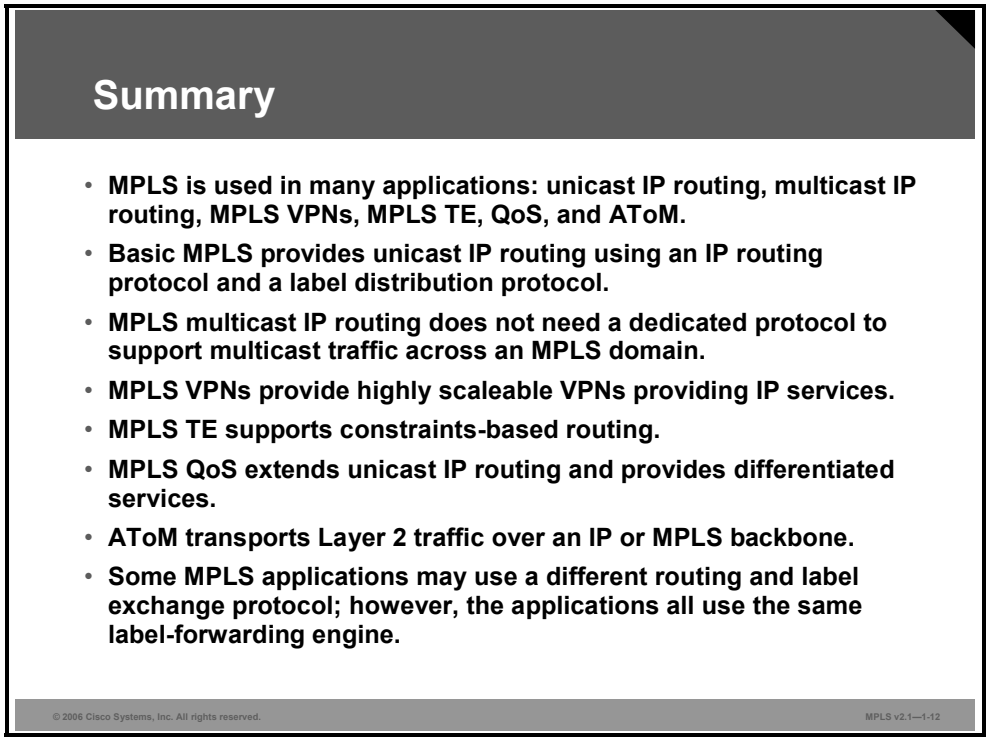

Module Summary

This topic summarizes the key points that were discussed in this module.

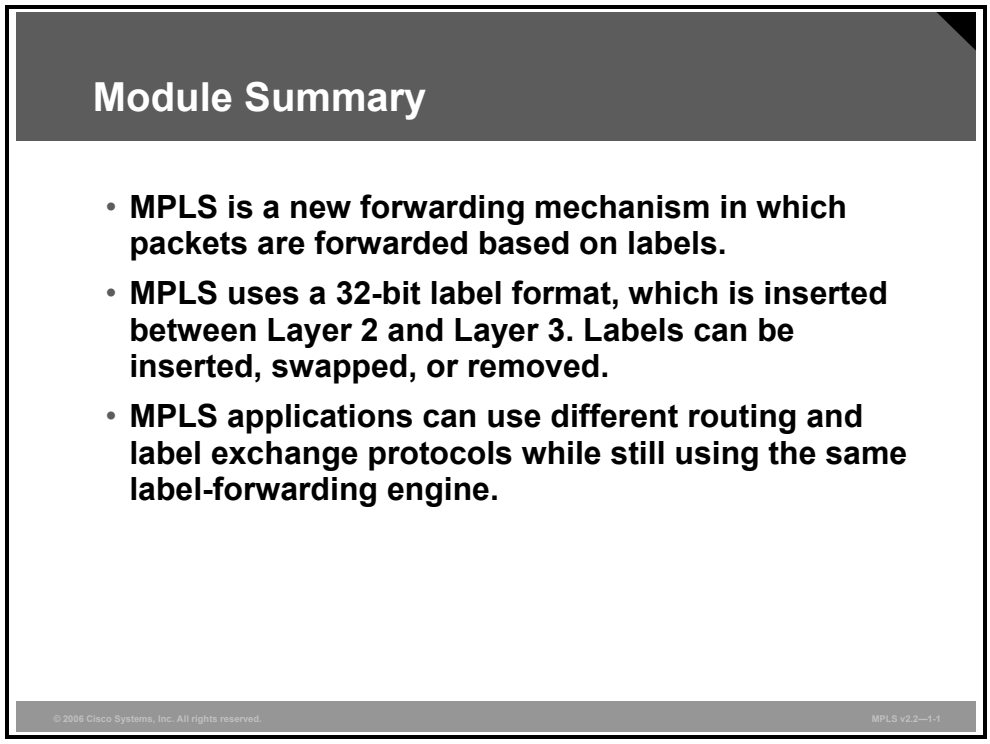

Multiprotocol Label Switching (MPLS) forwards packets based on labels. MPLS can be implemented in ATM networks to provide optimal routing across Layer 2 ATM switches. MPLS uses the concept of a label stack where multiple labels are supported in one packet. You can use MPLS in many applications. When many MPLS applications are being used, all applications use a single label-forwarding engine.

References

For additional information, refer to these resources:

- RFC 3031, *Multiprotocol Label Switching Architecture* *http://www.ietf.org/rfc/rfc3031.txt*
- RFC 3032, *MPLS Label Stack Encoding* *http://www.rfc-editor.org/rfc/rfc3032.tx.*

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) What are three foundations of traditional IP routing? (Choose three.) (Source: Introducing Basic MPLS Concepts)

A) Routing protocols are used on all devices to distribute routing information.
B) Regardless of protocol, routers always forward packets based on the IP destination address only (except for using PBR).
C) Routing lookups are performed on every router.
D) Routing is performed by assigning a label to an IP destination.

Q2) Which three statements are true? (Choose three.) (Source: Introducing Basic MPLS Concepts)

A) MPLS uses labels to forward packets.
B) MPLS works only in IP networks.
C) MPLS labels can correspond to a Layer 3 destination address, QoS, source address, or Layer 2 circuit.
D) MPLS does not require a routing table lookup on core routers.

Q3) In MPLS TE, which two statements are true? (Choose two.) (Source: Introducing Basic MPLS Concepts)

A) Traditional IP routing does not support traffic engineering.
B) Traditional IP routing would force all traffic to use the same path based on destination.
C) Using MPLS TE, traffic can be forwarded based on parameters such as QoS and source address.
D) MPLS does not support traffic engineering.

Q4) The label distribution protocol (either LDP or TDP) is the responsibility of the _____. (Source: Introducing Basic MPLS Concepts)

A) data plane
B) forwarding plane
C) system plane
D) control plane

Q5) The MPLS label field consists of how many bits? (Source: Introducing Basic MPLS Concepts)

A) 64 bits
B) 32 bits
C) 16 bits
D) 8 bits

Q6) Which two statements are true? (Choose two.) (Source: Introducing Basic MPLS Concepts)

A) An edge LSR is a device that inserts labels on packets or removes labels, and forwards packets based on labels.
B) An LSR is a device that primarily labels packets or removes labels.
C) An LSR is a device that forwards packets based on labels.
D) An end LSR is a device that primarily inserts labels on packets or removes labels.

Q7) MPLS labels can correspond to which type of addresses? (Source: Introducing Basic MPLS Concepts)

A) Layer 2 source addresses
B) Layer 3 source addresses
C) Layer 2 destination addresses
D) Layer 3 destination addresses

Q8) Which term is best described as "a simple label-based forwarding engine"? (Source: Introducing Basic MPLS Concepts)

A) control plane
B) ground plane
C) data plane
D) routing plane

Q9) Which three statements are true? (Choose three.) (Source: Introducing MPLS Labels and Label Stacks)

A) In frame-mode MPLS, labels are typically inserted between the Layer 2 header and the Layer 3 header.
B) MPLS labels are inserted after the Layer 3 header in frame-mode MPLS.
C) In cell-mode MPLS, MPLS uses the VPI/VCI fields as the label.
D) MPLS will not work in ATM networks.
E) MPLS labels are 32 bits.
F) MPLS labels are 64 bits.

Q10) How long is the actual MPLS label contained in the MPLS label field? (Source: Introducing MPLS Labels and Label Stacks)

A) 32 bits long
B) 8 bits long
C) 16 bits long
D) 20 bits long

Q11) Which two statements are true? (Choose two.) (Source: Introducing MPLS Labels and Label Stacks)

A) Usually one label is assigned to an IP packet.
B) Usually two labels are assigned to an IP packet.
C) Two labels will be assigned to an MPLS VPN packet.
D) One label will be assigned to an MPLS VPN packet.

Q12) What are two normal functions of an edge LSR? (Choose two.) (Source: Introducing MPLS Labels and Label Stacks)

A) impose labels at the ingress router
B) impose labels at the egress router
C) pop labels at the ingress router
D) pop labels at the egress router

Q13) Cisco routers automatically assign the IP precedence value to which field in the MPLS label? (Source: Introducing MPLS Labels and Label Stacks)

A) TTL field
B) experimental field
C) top-of-stack field
D) The IP precedence value is not copied to the MPLS field; this value remains in the IP packet.

Q14) What is NOT a valid Ethertype used to identify Layer 3 protocols with most Layer 2 encapsulations? (Source: Introducing MPLS Labels and Label Stacks)

A) unlabeled IP unicast (PID = 0x0800)
B) labeled IP unicast (PID = 0x0847)
C) unlabeled IP multicast (PID = 0x8846)
D) labeled IP multicast (PID = 0x8848)

Q15) Which two statements are true regarding RSVP? (Choose two.) (Source: Identifying MPLS Applications)

A) RSVP is used to create an LSP tunnel.
B) RSVP propagates labels for TE tunnels.
C) RSVP assigns labels for TE tunnels.
D) RSVP is not used to create an LSP tunnel.

Q16) When MPLS is used for QoS, which statement is true? (Source: Identifying MPLS Applications)

A) QoS is achieved by using the protocol bits in the MPLS label field.
B) QoS is achieved by using the TTL bits in the MPLS label field.
C) QoS is achieved by using the experimental bits in the MPLS label field.
D) At this time, QoS is not supported by MPLS.

Q17) In MPLS VPN networks, which statement is true? (Source: Identifying MPLS Applications)

A) Labels are propagated via LDP or TDP.
B) Next-hop addresses instead of labels are used in an MPLS VPN network.
C) Labels are propagated via MP-BGP.
D) Two labels are used; the top label identifies the VPN, and the bottom label identifies the egress router.

Q18) Which two statements are true regarding interactions between MPLS applications? (Choose two.) (Source: Identifying MPLS Applications)

A) The forwarding plane is the same for all applications.
B) Differences exist in the forwarding plane depending on the MPLS application.
C) The control plane is the same for all applications.
D) Differences exist in the control plane depending on the MPLS application.

Q19) In MPLS VPNs, what does the FEC refer to? (Source: Identifying MPLS Applications)

A) IP destination network
B) MPLS ingress router
C) core of the MPLS network
D) VPN destination network

Module Self-Check Answer Key

Q1) A, B, C

Q2) A, C, D

Q3) B, C

Q4) D

Q5) B

Q6) A, C

Q7) D

Q8) C

Q9) A, C, E

Q10) D

Q11) A, C

Q12) A, D

Q13) B

Q14) C

Q15) A, B

Q16) C

Q17) C

Q18) A, D

Q19) D

Module 2

Label Assignment and Distribution

Overview

This module describes the assignment and distribution of labels in a Multiprotocol Label Switching (MPLS) network, including neighbor discovery and session establishment procedures. Label distribution, control, and retention modes will also be covered. This module also covers the functions and benefits of penultimate hop popping (PHP).

Module Objectives

Upon completing this module, you will be able to describe how MPLS labels are assigned and distributed. This ability includes being able to meet these objectives:

- Describe how LDP neighbors are discovered
- Describe how the LIB, FIB, and LFIB tables are populated with label information
- Describe how convergence occurs in a frame-mode MPLS network
- Describe MPLS label allocation, distribution, and retention modes

Lesson 1

Discovering LDP Neighbors

Overview

This lesson takes a detailed look at the Label Distribution Protocol (LDP) neighbor discovery process via hello messages and the type of information that is exchanged. The lesson also describes the events that occur during the negotiation phase of LDP session establishment and concludes with the nonadjacent neighbor discovery process.

This lesson provides an understanding of how an LDP neighbor is discovered and what type of information is sent back and forth between two neighbors. The lesson also discusses situations in which the neighbor is not directly connected to a peer. This information will provide a further understanding of the Multiprotocol Label Switching (MPLS) technology.

Objectives

Upon completing this lesson, you will be able to describe how LDP neighbors are discovered. This ability includes being able to meet these objectives:

- Describe how LDP sessions are established between adjacent neighbors
- Describe the contents of an LDP hello message
- Describe negotiating label space as it applies to LDP session establishment
- Describe how LDP neighbors are discovered
- Describe the process of LDP session negotiation between LDP neighbors
- Describe how LDP sessions are established between nonadjacent neighbors

Establishing an Adjacent LDP Session

This topic describes how LDP sessions are established between neighbors.

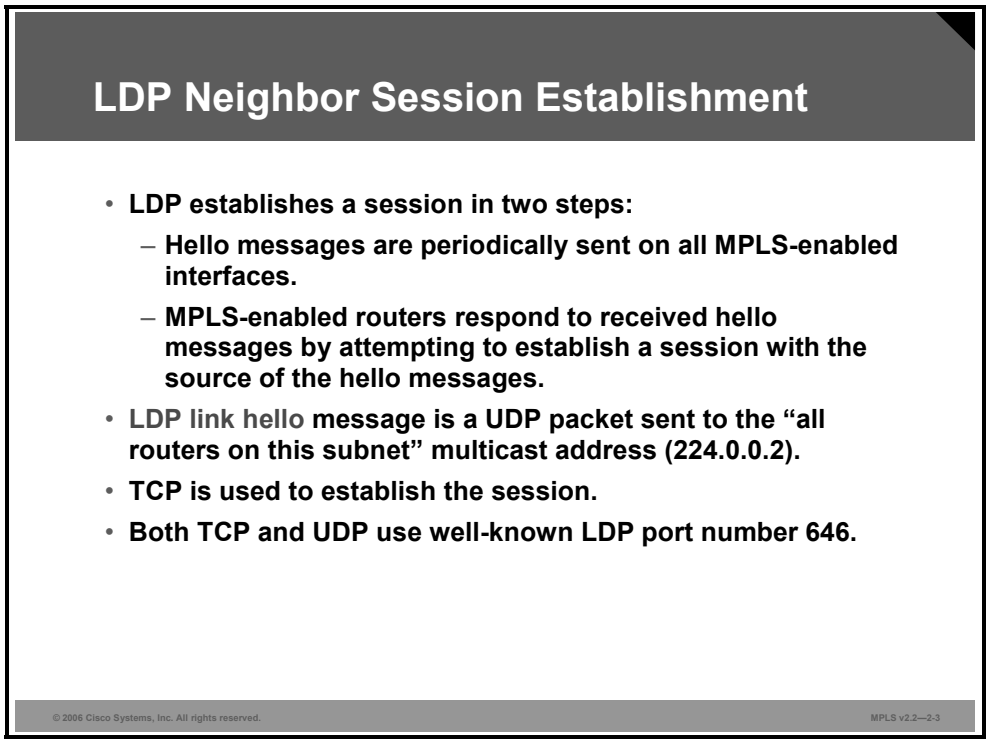

LDP is a standard protocol used to exchange labels between adjacent routers.

Note	Tag Distribution Protocol (TDP) is an older Cisco proprietary protocol that has the same functionality as LDP. Although the remainder of this lesson will focus on LDP, it should be noted that TDP, as the predecessor of LDP, works in a similar fashion.

LDP periodically sends hello messages (every 5 seconds). If the label switch router (LSR) is adjacent or one hop from its neighbor, the LSR sends out LDP link hello messages to all the routers on the subnet as User Datagram Protocol (UDP) packets with a multicast destination address of 224.0.0.2 ("all routers on a subnet") and destination port number of 646. (TDP uses destination port 711.)

A neighboring LSR enabled for LDP will respond by opening a TCP session with the same destination port number 646, and the two routers begin to establish an LDP session through unicast TCP.

What Are LDP Hello Messages?

This topic describes the contents of an LDP link hello message.

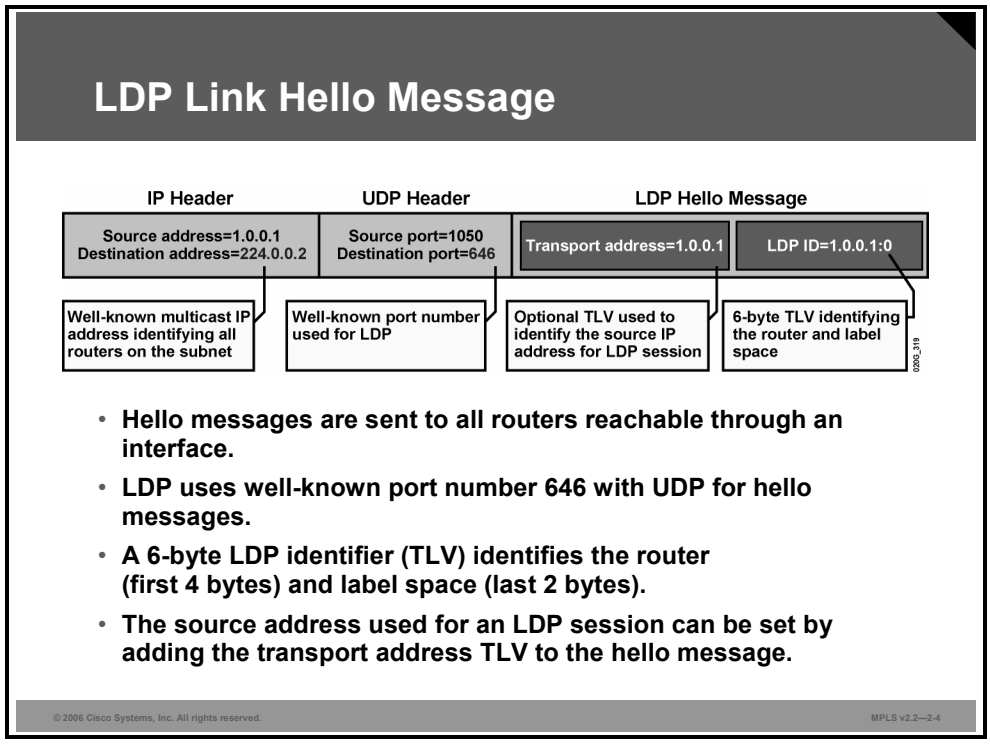

The contents of an LDP link hello message are as follows:

- Destination IP address (224.0.0.2), which reaches all routers on the subnetwork
- Destination port, which equals the LDP well-known port number 646
- The actual hello message, which may optionally contain a transport address type, length, value (TLV) to instruct the peer to open the TCP session to the transport address instead of the source address found in the IP header

The LDP identifier (LDP ID) is used to uniquely identify the neighbor and the label space.

Note	Label space defines the way MPLS assigns labels to destinations. Label space can either be per-platform or per-interface.

Multiple LDP sessions can be established between a pair of LSRs if they use multiple label spaces.

Example: Per-Platform Label Space

This example illustrates per-platform label space.

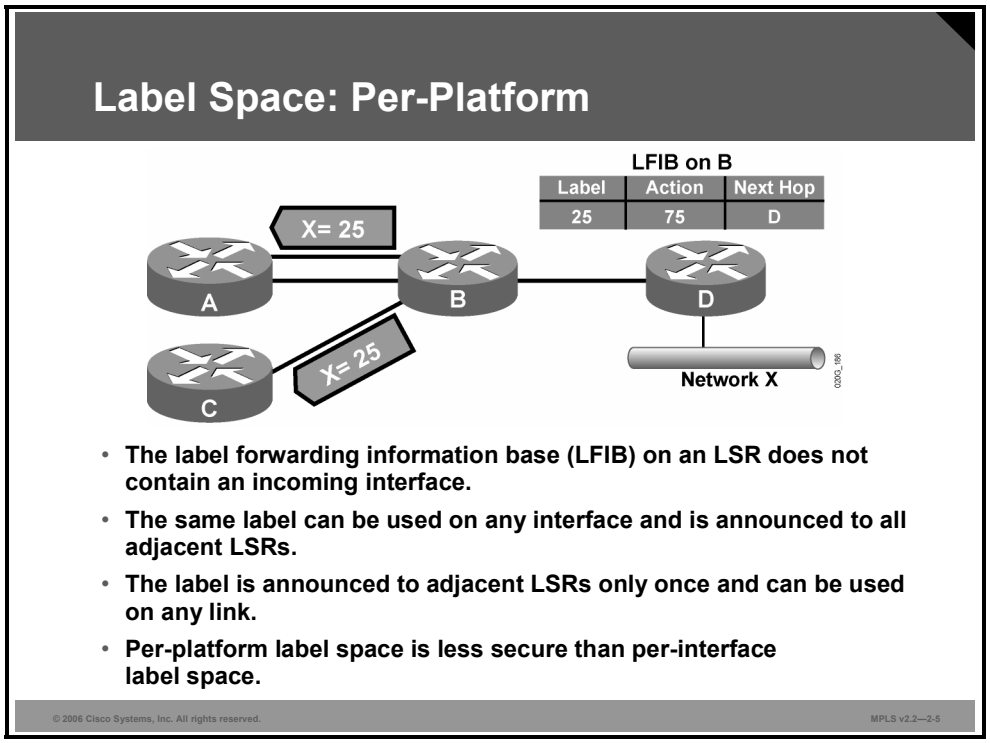

Per-platform label space is used with frame-mode MPLS, where one label is assigned to a destination network and sent to all LDP peers. This label can then be used on any incoming interface. The per-platform label space minimizes the number of LDP sessions and allows upstream label-switched path (LSP) tunnels to span parallel links, because the same label is used on all of those links. However, per-platform label space is less secure than per-interface label space, because untrusted routers could use labels that were never allocated to them.

Negotiating Label Space

This topic describes negotiating label space as it applies to LDP session establishment.

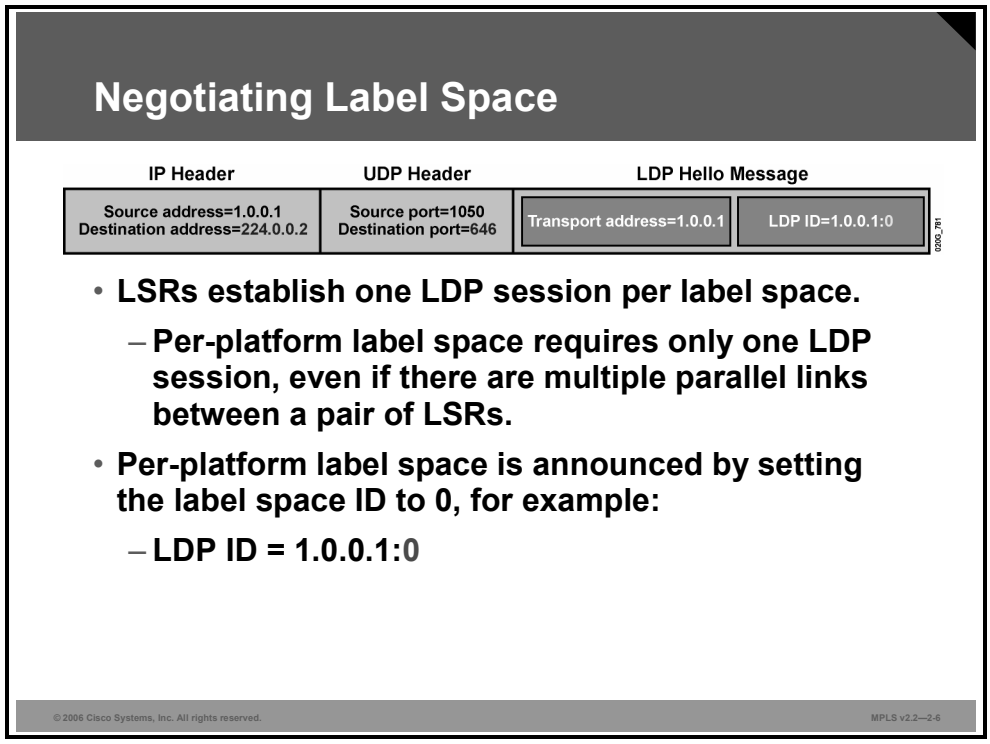

If a pair of routers is connected over two or more parallel links and uses frame-mode MPLS, the routers try to establish multiple sessions by using the same LDP ID. Because the routers are using per-platform label space, this action will result in only one session remaining; the other session will be broken.

Per-platform label space is identified by setting the label space ID to 0 in the LDP ID field.

For all frame-mode interfaces, only one LDP session between a pair of LSRs is used because frame-mode MPLS uses per-platform label space.

Discovering LDP Neighbors

This topic describes how LDP neighbors are discovered.

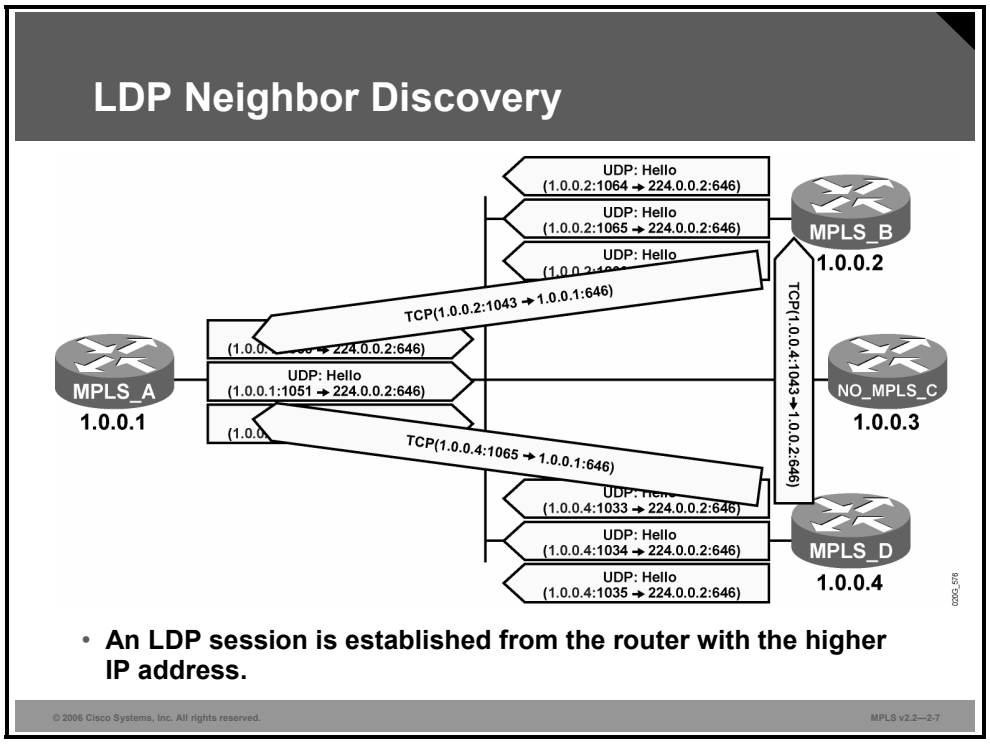

Example: LDP Neighbor Discovery

In the figure, three out of four routers periodically send out LDP hello messages (the fourth router is not MPLS-enabled).

Routers that have the higher IP addresses must initiate the TCP session.

Note	The highest IP address of all loopback interfaces on a router is used. If no loopback interfaces are configured on the router, the highest IP address of a configured interface that was operational at LDP startup is used.

After the TCP session is established, routers will keep sending LDP hello messages to potentially discover new peers or to identify failures.

Negotiating LDP Sessions

This topic describes the process of LDP neighbor session negotiation between LDP neighbors.

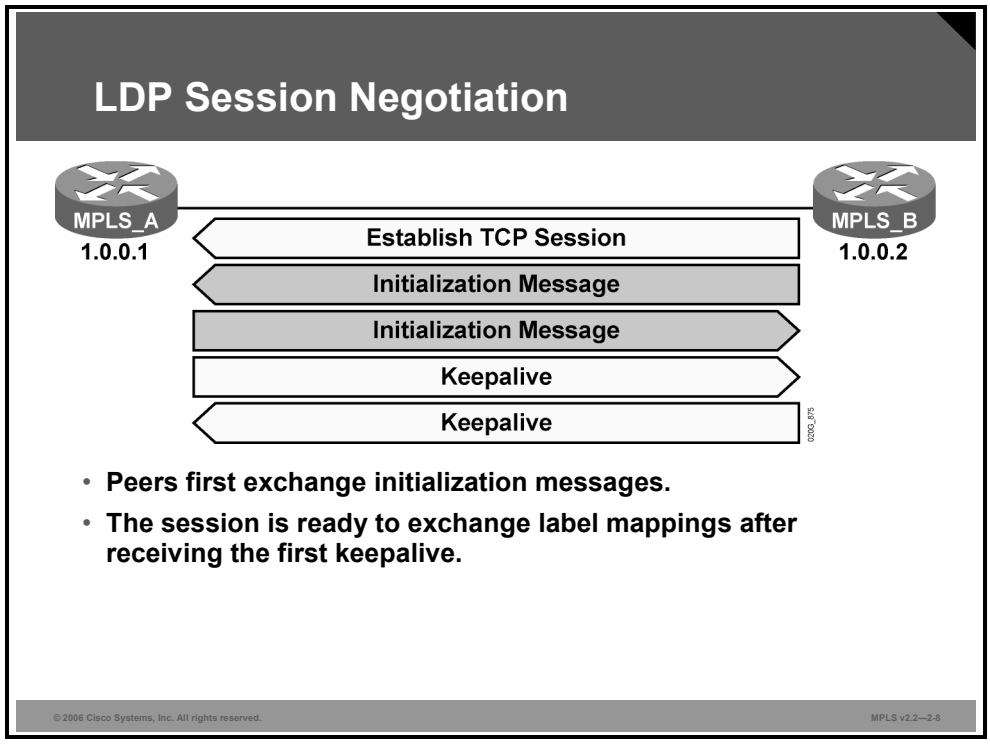

LDP session negotiation is a three-step process, as follows:

Step 1 Establish the TCP session.

Step 2 Exchange initialization messages.

Step 3 Exchange initial keepalive messages.

Note LDP keepalives are sent every 60 seconds.

After these steps have occurred, the two peers will start exchanging labels for networks that they have in their main routing tables.

Discovering Nonadjacent Neighbors

This topic describes how LDP discovers nonadjacent neighbors.

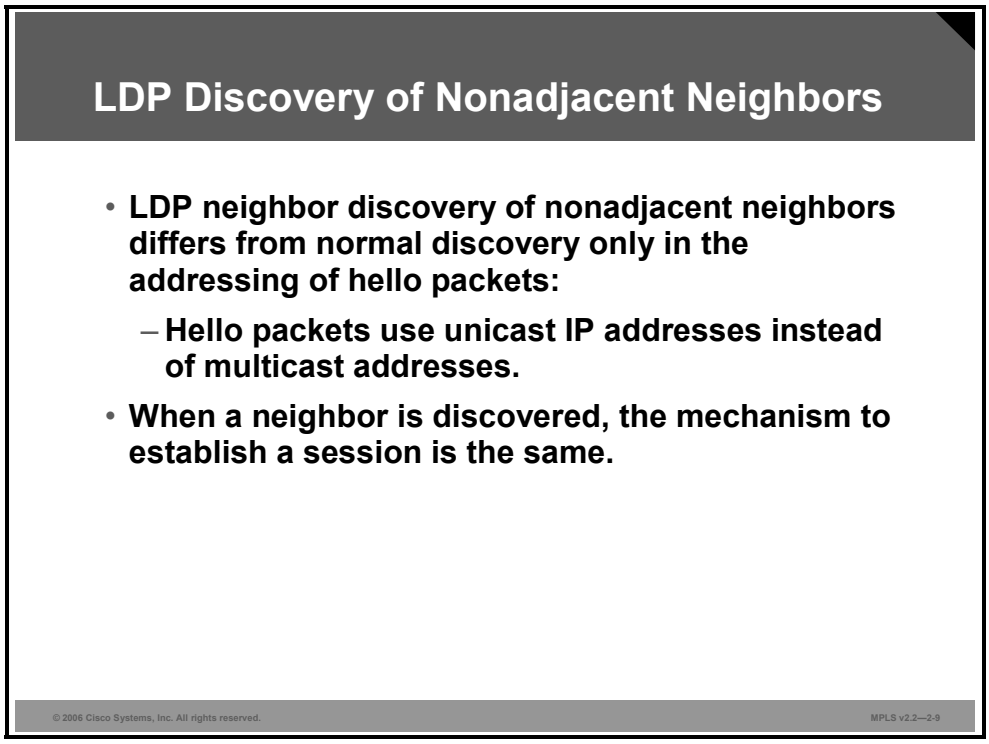

If the LSR is more than one hop from its neighbor, it is not directly connected or adjacent to its neighbor. The LSR can be configured with the **mpls ldp neighbor** [**vrf** *vrf-name*] *ip-address* **targeted** command to send a directed hello message as a unicast UDP packet specifically addressed to the nonadjacent neighbor LSR. The directed hello message is called an LDP targeted hello.

The rest of the session negotiation is the same as for adjacent routers. The nondirectly connected LSR will respond to the hello message by opening a unicast TCP session with the same destination port number 646, and the two routers begin to establish an LDP session. (If the path between two LSRs has been traffic engineered and has LDP enabled, the LDP session between them is called a targeted session.)

Example: Applications Using Targeted LDP Sessions

This figure lists some applications that use targeted LDP sessions.

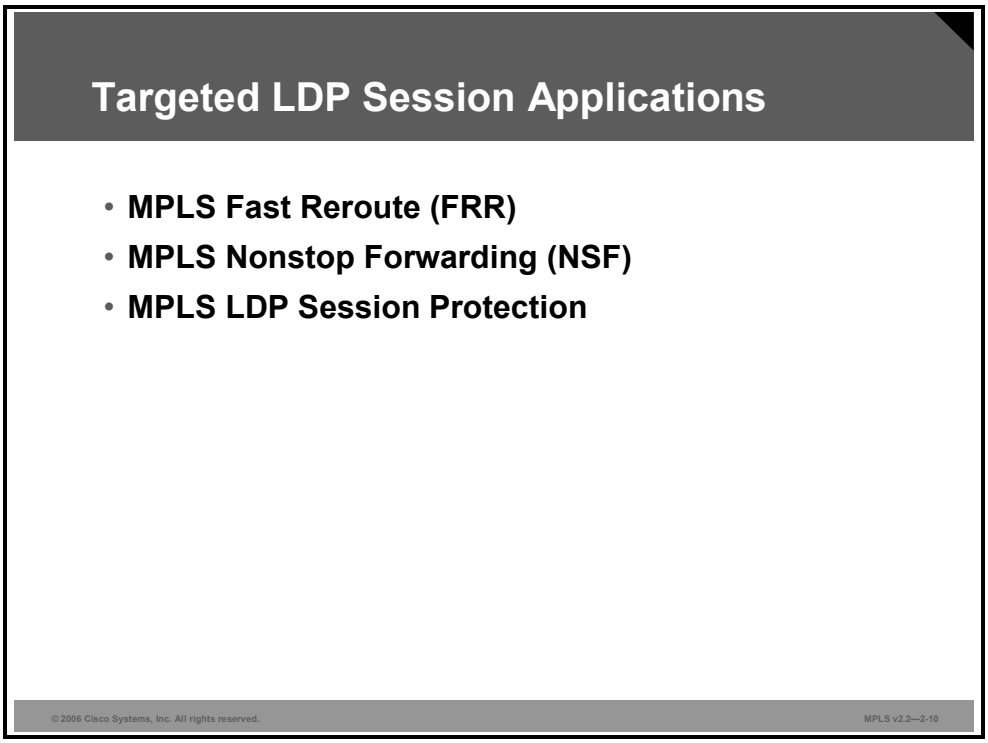

Here are some applications that use targeted LDP sessions:

- MPLS Fast Reroute (FRR), which is the ability to locally patch traffic onto a backup tunnel in case of a link or node failure with a failover time of 50 ms or lower
- MPLS Nonstop Forwarding (NSF), which allows a router to recover from disruption in control plane service (specifically, the LDP component) without losing its MPLS forwarding state
- MPLS LDP Session Protection, which provides faster LDP convergence when a link recovers following an outage by maintaining LDP bindings for a period of time

Summary

This topic summarizes the key points that were discussed in this lesson.

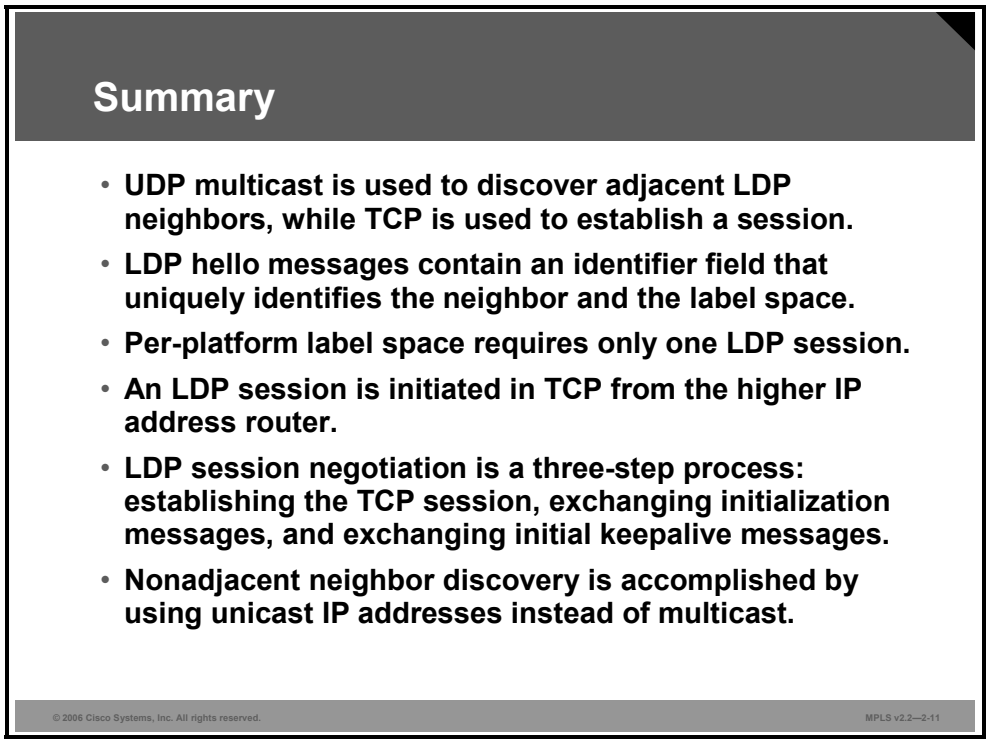

Lesson 2

Introducing Typical Label Distribution in Frame-Mode MPLS

Overview

This lesson discusses how label allocation and distribution function in a frame-mode network. Also covered are penultimate hop popping (PHP) and how the Multiprotocol Label Switching (MPLS) data structures are built. This lesson is essential to understanding the basic fundamentals of how information gets distributed and placed into the appropriate tables for both labeled and unlabeled packet usage.

Objectives

Upon completing this lesson, you will be able to describe how the Label Information Base (LIB), Forwarding Information Base (FIB), and label forwarding information base (LFIB) tables are populated with label information. This ability includes being able to meet these objectives:

- Describe how labels are propagated across a network
- Describe the function of LSPs
- Describe the function of PHP
- Describe the impact that IP aggregation has on LSPs
- Describe how labels are allocated in a frame-mode MPLS network
- Describe how MPLS labels are distributed and advertised in a frame-mode network
- Describe how the LFIB table is populated in an MPLS network
- Describe how IP packets cross an MPLS network
- Describe how frame-mode loops are detected
- Describe the approaches for assigning labels to networks

Propagating Labels Across a Network

This topic describes how labels are propagated across a network.

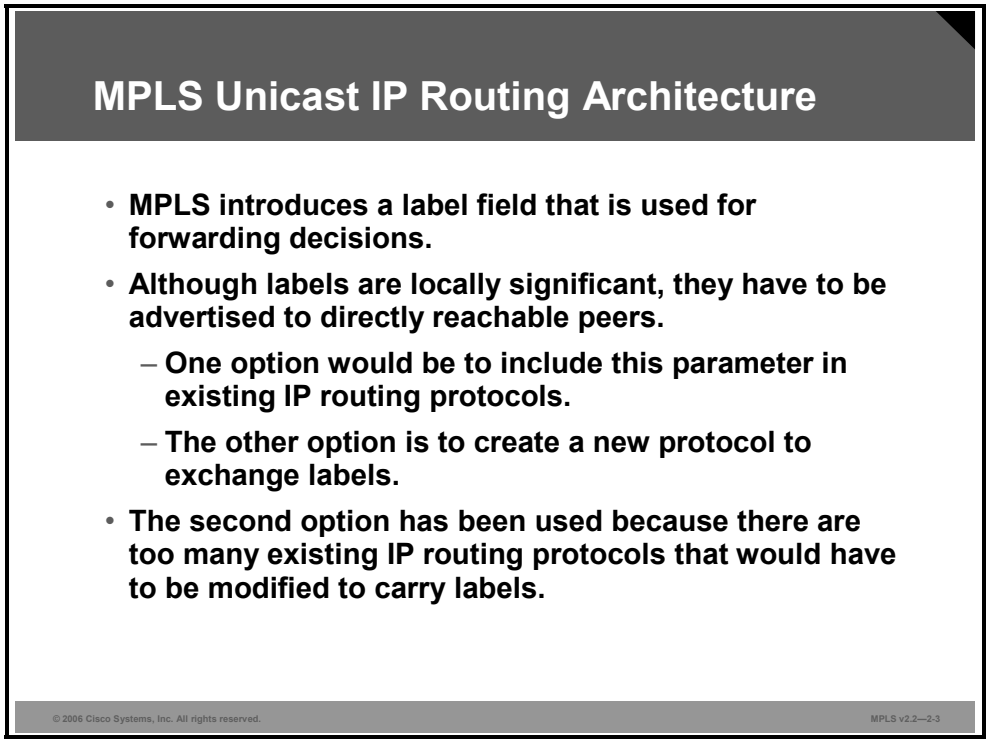

One application of MPLS is unicast IP routing. A label is assigned to destination IP networks and is later used to label packets sent toward those destinations.

Note	In MPLS terminology, a forwarding equivalence class (FEC) in MPLS unicast IP routing equals an IP destination network.

Standard or vendor-specific routing protocols are used to advertise IP routing information. MPLS adds a new piece of information that must be exchanged between adjacent routers.

Here are the two possible approaches to propagating this additional information (labels) between adjacent routers:

- Extend the functionality of existing routing protocols
- Create a new protocol dedicated to exchanging labels

The first approach requires much more time and effort because of the large number of different routing protocols: Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Enhanced Interior Gateway Routing Protocol (EIGRP), Interior Gateway Routing Protocol (IGRP), Routing Information Protocol (RIP), and so on. The first approach also causes interoperability problems between routers that support this new functionality and those that do not. Therefore, the Internet Engineering Task Force (IETF) selected the second approach.

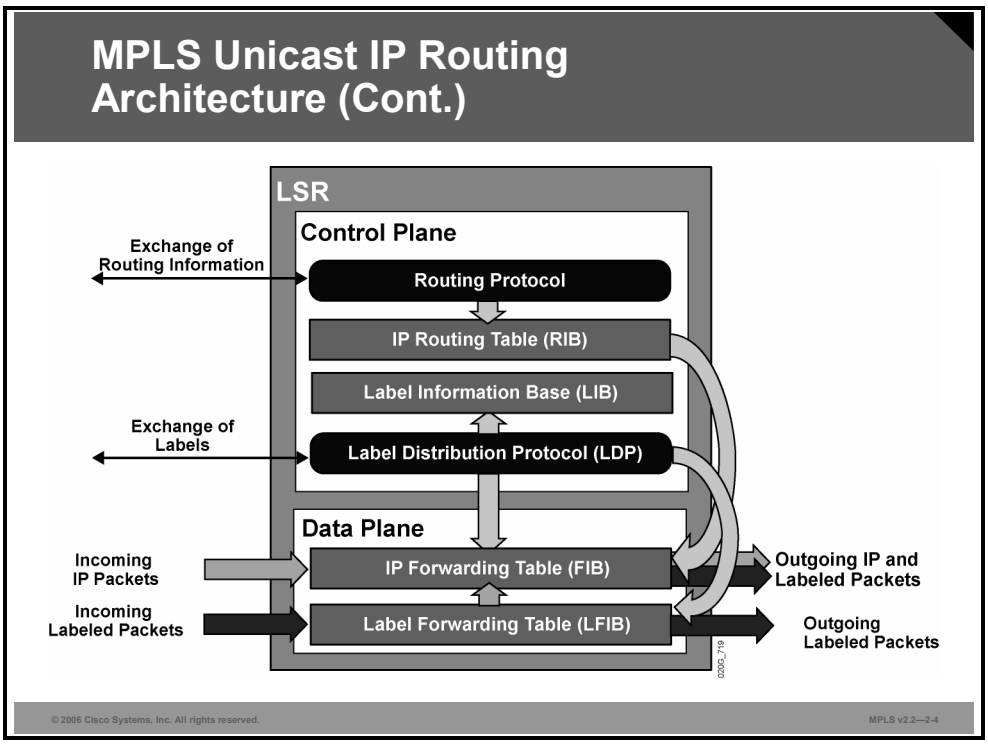

Example: Building Blocks for IP Forwarding

The figure shows the building blocks used by routers to perform traditional IP forwarding.

The control plane consists of a routing protocol that exchanges routing information and maintains the contents of the main routing table. When combined with Cisco Express Forwarding (CEF), the IP forwarding table in the data plane forwards the packets through the router.

The Label Distribution Protocol (LDP) in the control plane exchanges labels and stores them in the Label Information Base (LIB). This information is then used in the data plane to provide MPLS functionality, as follows:

- A label is added to the IP forwarding table (FIB) to map an IP prefix to a next-hop label.
- A locally generated label is added to the LFIB and mapped to a next-hop label.

These forwarding scenarios are possible when MPLS is enabled on a router:

- An incoming IP packet is forwarded by using the FIB table and sent out as an IP packet (the usual CEF switching).
- An incoming IP packet is forwarded by using the FIB table and sent out as a labeled IP packet (the default action if there is a label assigned to the destination IP network).
- An incoming labeled packet is forwarded by using the LFIB table and sent out as a labeled IP packet.
- An incoming labeled packet has its label removed, is inspected against the FIB table, and is forwarded as an IP packet.

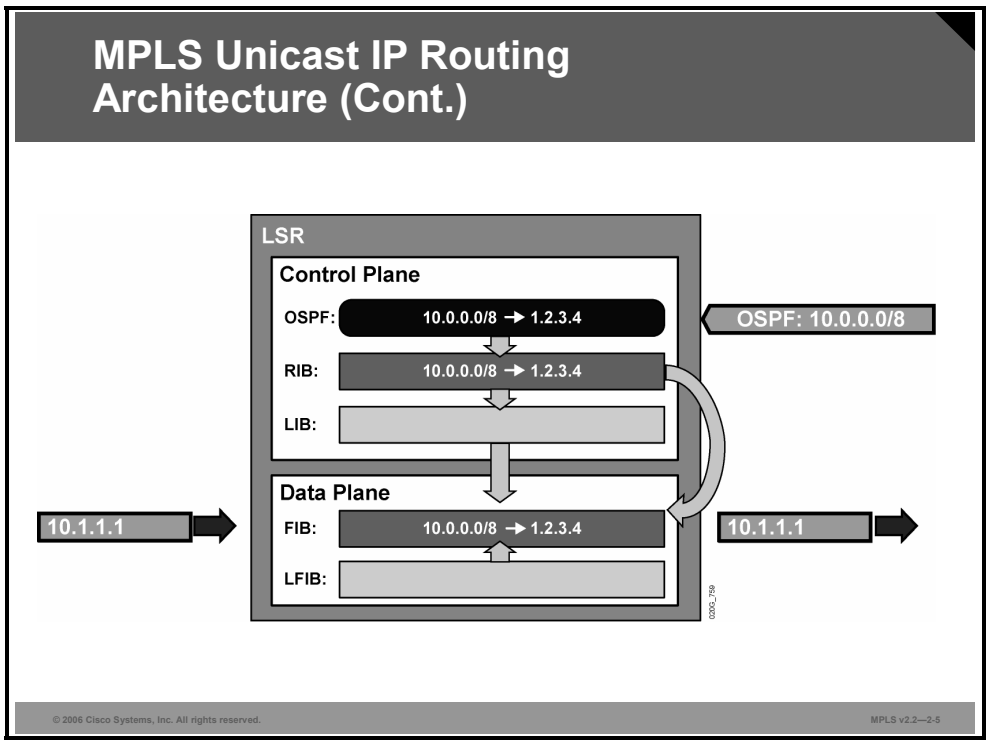

Example: Using the FIB Table to Forward Packets

The figure shows a scenario in which IP packets are successfully forwarded by using the FIB table.

Note	CEF is the only Layer 3 switching mechanism that uses the FIB table. CEF must be enabled on all routers running MPLS, and on all ingress interfaces receiving unlabeled IP packets that are to be propagated as labeled packets.

Labeled packets, on the other hand, are not forwarded because of a lack of information in the LFIB table. Normal MPLS functionality prevents the forwarding from happening, because no adjacent router is going to use a label unless this router previously advertised the label.

The example illustrates that label switching tries to use the LFIB table only if the incoming packet is labeled, even if the destination address is reachable by using the FIB table.

Note	The LIB table will hold all locally generated labels by a label switch router (LSR). The LFIB table contains labels that are used to switch packets.

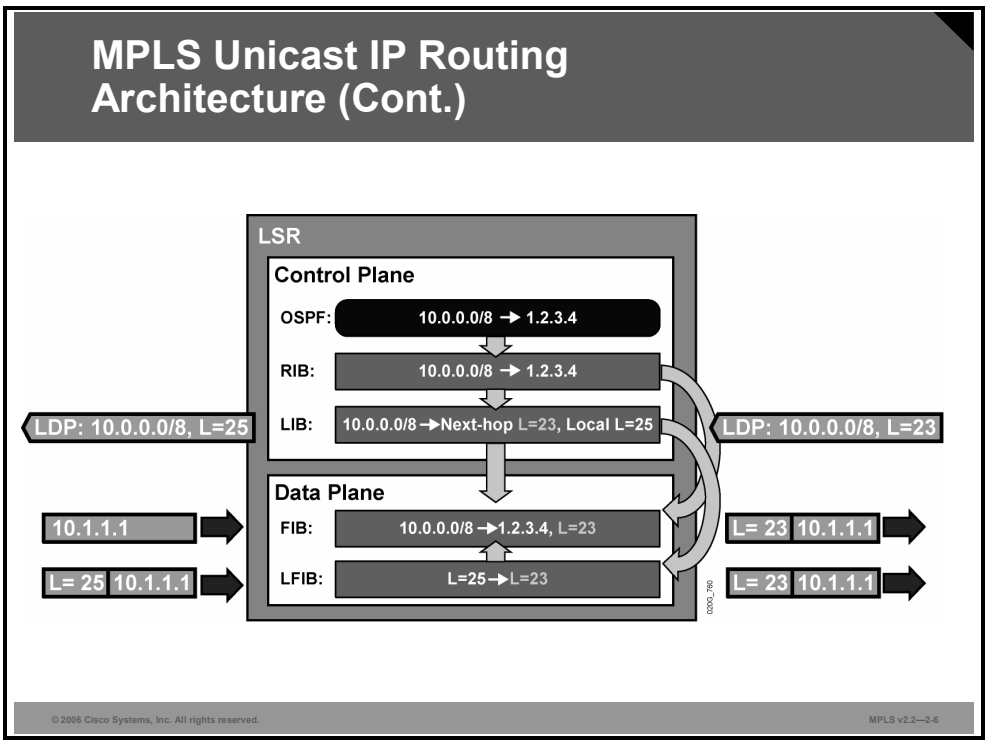

Example: Using LDP

The figure shows a router where OSPF is used to exchange IP routing information, and LDP is used to exchange labels. The router has attached the locally significant label 25 to the prefix 10.0.0.0/8 and advertised it to its neighbors. The label 23 has been assigned to prefix 10.0.0.0/8 by the upstream neighbor of the router (to the right in the diagram).

When an incoming IP packet to 10.1.1.1 arrives, it is forwarded by using the FIB table, where a next-hop label dictates that the outgoing packet should be labeled with label 23.

When a downstream router participating in MPLS receives a packet for 10.1.1.1, it will attach the label 25 and forward it to this router. When this router receives the labeled packet, it will swap the label value of 25 with a label value of 23 based on the LFIB table. With this process, the incoming (locally significant) label 25 is swapped with the next-hop label 23.

What Are LSPs?

This topic describes the function of label-switched paths (LSPs).

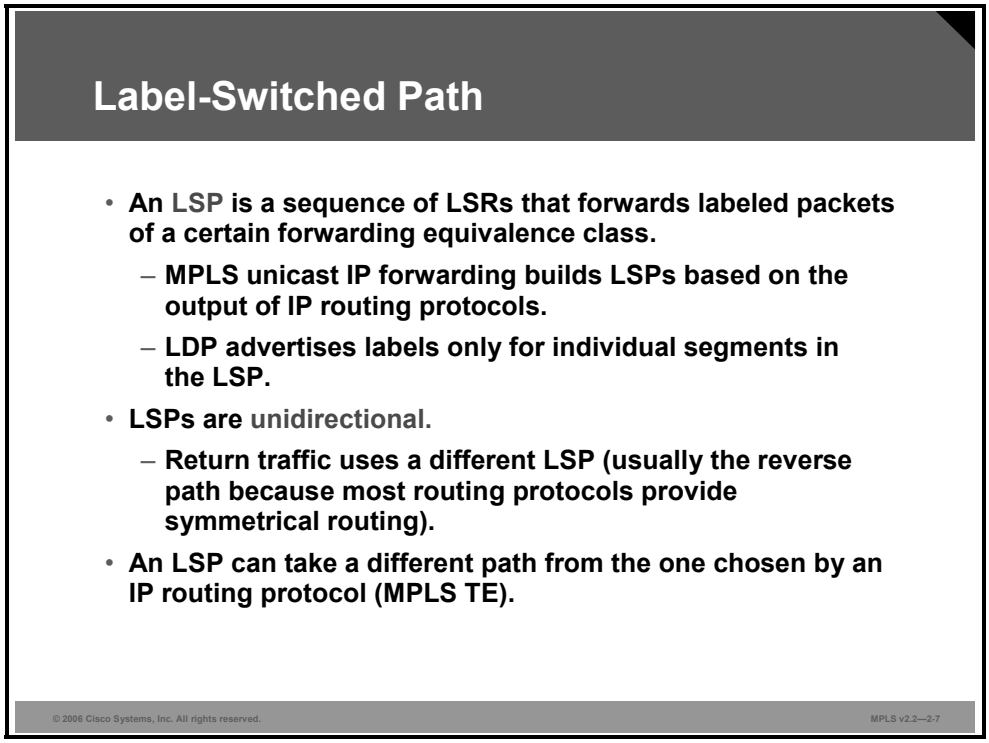

An LSP is a sequence of LSRs that forwards labeled packets for a particular FEC. Each LSR swaps the top label in a packet traversing the LSP. An LSP is similar to Frame Relay or ATM virtual circuits.

In MPLS unicast IP forwarding, the FECs are determined by destination networks found in the main routing table. Therefore, an LSP is created for each entry found in the main routing table. Border Gateway Protocol (BGP) entries are the only exceptions and are covered in the “MPLS Virtual Private Network Technology” module.

An Interior Gateway Protocol (IGP) is used to populate the routing tables in all routers in an MPLS domain. LDP is used to propagate labels for these networks and build LSPs.

LSPs are unidirectional. Each LSP is created over the shortest path, selected by the IGP, toward the destination network. Packets in the opposite direction use a different LSP. The return LSP is usually over the same LSRs, except that packets form the LSP in the opposite order.

Cisco MPLS Traffic Engineering (MPLS TE) can be used to change the default IGP shortest path selection.

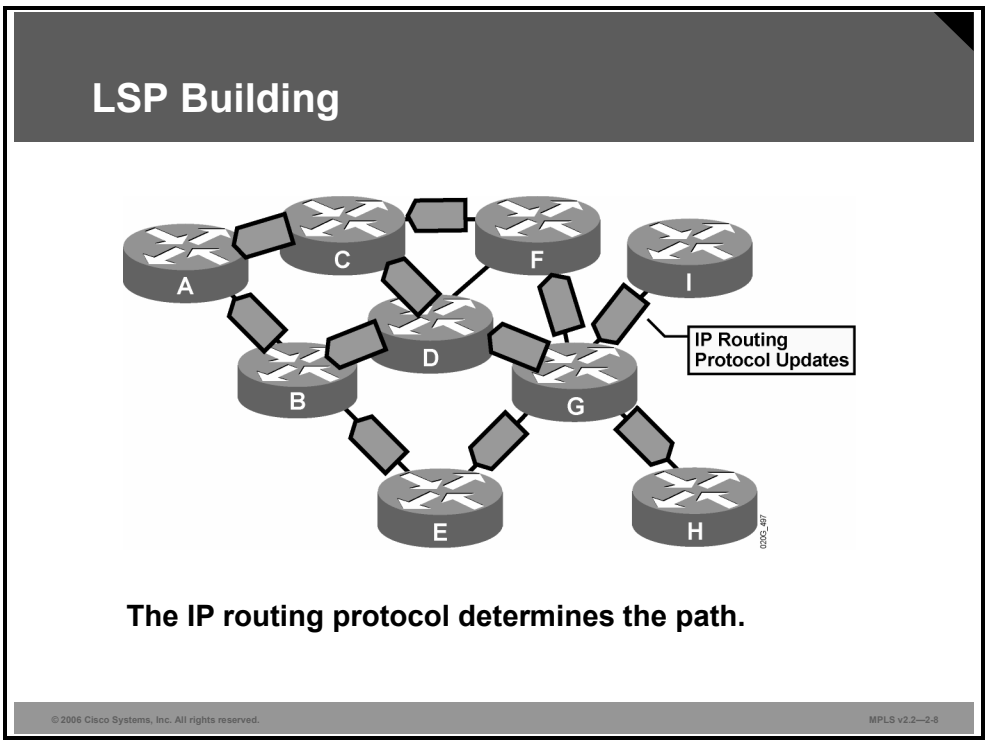

Example: IGP Propagates Routing Information

The figure illustrates how an IGP, such as OSPF, IS-IS, or EIGRP, propagates routing information to all routers in an MPLS domain. Each router determines its own shortest path.

LDP, which propagates labels for those networks and routers, adds labels to the FIB and LFIB tables.

In the figure below, an LSP is created for a particular network. This LSP starts on router A and follows the shortest path, determined by the IGP.

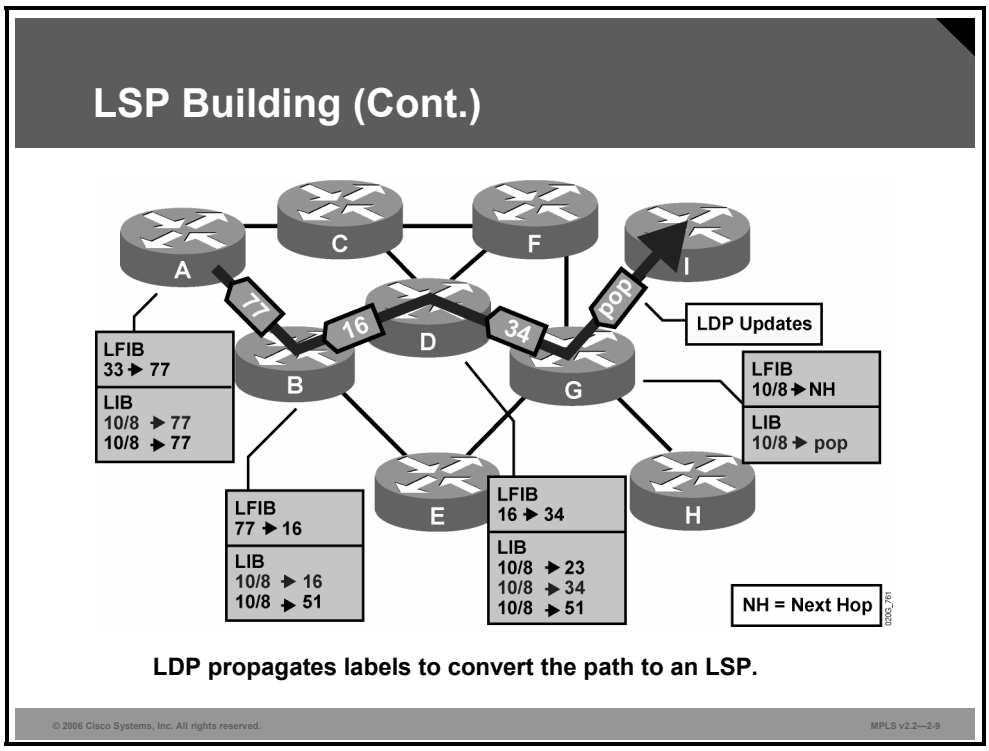

Example: LFIB and LIB Tables

The figure shows the contents of LFIB and LIB tables. Frame-mode MPLS uses a liberal label retention mode, which means that each LSR keeps all labels received from LDP peers, even if they are not the downstream peers for network X.

Liberal retention mode is evident after comparing the contents of the LIB and LFIB tables. Only those labels that come from the next-hop router are inserted into the LFIB table.

Note	Notice that router G receives a pop label from final destination router I. The pop action results in the removal of the label rather than swapping labels. This allows the regular IP packet to be forwarded.

Propagating Labels Using PHP

This topic describes the function of PHP.

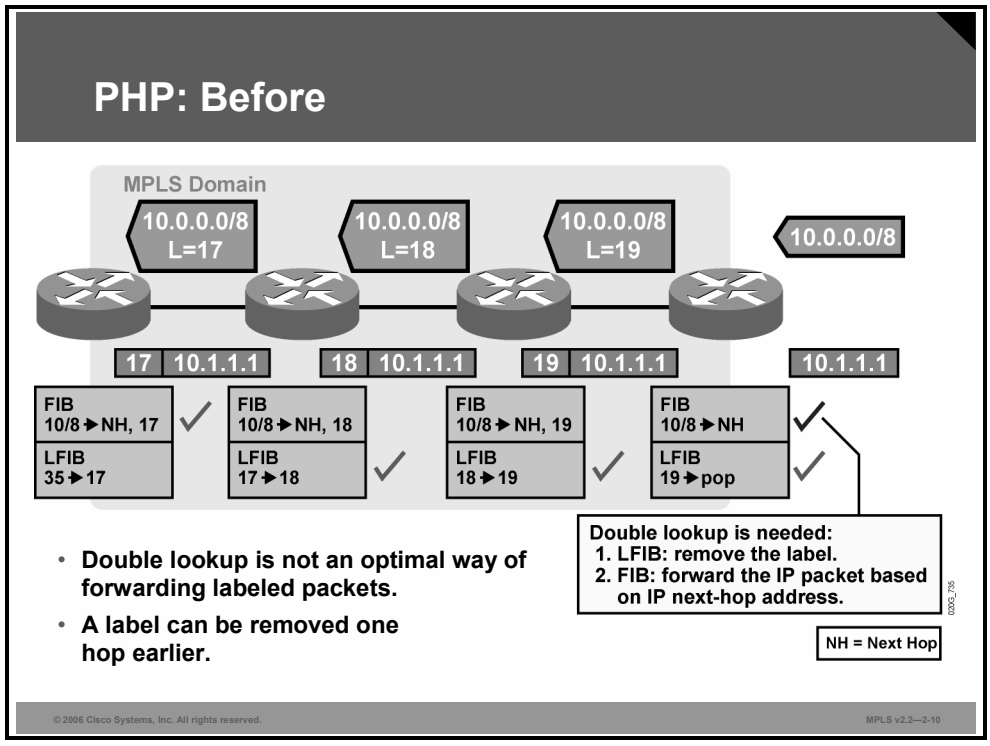

Example: PHP—Before

The figure illustrates how labels are propagated and used in a typical frame-mode MPLS network. The check marks show which tables are used on individual routers. The egress router in this example must do a lookup in the LFIB table to determine whether the label must be removed and if a further lookup in the FIB table is required.

PHP removes the requirement for a double lookup to be performed on egress LSRs.

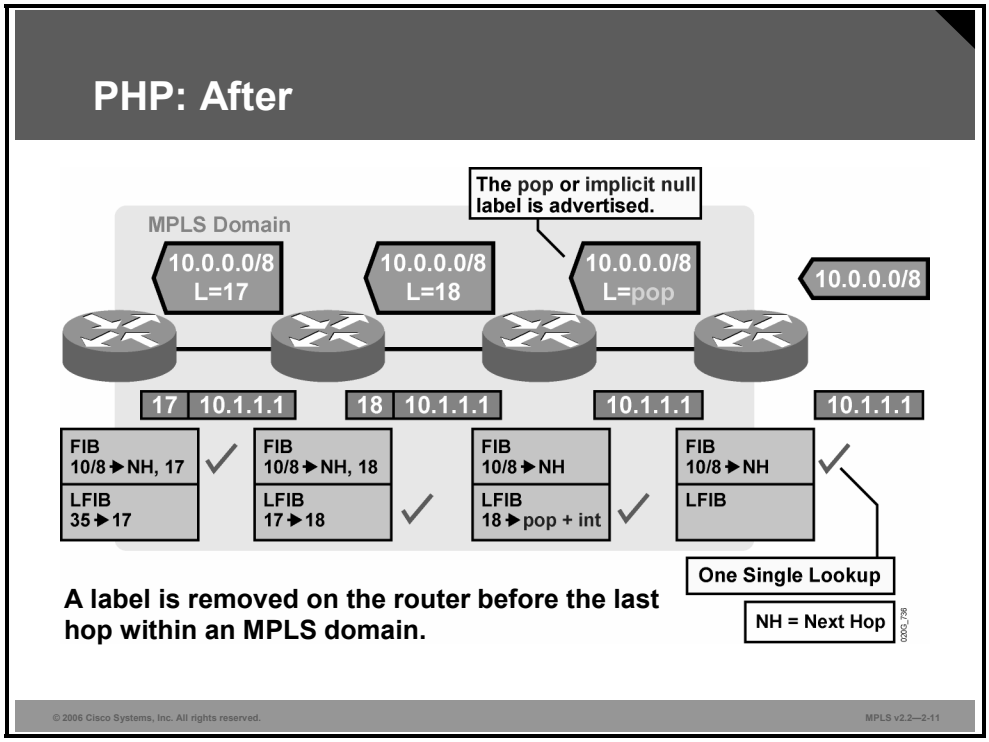

Example: PHP—After

The figure illustrates how a predefined label pop, which corresponds to the pop action in the LFIB, is propagated on the first hop or the last hop, depending on the perspective. The term "pop" means to remove the top label in the MPLS label stack instead of swapping it with the next-hop label. The last router before the egress router therefore removes the top label.

PHP slightly optimizes MPLS performance by eliminating one LFIB lookup.

PHP

- **Penultimate hop popping optimizes MPLS performance (one less LFIB lookup).**
- **PHP does not work on ATM. (virtual path identifier/virtual channel identifier cannot be removed.)**
- **The pop or implicit null label uses a reserved value when being advertised to a neighbor.**

PHP optimizes MPLS performance by reducing the number of table lookups on the egress router.

Note	A pop label is encoded with a value of 3 for LDP or a value of 1 for Tag Distribution Protocol (TDP). This label instructs upstream routers to remove the label instead of swapping it. What will be displayed in the LIB table of the router will be “imp-null” rather than the value of 3 or 1.

What Is the Impact of IP Aggregation on LSPs?

This topic describes the impact that IP aggregation has on LSPs.

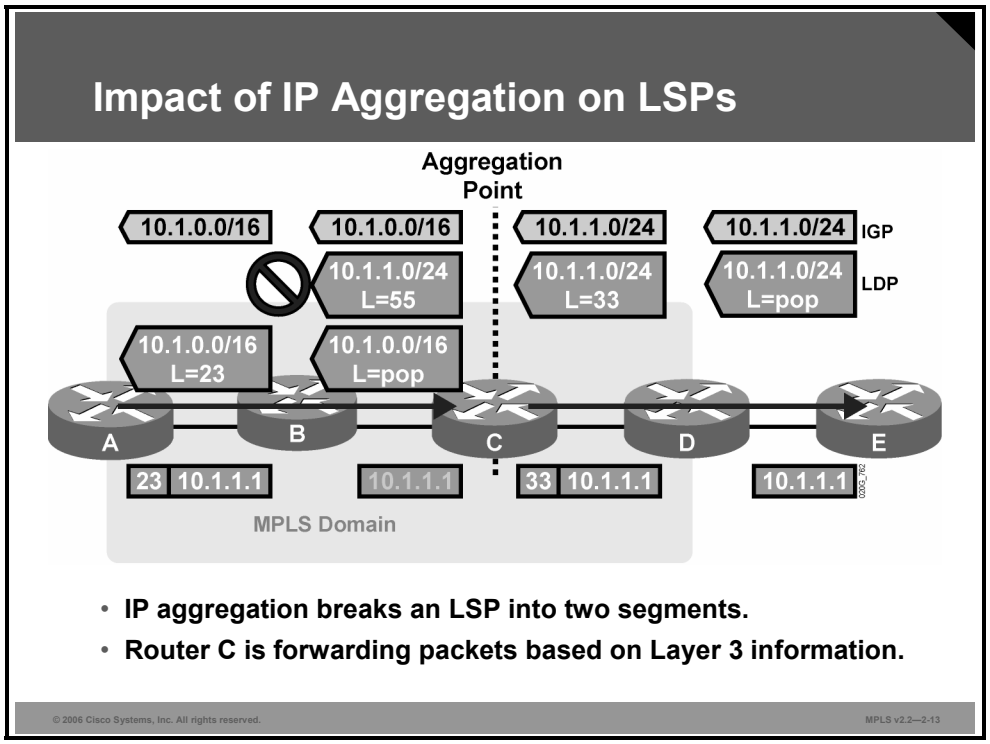

Example: MPLS IP Aggregation Problem

The figure illustrates a potential problem in an MPLS domain.

An IGP propagates the routing information for network 10.1.1.0/24 from router E to other routers in the network. Router C uses a summarization mechanism to stop the proliferation of all subnetworks of network 10.1.0.0/16. Only the summary network 10.1.0.0/16 is sent to routers B and A.

LDP propagates labels concurrently with the IGP. The LSR that is the endpoint of an LSP always propagates the “pop” label.

Router C has both networks in the routing table, as listed here:

- 10.1.1.0/24 (the original network)
- 10.1.0.0/16 (the summary)

Router C, therefore, sends a label, 55 in the example, for network 10.1.1.0/24 to router B. Router C also sends a pop label for the new summary network 10.1.0.0/16 that originates on this router. Router B, however, can use the pop label only for the summary network 10.1.0.0/16 because it has no routing information about the more specific network 10.1.1.0/24 because this information was suppressed on router C.

The summarization results in two LSPs for destination network 10.1.1.0/24. The first LSP ends on router C, where a routing lookup is required to assign the packet to the second LSP.

Impact of IP Aggregation on LSPs (Cont.)

- **IP aggregation breaks an LSP into two segments.**
- **Aggregation should not be used where end-to-end LSPs are required, such as with:**
  - **MPLS VPNs**
  - **MPLS TEs**
  - **MPLS-enabled ATM network**
  - **Transit BGP where core routers are not running BGP**

Aggregation should also not be used where an end-to-end LSP is required. Typical examples of networks that require end-to-end LSPs are as follows:

- An MPLS Virtual Private Network (VPN) backbone
- A network that uses MPLS TE
- An MPLS-enabled ATM network
- A transit BGP autonomous system (AS) where core routers are not running BGP

Allocating Labels in a Frame-Mode MPLS Network

This topic describes how labels are allocated and distributed in a frame-mode MPLS network.

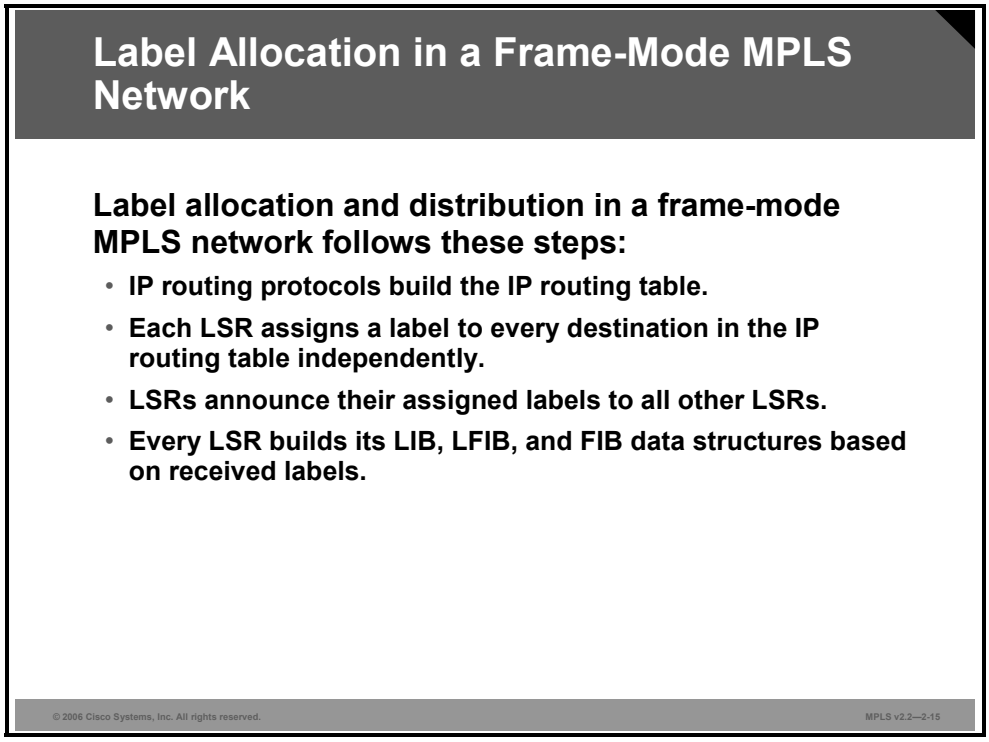

Unicast IP routing and MPLS functionality can be divided into these steps:

- Routing information exchange using standard or vendor-specific IP routing protocols (OSPF, IS-IS, EIGRP, and so on)
- Generation of local labels (One locally unique label is assigned to each IP destination found in the main routing table and stored in the LIB table.)
- Propagation of local labels to adjacent routers, where these labels might be used as next-hop labels (stored in the FIB and LFIB tables to enable label switching)

These data structures contain label information:

- The LIB, in the control plane, is the database used by LDP where an IP prefix is assigned a locally significant label that is mapped to a next-hop label that has been learned from a downstream neighbor.
- The LFIB, in the data plane, is the database used to forward labeled packets received by the router. Local labels, previously advertised to upstream neighbors, are mapped to next-hop labels, previously received from downstream neighbors.
- The FIB, in the data plane, is the database used to forward unlabeled IP packets received by the router. A forwarded packet is labeled if a next-hop label is available for a specific destination IP network. Otherwise, a forwarded packet is not labeled.

Example: Building the FIB Table

The figure illustrates that all routers learn about network X via an IGP such as OSPF, IS-IS, or EIGRP.

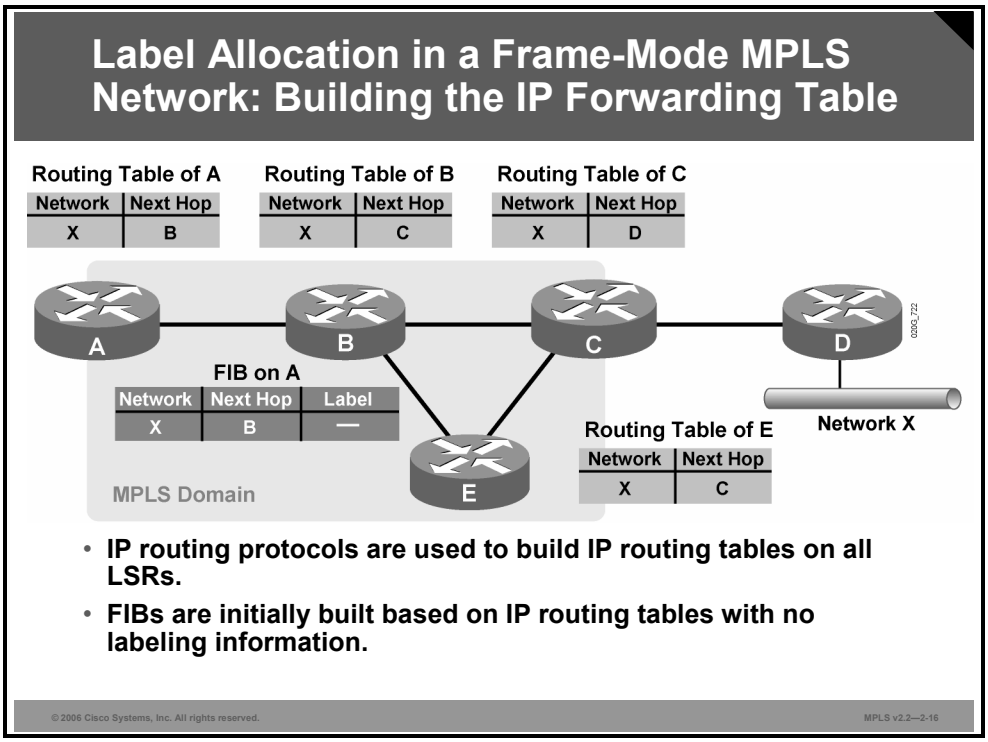

The FIB table on router A contains the entry for network X that is mapped to the IP next-hop address B. At this time, a next-hop label is not available, which means that all packets are forwarded in a traditional fashion (as unlabeled packets).

Example: Label Allocation

This example illustrates label allocation in a frame-mode MPLS network.

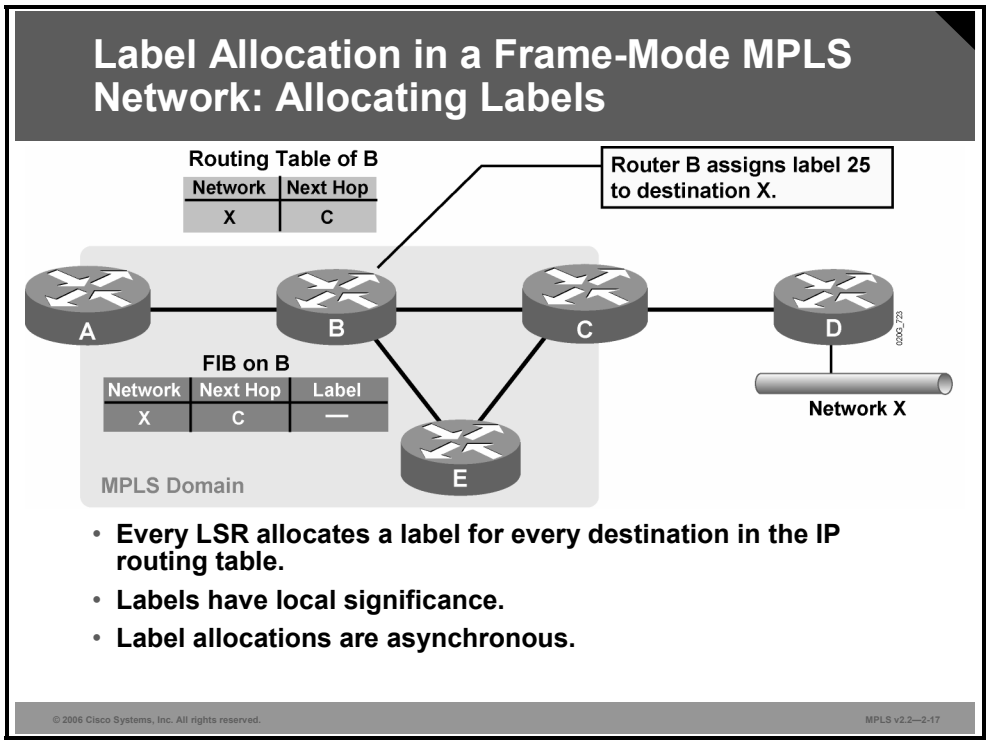

In this example, router B generates a locally significant and locally unique label 25 assigned to IP network X. Router B generates this label independently of other routers (asynchronous allocation of labels).

Note	Labels 0 to 15 are reserved. Each LSR independently assigns a local label to each non-BGP IP prefix in its routing table. Labels are not assigned to BGP routes in the routing table.

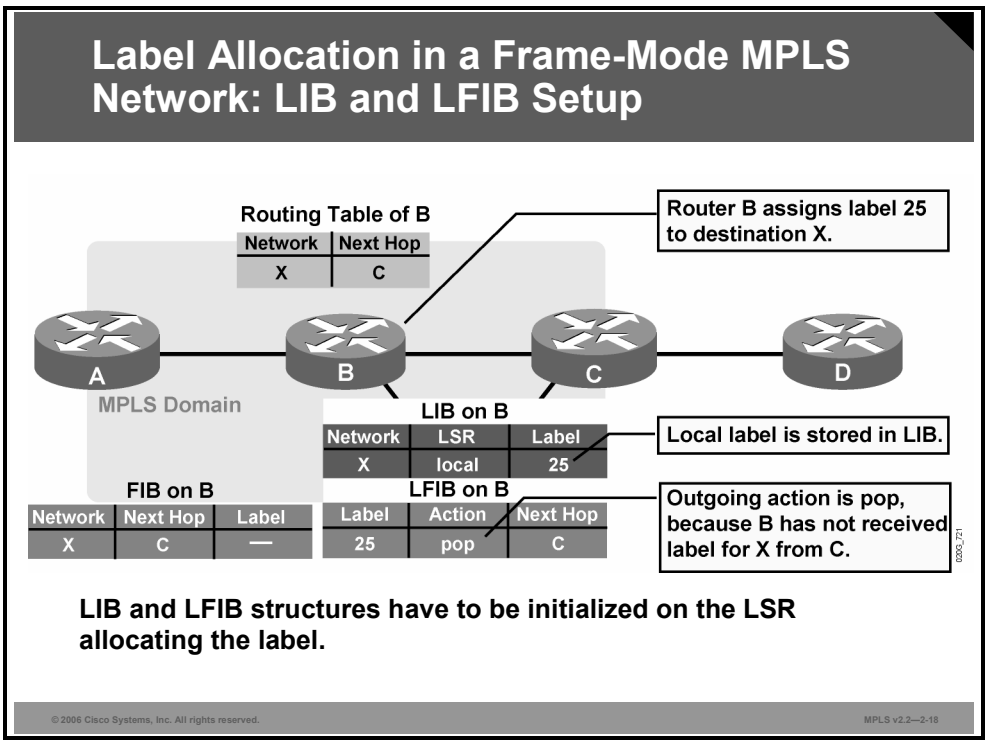

When a label is assigned to an IP prefix, it is stored in these two tables:

- The LIB table is used to maintain the mapping between the IP prefix (network X), the local label (25), and the next-hop label (not available yet).
- The LFIB table is modified to contain the local label mapped to the pop action (label removal). The pop action is used until the next-hop label is received from the downstream neighbor.

Note	The FIB table does not yet contain a label for forwarding packets to network X.

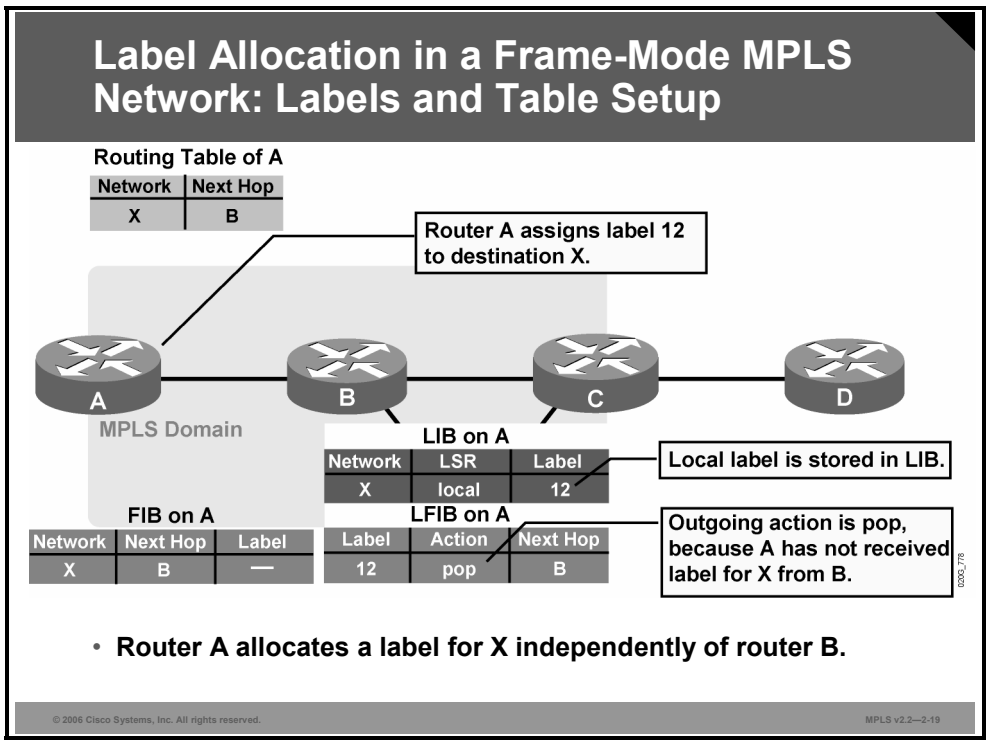

The FIB, LIB, and LFIB tables are updated similarly on router A after it allocates local label 12 for network X. The tables have specific roles:

- The LIB table is used to maintain the mapping between the IP prefix (network X), the local label (12), and the next-hop label (not available yet).
- The LFIB table is modified to contain the local label mapped to the pop action (label removal). The pop action is used until the next-hop label is received from the downstream neighbor.
- The FIB table does not yet contain a label for forwarding packets to network X.

Distributing and Advertising Labels

This topic describes how MPLS labels are distributed and advertised within an MPLS network.

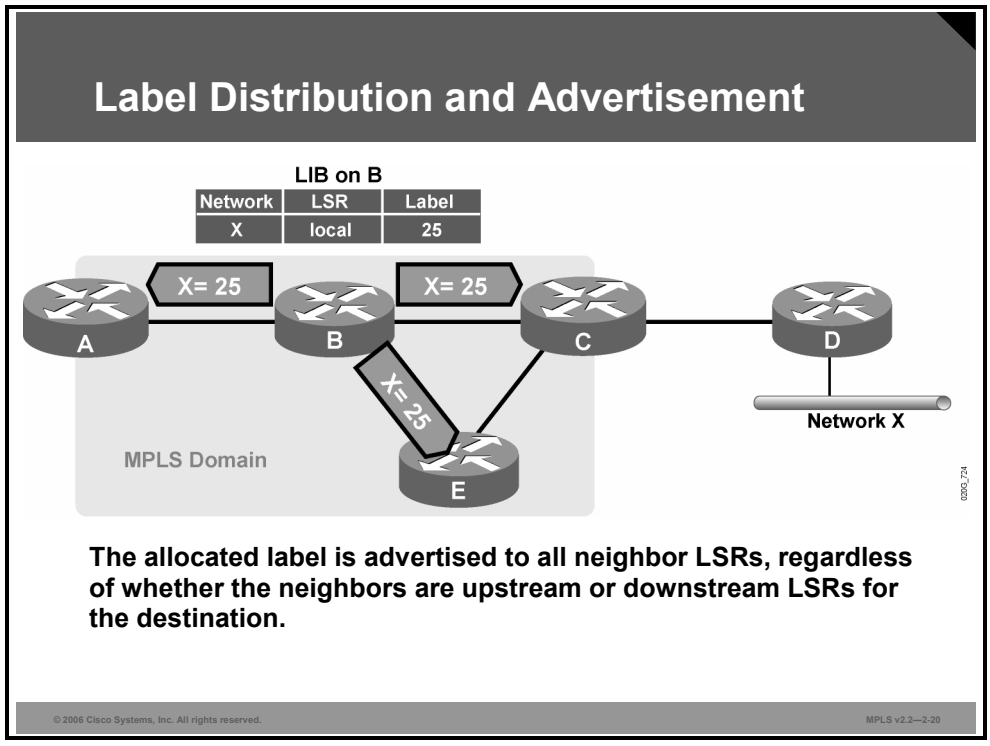

Example: Label Distribution and Advertisement

The figure illustrates the next step after a local label has been generated on router B. Router B propagates this label, 25, to all adjacent neighbors where this label can be used as a next-hop label.

Note	Because router B cannot predict which routers might use it as the downstream neighbor, router B sends its local mappings to all LDP neighbors.

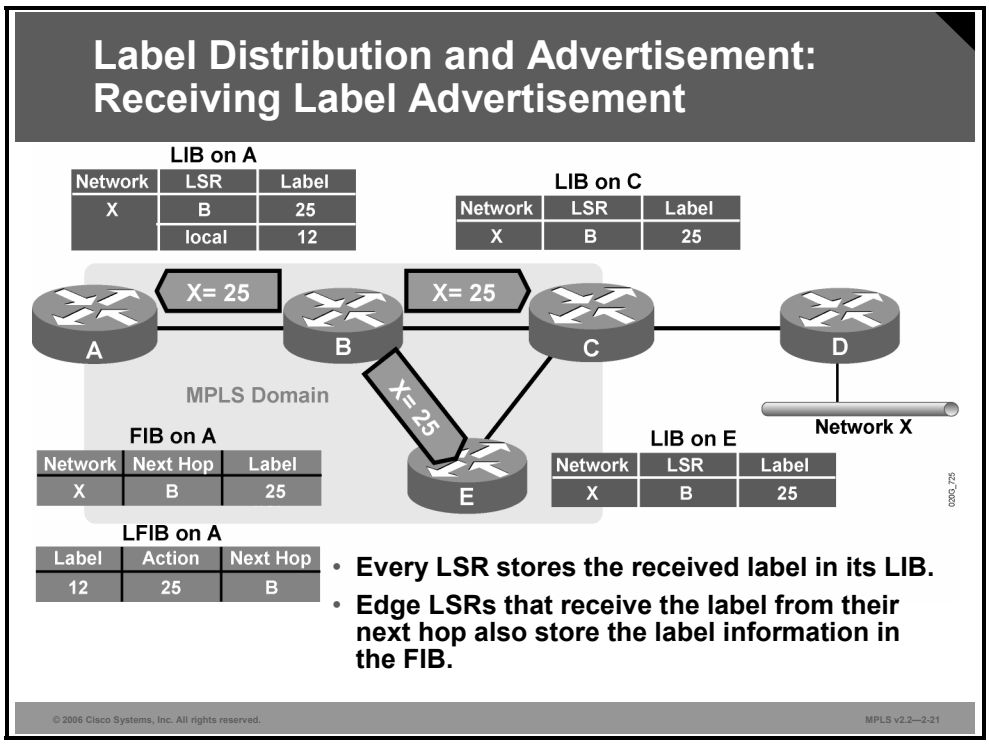

Upon receiving an LDP update, router A can fill in the missing piece in its LIB, LFIB, and FIB tables, as listed here:

- Label 25 is stored in the LIB table as the label for network X received from LSR B. (Label 25 is also stored in the LIB tables on routers C and E.)
- Label 25 is attached to the IP forwarding entry in the FIB table to enable the MPLS edge functionality (incoming IP packets are forwarded as labeled packets).
- The local label in the LFIB table is mapped to outgoing label 25 instead of the pop action (incoming labeled packets can be forwarded as labeled packets).

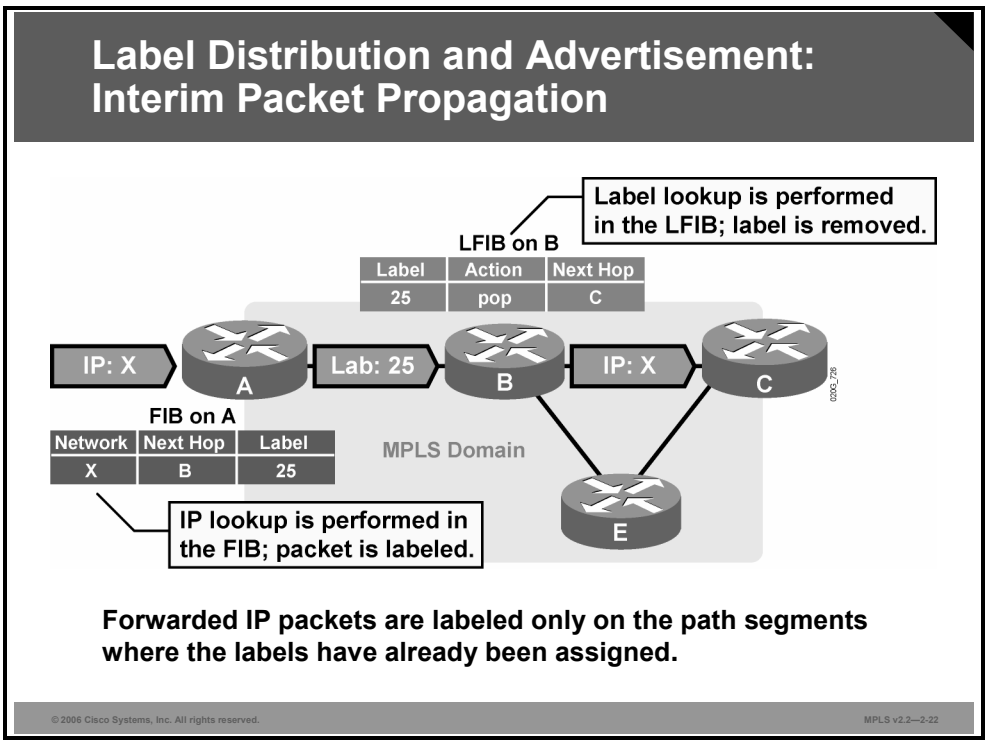

Example: Interim Packet Propagation Through an MPLS Network

The figure shows how an unlabeled IP packet is forwarded based on the information found in the FIB table on router A. Label 25, found in the FIB table, is used to label the packet.

Router B must remove the label because LSR B has not yet received any next-hop label from router C (the action in the LFIB is “pop”).

Note	The LFIB on router B is not yet complete.

Router A performs an IP lookup (CEF switching), whereas router B performs a label lookup (label switching) in which the label is removed and a normal IP packet is sent out of router B.

Example: LDP Update Sent to All Adjacent Routers

The figure illustrates how an LDP update, advertising label 47 for network X, from router C is sent to all adjacent routers, including router B.

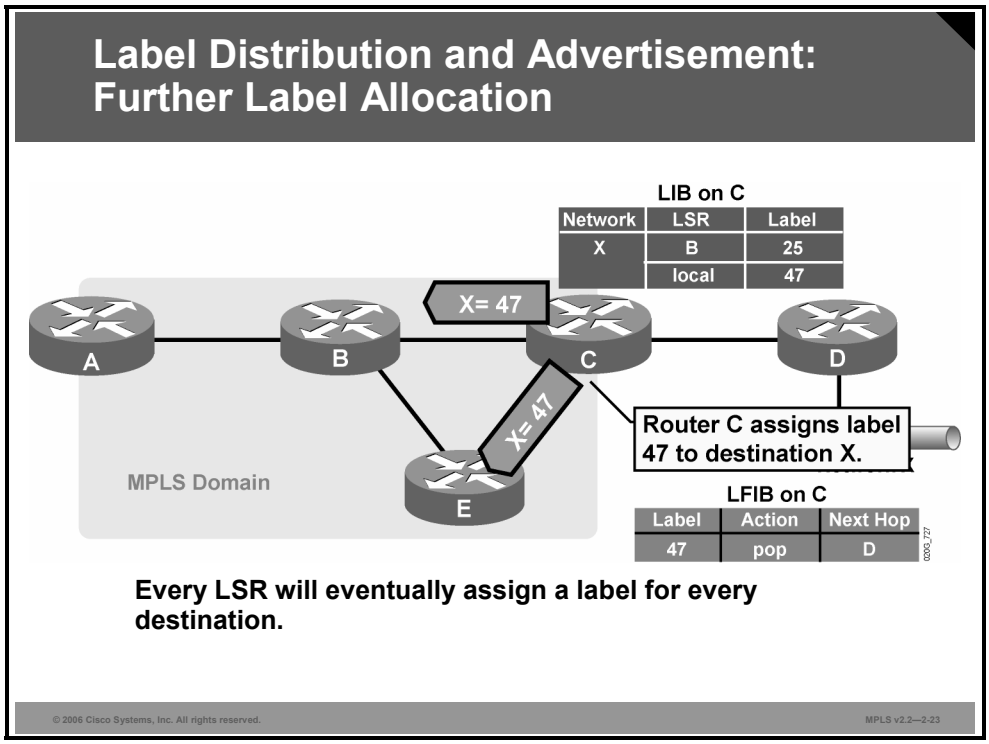

After all routers in an MPLS domain independently distribute their labels as routers A and B did, an LSP tunnel exists for network X spanning from router A to router C.

Note	This example discussed only the label allocation and distribution process for one prefix. In actual practice, the LSRs and edge LSRs would concurrently allocate and distribute labels for all of the prefixes in their routing table.

Label Distribution and Advertisement: Receiving Label Advertisement

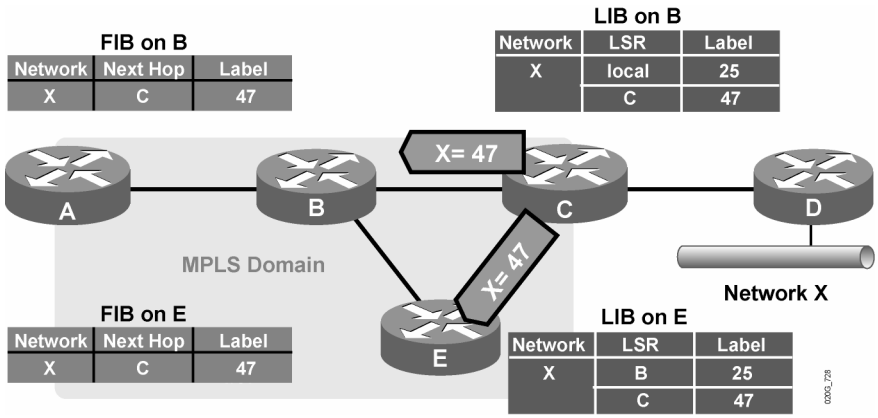

- **Every LSR stores received information in its LIB.**
- **LSRs that receive their label from their next-hop LSR will also populate the IP forwarding table.**

Router B can now map the entry for network X in its FIB, and the local label 25 in its LIB, to the next-hop label 47 received from the downstream neighbor router C.

Populating the LFIB

This topic describes how the LFIB table is populated in an MPLS network.

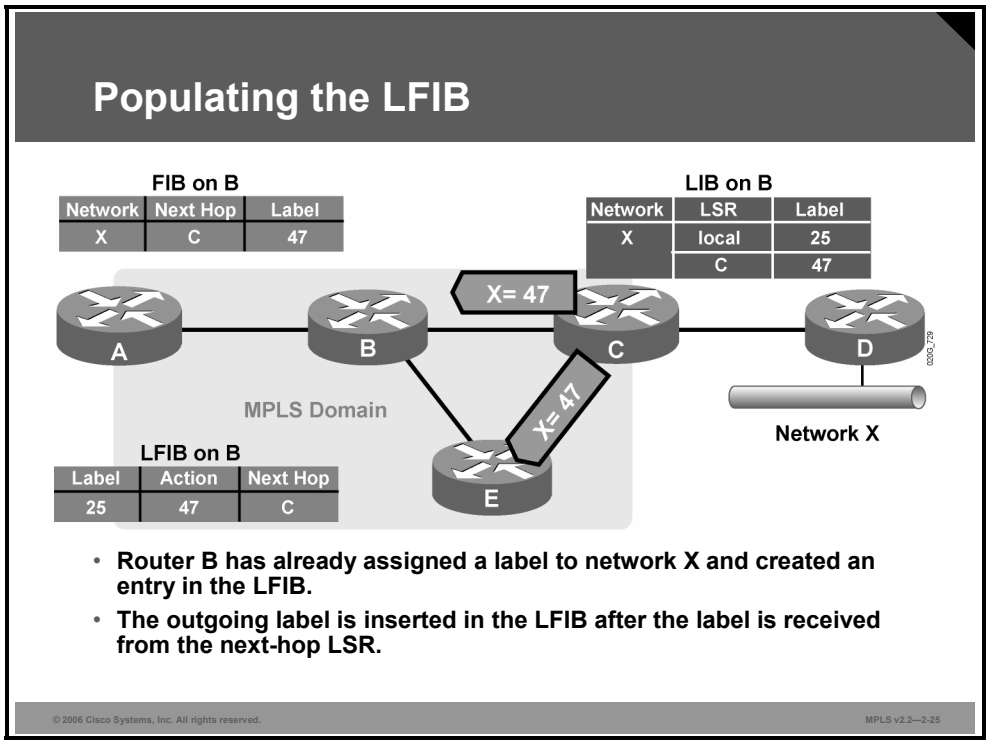

Example: LFIB Population

After router C advertises label 47 to adjacent routers, the LSP tunnel for network X has two hops. The components are as follows:

- On router A, network X is mapped to the next-hop label 25 (router B).
- On router B, label 25 is mapped to the next-hop label 47 (router C).
- Router C still has no next-hop label. Label 47 is therefore mapped to the pop action.

Note	In the figure, label distribution is from right to left, and packet forwarding is from left to right.

Propagating Packets Across an MPLS Network

This topic describes how IP packets cross an MPLS network.

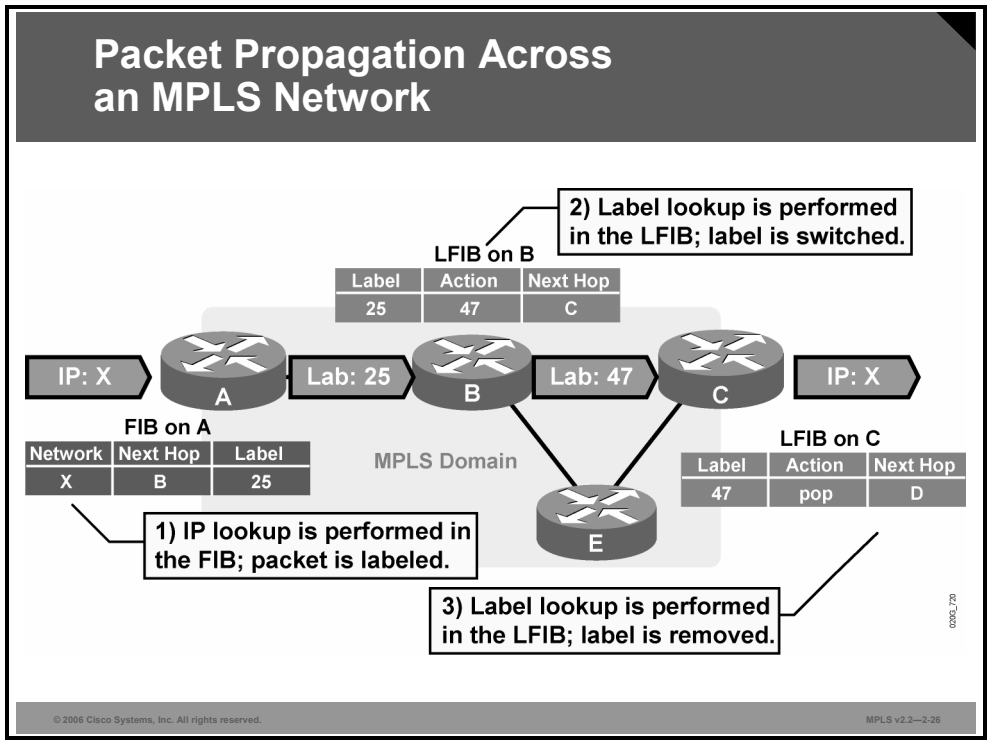

Example: Packet Propagation Through an MPLS Network

The figure illustrates how IP packets are propagated across an MPLS domain. The steps are as follows:

Step 1 Router A labels a packet destined for network X by using the next-hop label 25 (CEF switching by using the FIB table).

Step 2 Router B swaps label 25 with label 47 and forwards the packet to router C (label switching by using the LFIB table).

Step 3 Router C removes the label and forwards the packet to router D (label switching by using the LFIB table).

Detecting Frame-Mode Loops

This topic describes how frame-mode loops are detected.

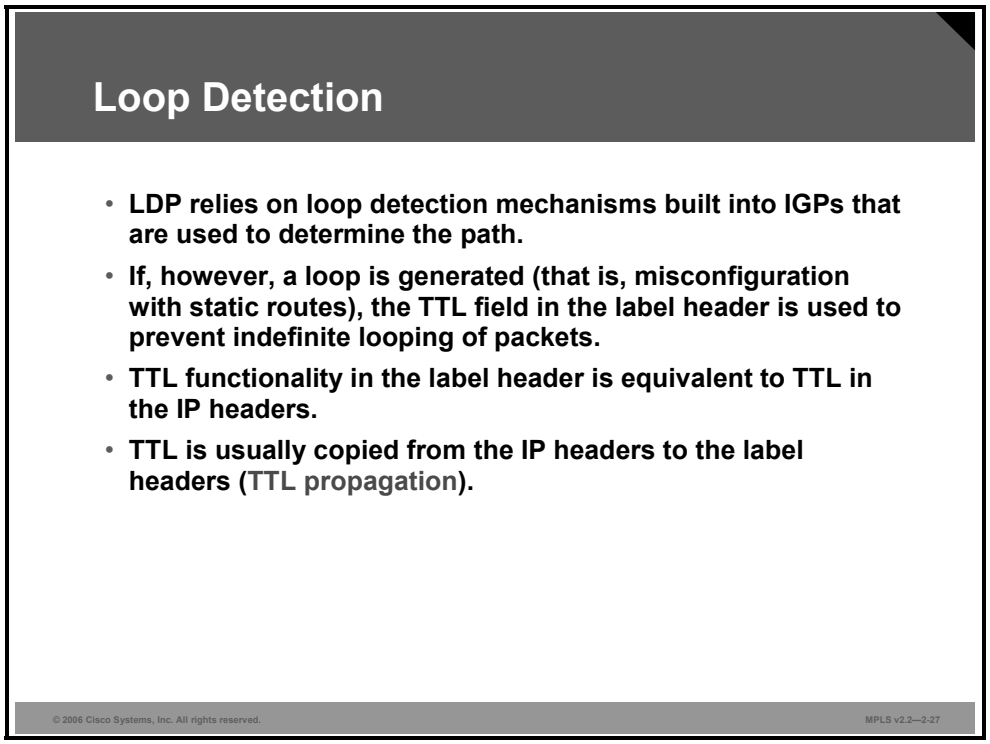

Loop detection in an MPLS-enabled network relies on more than one mechanism.

Most routing loops are prevented by the IGP used in the network. MPLS for unicast IP forwarding simply uses the shortest paths determined by the IGP. These paths are typically loop-free.

If, however, a routing loop does occur (for example, because of misconfigured static routes), MPLS labels also contain a time-to-live (TTL) field that prevents packets from looping indefinitely.

The TTL functionality in MPLS is equivalent to that of traditional IP forwarding. Furthermore, when an IP packet is labeled, the TTL value from the IP header is copied into the TTL field in the label. This is called TTL propagation.

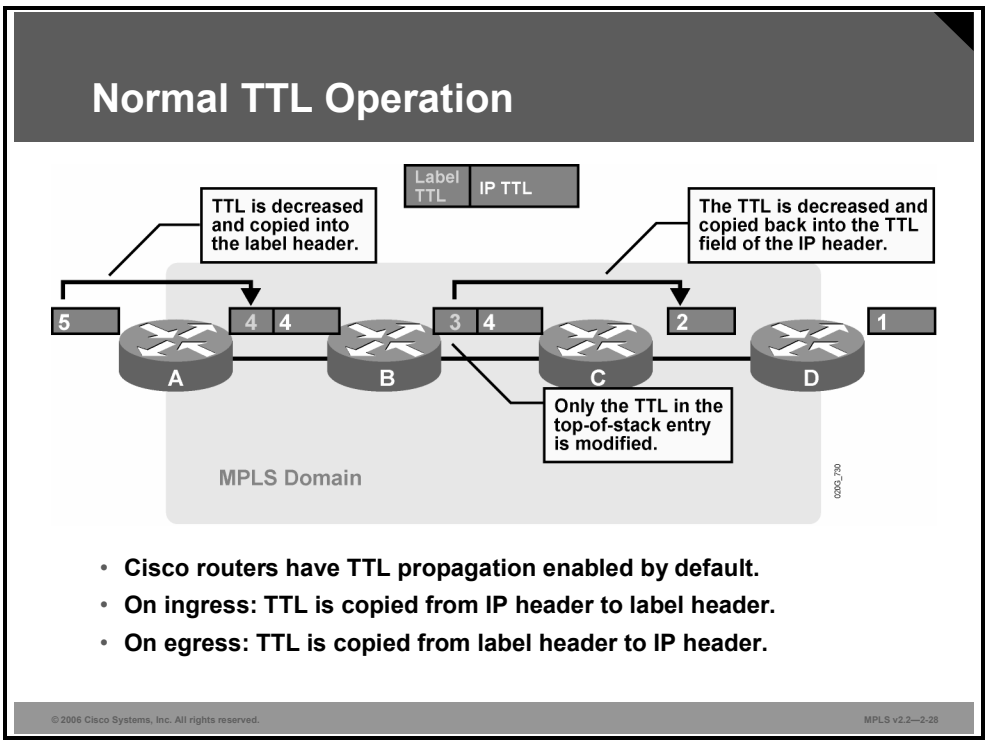

Example: Normal TTL Operation

The figure illustrates how the TTL value 5 in the IP header is decreased and copied into the TTL field of the label when a packet enters an MPLS domain.

All other LSRs decrease the TTL field only in the label. The original TTL field is not changed until the last label is removed when the label TTL is copied back into the IP TTL.

TTL propagation provides a transparent extension of IP TTL functionality into an MPLS-enabled network.

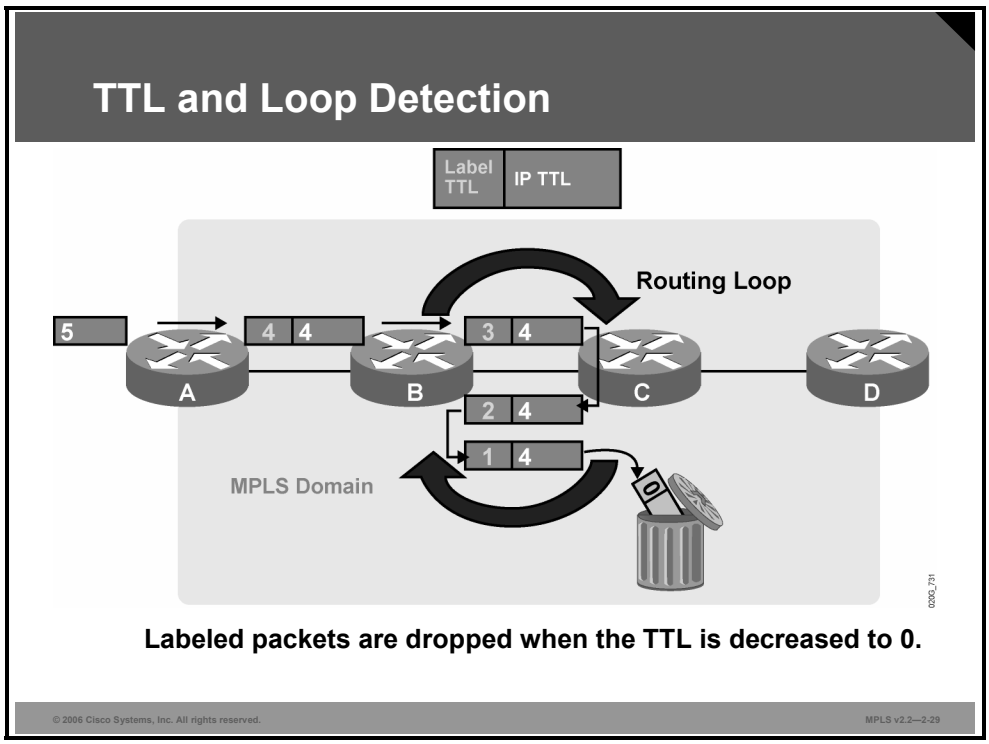

Example: TTL and Loop Detection

The figure illustrates a routing loop between routers B and C. The packet looping between these two routers is eventually dropped because the value of its TTL field reaches 0.

Disabling TTL Propagation

- **TTL propagation can be disabled.**
- **The IP TTL value is not copied into the TTL field of the label, and the label TTL is not copied back into the IP TTL.**
- **Instead, the value 255 is assigned to the label header TTL field on the ingress LSR.**
- **Disabling TTL propagation hides core routers in the MPLS domain.**
- **Traceroute across an MPLS domain does not show any core routers.**

TTL propagation can be disabled to hide the core routers from the end users. Disabling TTL propagation causes routers to set the value 255 into the TTL field of the label when an IP packet is labeled.

The network is still protected against indefinite loops, but it is unlikely that the core routers will ever have to send an Internet Control Message Protocol (ICMP) reply to user-originated traceroute packets.

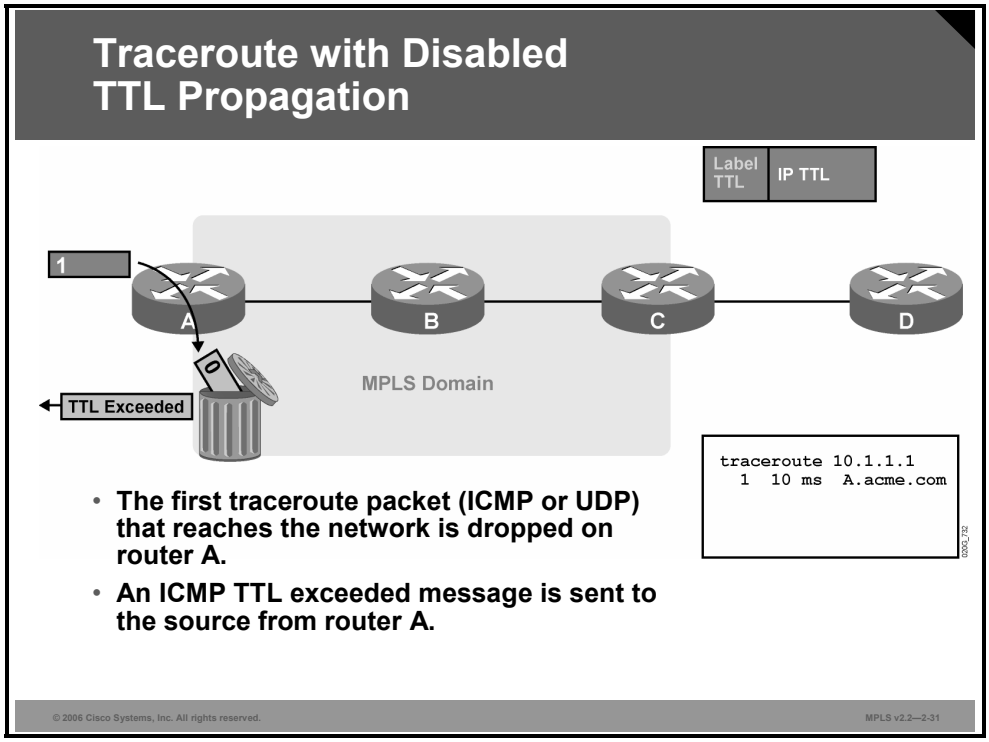

Example: Traceroute with Disabled TTL Propagation

These figures illustrate the result of a traceroute across an MPLS network that does not use TTL propagation.

The first traceroute packet—ICMP or User Datagram Protocol (UDP)—that reaches the MPLS network is dropped on the first router (A), and an ICMP reply is sent to the source. This action results in an identification of router A by the traceroute application.

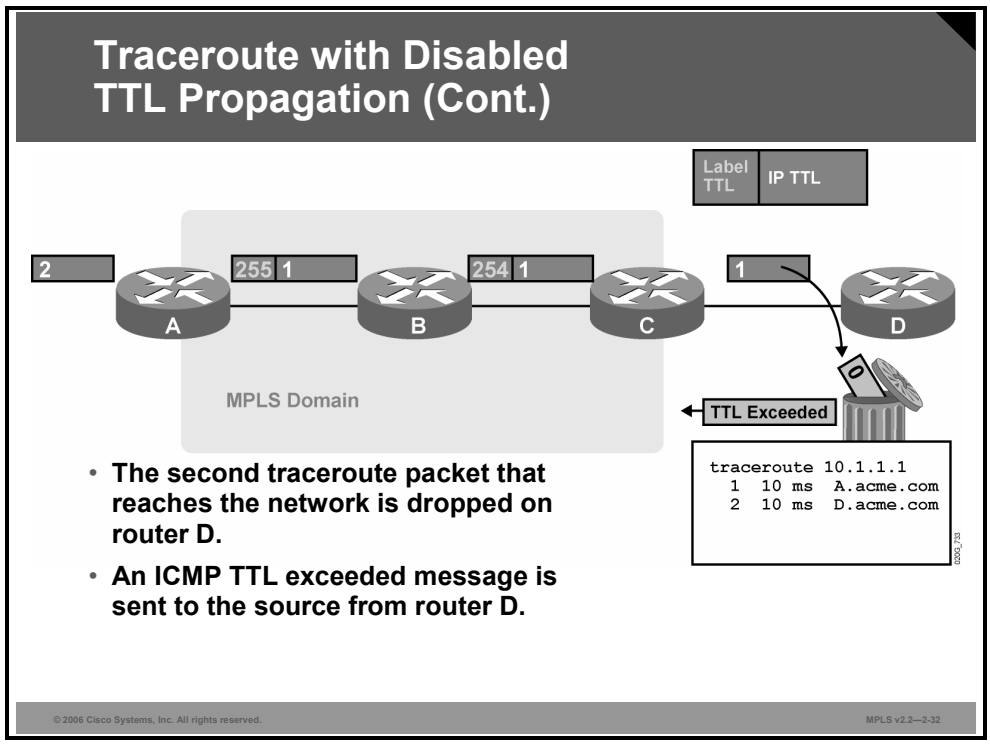

The traceroute application increases the initial TTL for every packet that it sends. The second packet, therefore, would be able to reach one hop farther (router B in the example). However, the TTL value is not copied into the TTL field of the label. Instead, router A sets the TTL field of the label to 255. Router B decreases the TTL of the label, and router C removes the label without copying it back into the IP TTL. Router D then decreases the original IP TTL, drops the packet because the TTL has reached zero, and sends an ICMP reply to the source.

The traceroute application has identified router D. The next packets would simply pass through the network.

The final result is that a traceroute application was able to identify the edge LSRs, but not the core LSRs.

Impact of Disabling TTL Propagation

- **Traceroute across an MPLS domain does not show core routers.**
- **TTL propagation has to be disabled on all label switch routers.**
- **Mixed configurations (some LSRs with TTL propagation enabled and some with TTL propagation disabled) could result in faulty traceroute output.**
- **TTL propagation can be enabled for forwarded traffic only—traceroute from LSRs does not use the initial TTL value of 255.**

Cisco routers have TTL propagation enabled by default.

If TTL propagation is disabled, it must be disabled on all routers in an MPLS domain to prevent unexpected behavior.

TTL can be optionally disabled for forwarded traffic only, which allows administrators to use traceroute from routers to troubleshoot problems in the network.

Allocating Per-Platform Labels

This topic describes the approaches for assigning labels to networks.

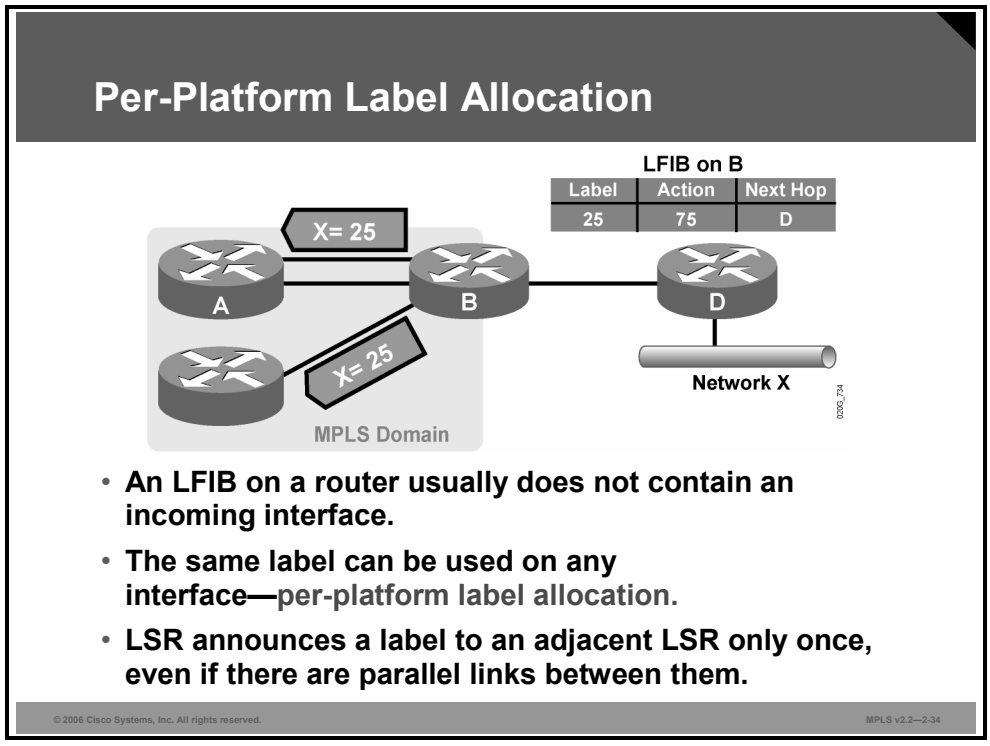

Here are the two possible approaches for assigning labels to networks:

- **Per-platform label allocation:** One label is assigned to a destination network and announced to all neighbors. The label must be locally unique and valid on all incoming interfaces. This is the default operation in frame-mode MPLS.
- **Per-interface label allocation:** Local labels are assigned to IP destination prefixes on a per-interface basis. These labels must be unique on a per-interface basis.

Example: Per-Platform Label Allocation

The figure illustrates how one label (25) is assigned to a network and used on all interfaces. The same label is propagated to both routers A and C.

The figure also shows how one label is sent across one LDP session between routers A and B even though there are two parallel links between the two routers.

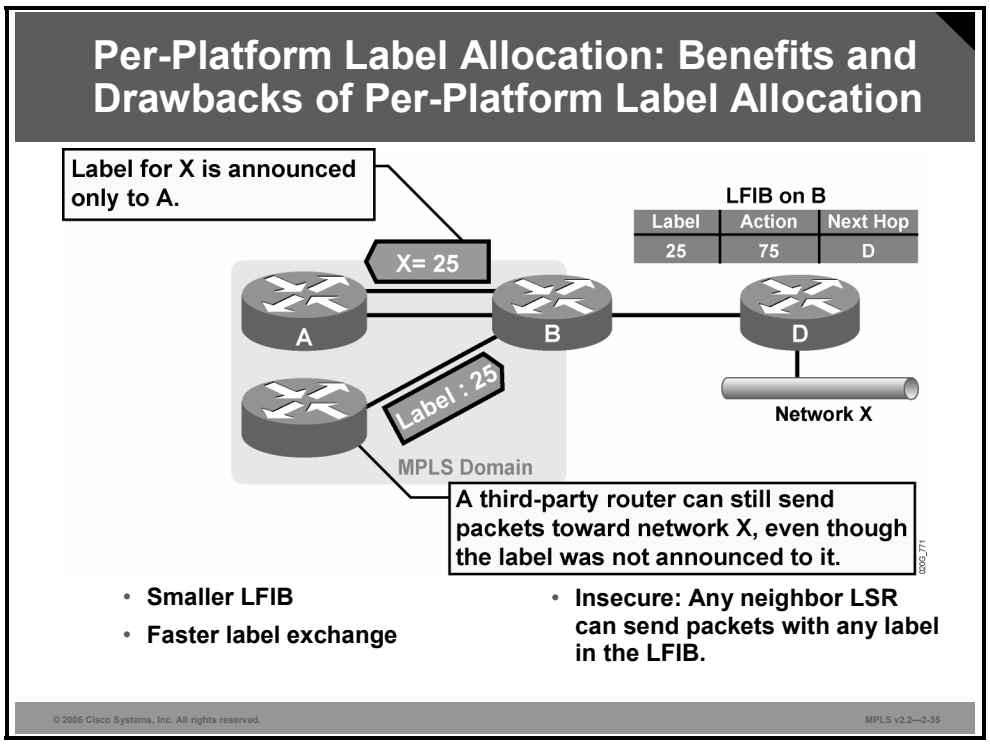

A potential drawback of per-platform label allocation is illustrated in the figure, which shows how an adjacent router can send a labeled packet with a label that has not been previously advertised to this router (label spoofing). If label switching has not been enabled on that interface, the packet will be discarded. If label switching has been enabled on this interface, the packet would be forwarded, causing a possible security issue.

On the other hand, per-platform label allocation results in smaller LIB and LFIB tables and a faster exchange of labels.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Labels are propagated across a network either by extending the functionality of existing routing protocols or by creating a new protocol that is dedicated to exchanging labels.
- An LSP is a sequence of LSRs that forward labeled packets of a certain forwarding equivalence class.
- Penultimate hop popping optimizes MPLS performance (one less LFIB lookup).
- IP aggregation can break an LSP into two segments.
- Every LSR assigns a label for every destination in the IP routing table.

Summary (Cont.)

- Although labels are locally significant, they have to be advertised to directly reachable peers.
- Outgoing labels are inserted in the LFIB after the label is received from the next-hop LSR.
- Packets are forwarded using labels from the LFIB table rather than the IP routing table.
- If TTL propagation is disabled, traceroute across an MPLS domain does not show core routers.
- LSR announces a label to an adjacent LSR only once, even if there are parallel links between them.

Lesson 3

Introducing Convergence in Frame-Mode MPLS

Overview

This lesson presents Label Distribution Protocol (LDP) convergence issues and describes how routing protocols and Multiprotocol Label Switching (MPLS) convergence interact. This lesson concludes with a look at link failure, convergence after a link failure, and link recovery.

It is important to understand the convergence times for LDP. It also is important to understand how routing protocols interact with MPLS. This information will ensure a clear understanding of how the various routing tables are built and refreshed during and after a link failure and how recovery in an MPLS network takes place.

Objectives

Upon completing this lesson, you will be able to describe how convergence occurs in a frame-mode MPLS network. This ability includes being able to meet these objectives:

- Describe the MPLS steady-state environment
- Describe what happens in the routing tables when a link failure occurs
- Describe routing protocol convergence after a link failure
- Describe frame-mode MPLS convergence after a link failure
- Describe IP and MPLS convergence actions after a link failure has been resolved

What Is the MPLS Steady-State Operation?

This topic describes an MPLS network steady-state operation.

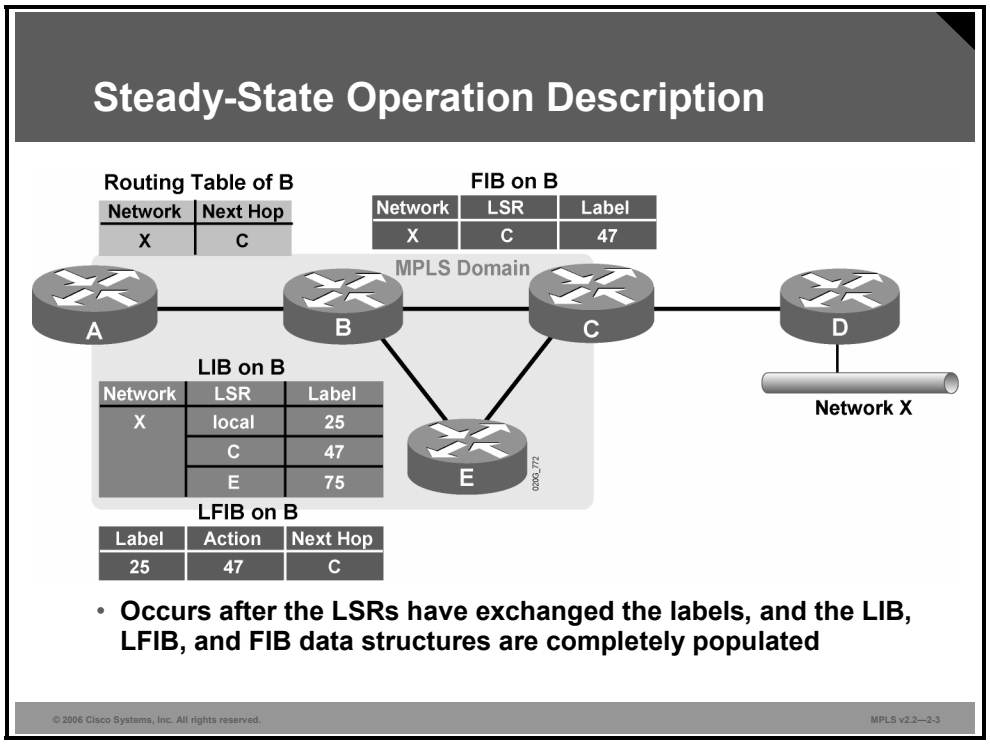

MPLS is fully functional when the Interior Gateway Protocol (IGP) and LDP have populated all the tables, as listed here:

- Main IP routing (routing information base [RIB]) table
- Label Information Base (LIB) table
- Forwarding Information Base (FIB) table
- Label forwarding information base (LFIB) table

Although it takes longer for LDP to exchange labels (compared with an IGP), a network can use the FIB table in the meantime; therefore, there is no routing downtime while LDP exchanges labels between adjacent LSRs.

What Happens in a Link Failure?

This topic describes what happens in the routing tables when a link failure occurs.

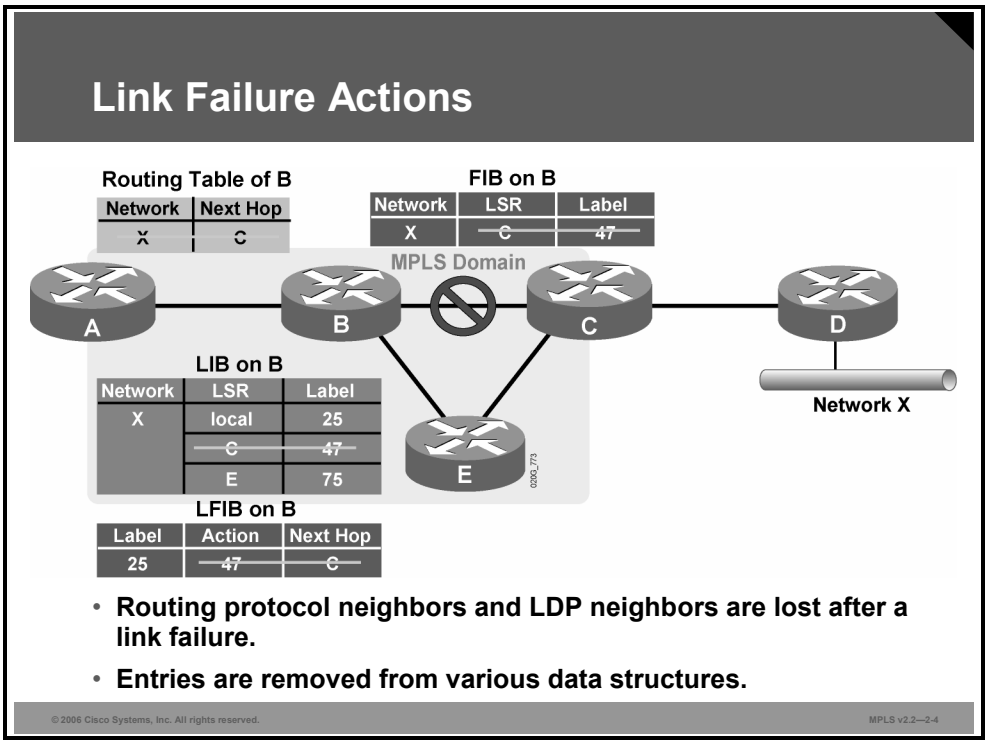

Example: Link Failure Actions

The figure illustrates how a link failure is handled in an MPLS domain. The steps are as follows:

- The overall convergence fully depends on the convergence of the IGP used in the MPLS domain.
- When router B determines that router E should be used to reach network X, the label learned from router E can be used to label-switch packets.

LDP stores all labels in the LIB table, even if the labels are not used, because the IGP has decided to use another path.

This label storage is shown in the figure, where two next-hop labels were available in the LIB table on router B. The label status of router B just before MPLS label convergence is as follows:

- Label 47 was learned from router C and is currently unavailable; therefore, because of the failure, label 47 has to be removed from the LIB table.
- Label 75 was learned from router E and can now be used at the moment that the IGP decides that router E is the next hop for network X.

What Is the Routing Protocol Convergence After a Link Failure?

This topic describes the routing protocol convergence that occurs in an MPLS network after a link failure.

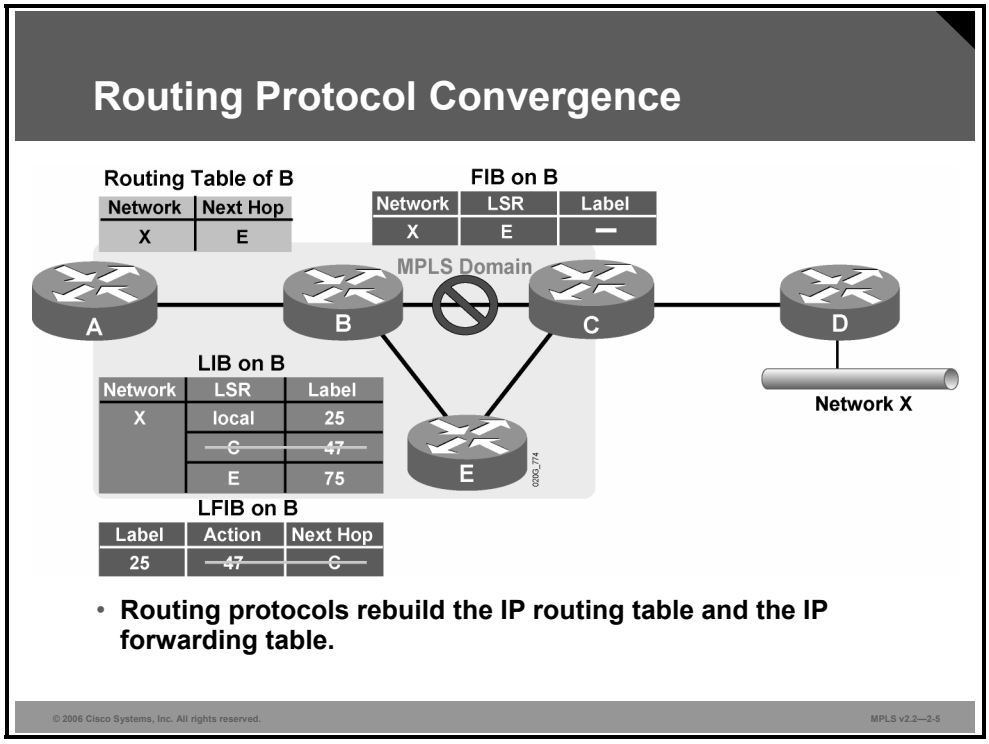

Example: Routing Protocol Convergence

The figure illustrates how two entries are removed, one from the LIB table and one from the LFIB table, when the link between routers B and C fails. This can be described as follows:

- Router B has already removed the entry from the FIB table, once the IGP determined that the next hop was no longer reachable.
- Router B has also removed the entry from the LIB table and the LFIB table given that the LDP has determined that router C is no longer reachable.

What Is the MPLS Convergence After a Link Failure?

This topic describes MPLS convergence that occurs in an MPLS network after a link failure.

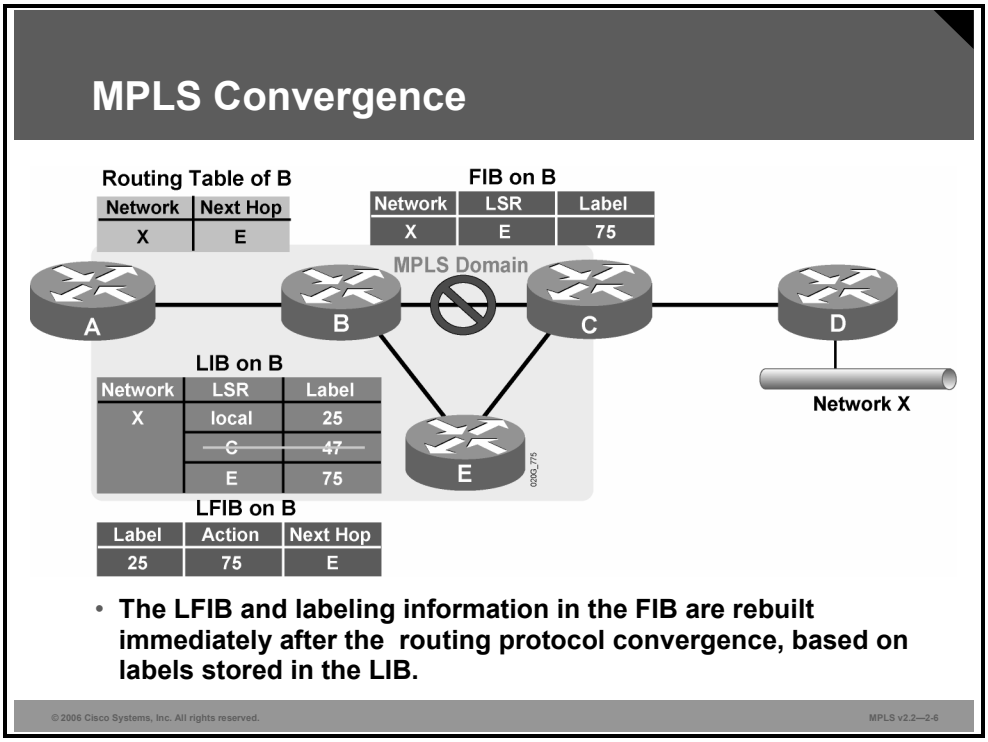

After the IGP determines that there is another path available, a new entry is created in the FIB table.

This new entry points toward router E, and there is already a label available for network X via router E.

This information is then used in the FIB table and the LFIB table to reroute the LSP tunnel via router E.

MPLS Convergence After a Link Failure

- **MPLS convergence in frame-mode MPLS does not affect the overall convergence time.**
- **MPLS convergence occurs immediately after the routing protocol convergence, based on labels already stored in the LIB.**

The overall convergence in an MPLS network is not affected by LDP convergence when there is a link failure.

Frame-mode MPLS uses liberal label retention mode, which enables routers to store all received labels, even if the labels are not being used.

These labels can be used, after the network convergence, to enable immediate establishment of an alternative LSP tunnel.

What Actions Occur in Link Recovery?

This topic describes actions in IP and MPLS convergence after a failure has been resolved.

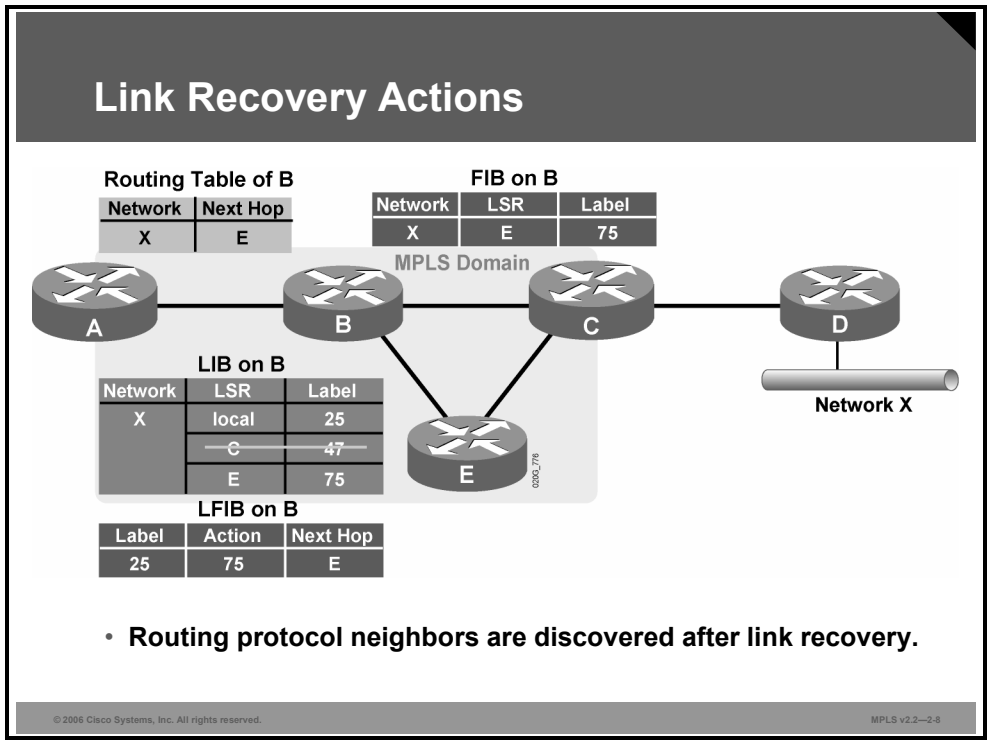

Example: Link Recovery Actions

The figure illustrates the state of the routing tables at the time the link between routers B and C becomes available again.

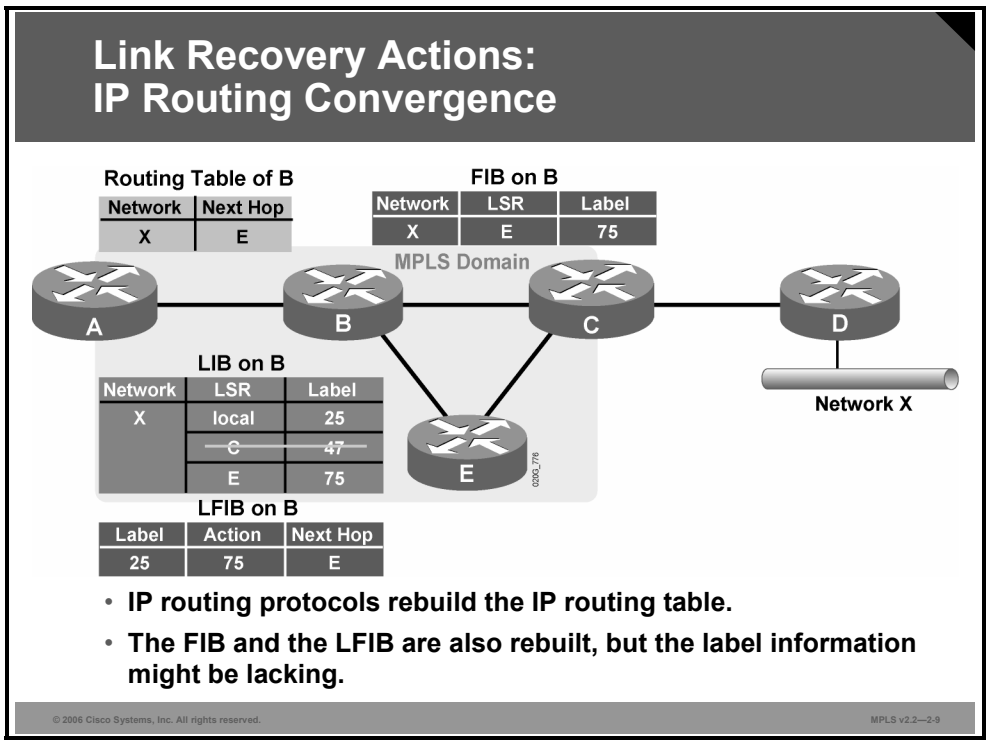

The IGP determines that the link is available again and changes the next-hop address for network X to point to router C. However, when router B also tries to set the next-hop label for network X, it has to wait for the LDP session between routers B and C to be reestablished.

A pop action is used in the LFIB table on router B while the LDP establishes the session between routers B and C. This process adds to the overall convergence time in an MPLS domain. The downtime for network X is not influenced by LDP convergence because normal IP forwarding is used until the new next-hop label is available.

Note	Although this behavior has no significant effect on traditional IP routing, it can significantly influence MPLS Virtual Private Networks (VPNs). This is because the VPN traffic cannot be forwarded before the LDP session is fully operational.

Link Recovery Actions: MPLS Convergence

- **Routing protocol convergence optimizes the forwarding path after a link recovery.**
- **The LIB might not contain the label from the new next hop by the time the IGP convergence is complete.**
- **End-to-end MPLS connectivity might be intermittently broken after link recovery.**
- **Use MPLS TE for make-before-break recovery.**

Link recovery requires that an LDP session be established (reestablished), which adds to the convergence time of LDP.

Networks may be temporarily unreachable because of the convergence limitations of routing protocols.

Cisco MPLS Traffic Engineering (MPLS TE) can be used to prevent longer downtime when a link fails or is recovering.

Summary

This topic summarizes the key points that were discussed in this lesson.

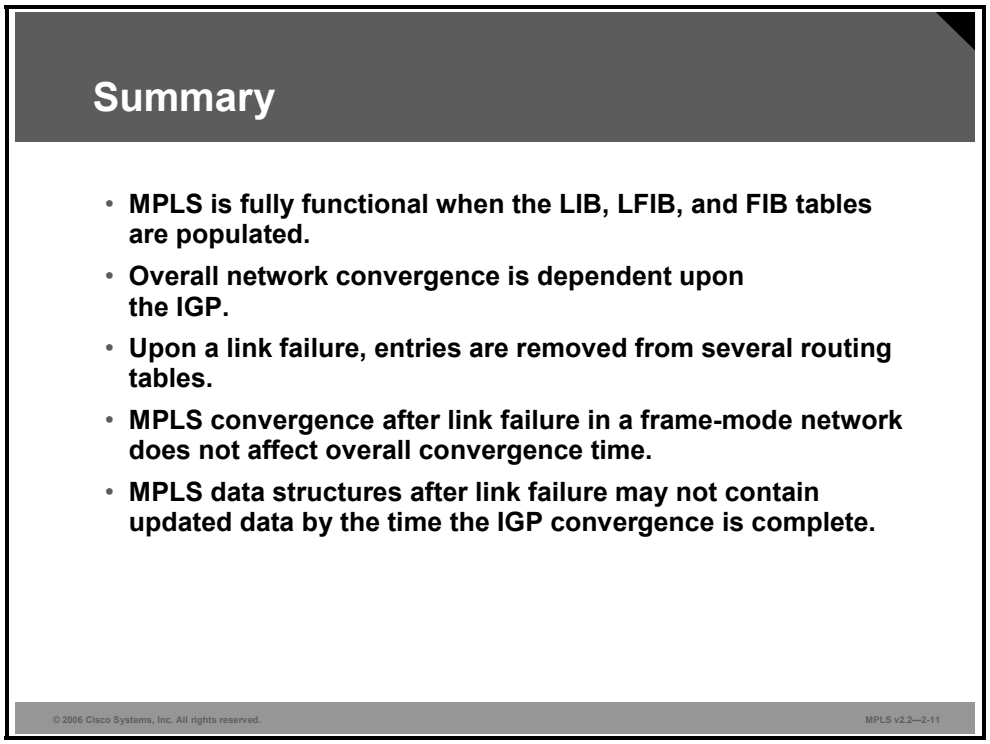

Lesson 4

Introducing MPLS Label Allocation, Distribution, and Retention Modes

Overview

In this lesson, label distribution parameters are discussed. The differences between label distribution parameters are covered, and the default Cisco parameter sets are identified.

There are different modes of operation for Multiprotocol Label Switchingt (MPLS). It is important to have a clear idea of what mode of operation is used under what condition, and if some situations will allow for multiple combinations of these modes.

Objectives

Upon completing this lesson, you will be able to describe the MPLS label allocation, distribution, and retention modes used in Cisco MPLS networks. This ability includes being able to meet these objectives:

- Describe the parameters used in Cisco MPLS label distribution and allocation
- Describe the way in which labels are distributed to neighbors in frame-mode MPLS
- Describe the way in which labels are allocated to neighbors in frame-mode MPLS
- Describe the way in which labels are retained in frame-mode MPLS

Label Distribution Parameters

This topic describes the parameters used in frame-mode MPLS label distribution and allocation.

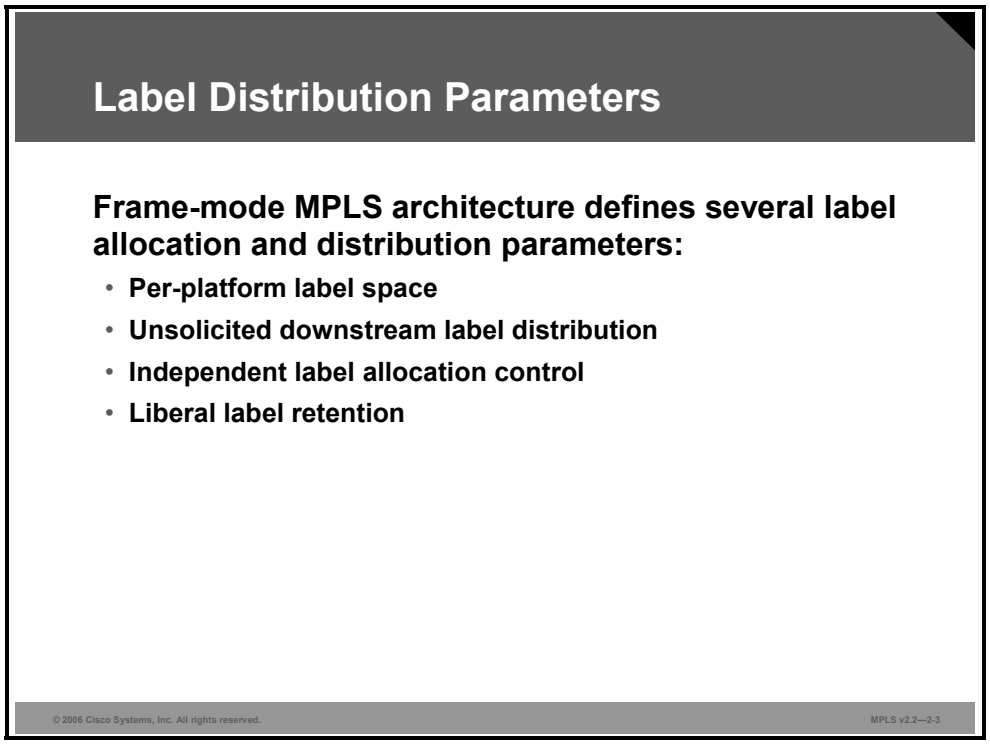

Frame-mode MPLS parameters include:

- Per-platform label space, where labels must be unique for the entire platform (router)

Note	Per-platform label space was discussed in the “Introducing Convergence in Frame-Mode MPLS” lesson.

- Unsolicited downstream distribution of labels, where all routers can asynchronously generate local labels and propagate those labels to adjacent routers
- Independent control mode, where all routers can start propagating labels independently of one another
- Liberal label retention mode, where unused labels are kept (This applies because multiple labels may be received for a prefix, but only one is used.)

Distributing Labels

This topic describes the way in which labels are distributed to neighbors in frame-mode MPLS.

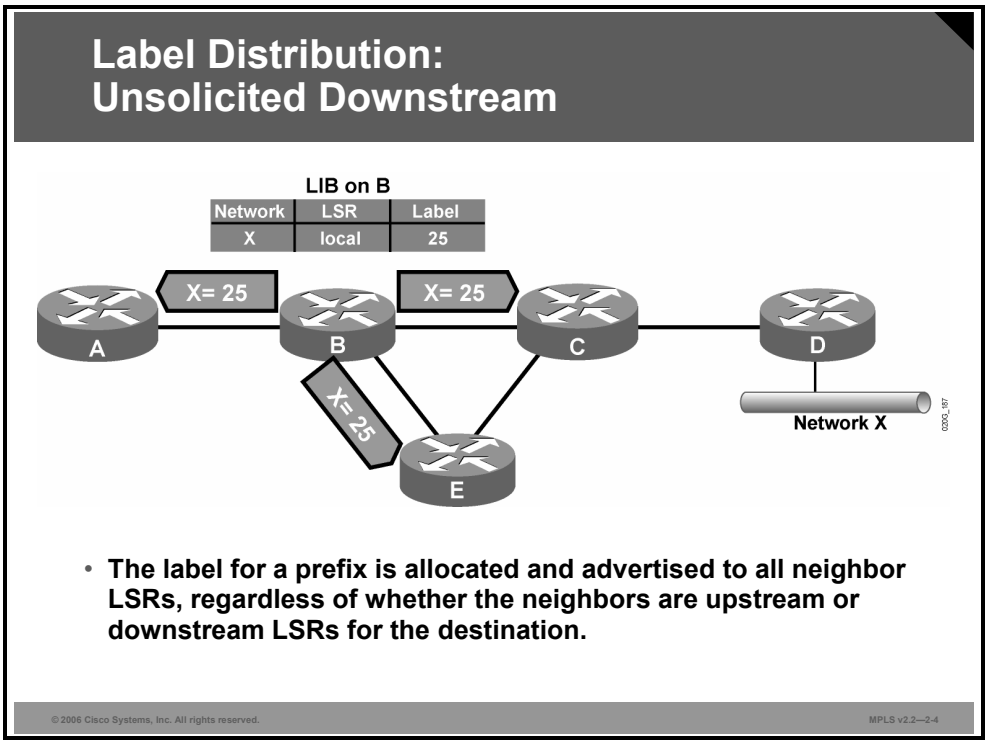

Unsolicited downstream distribution of labels is a method where each router independently assigns a label to each destination IP prefix in its routing table. This mapping is stored in the Label Information Base (LIB) table, which sends it to all Label Distribution Protocol (LDP) peers. There is no control mechanism to govern the propagation of labels in an ordered fashion.

Example: Unsolicited Downstream

The figure illustrates how router B creates a local label (25) and sends that label to all its neighbors. The same action is taken on other routers after the Interior Gateway Protocol (IGP) has put network X into the main routing table.

Each neighbor then decides upon one of these options regarding the label:

- Use the label (if router B is the closest next hop for network X)
- Keep the label in the LIB table
- Ignore the label

Allocating Labels

This topic describes independent control to allocate labels to neighbors in frame-mode MPLS.

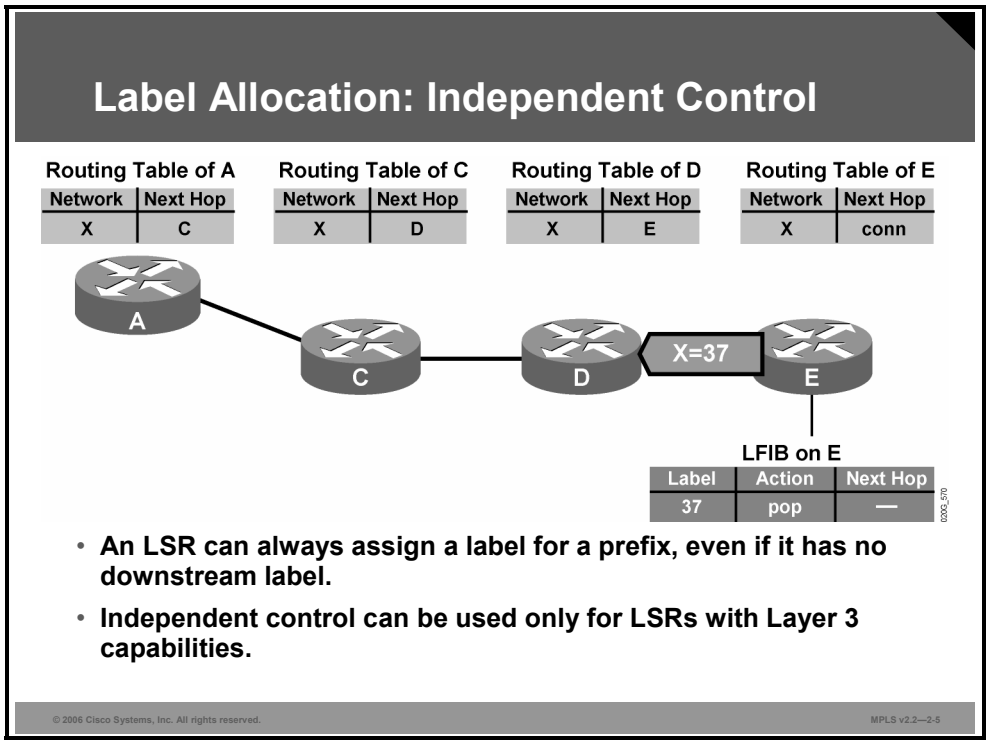

Independent control mode for allocating labels is usually combined with unsolicited downstream propagation of labels, where labels can be created and propagated independently of any other label switch router (LSR). When independent control mode is used, an LSR might be faced with an incoming labeled packet where there is no corresponding outgoing label in the label forwarding information base (LFIB) table. An LSR using independent control mode must therefore be able to perform full Layer 3 lookups. Independent control mode can be used only on LSRs with edge LSR functionality.

Retaining Labels

This topic describes the liberal label retention mode used in frame-mode MPLS.

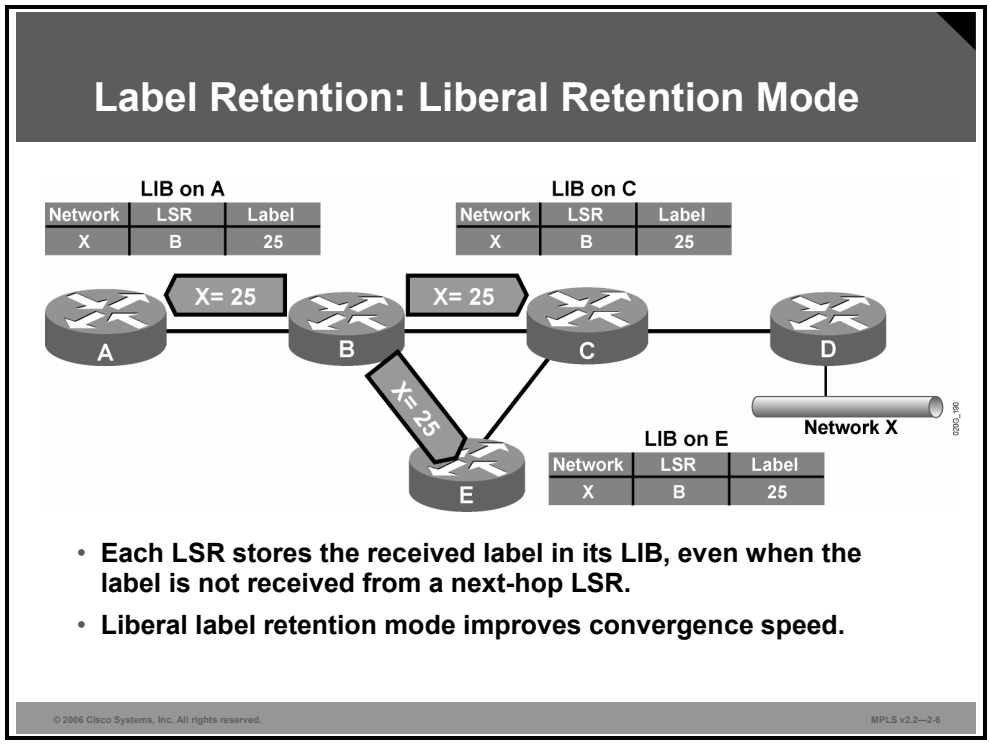

Liberal label retention mode dictates that each LSR keep all labels received from LDP peers, even if they are not the downstream peers for network X.

Example: Liberal Retention Mode

The figure shows how router C receives and keeps the label received from router B for network X, even though router D is the downstream peer.

Summary

This topic summarizes the key points that were discussed in this lesson.

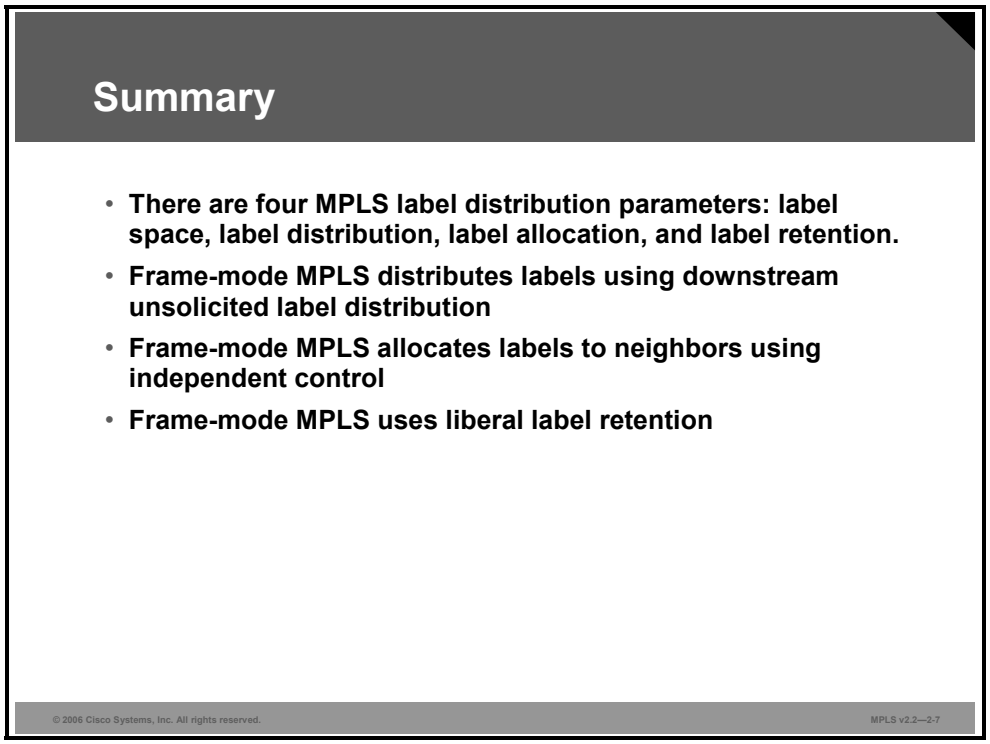

Module Summary

This topic summarizes the key points that were discussed in this module.

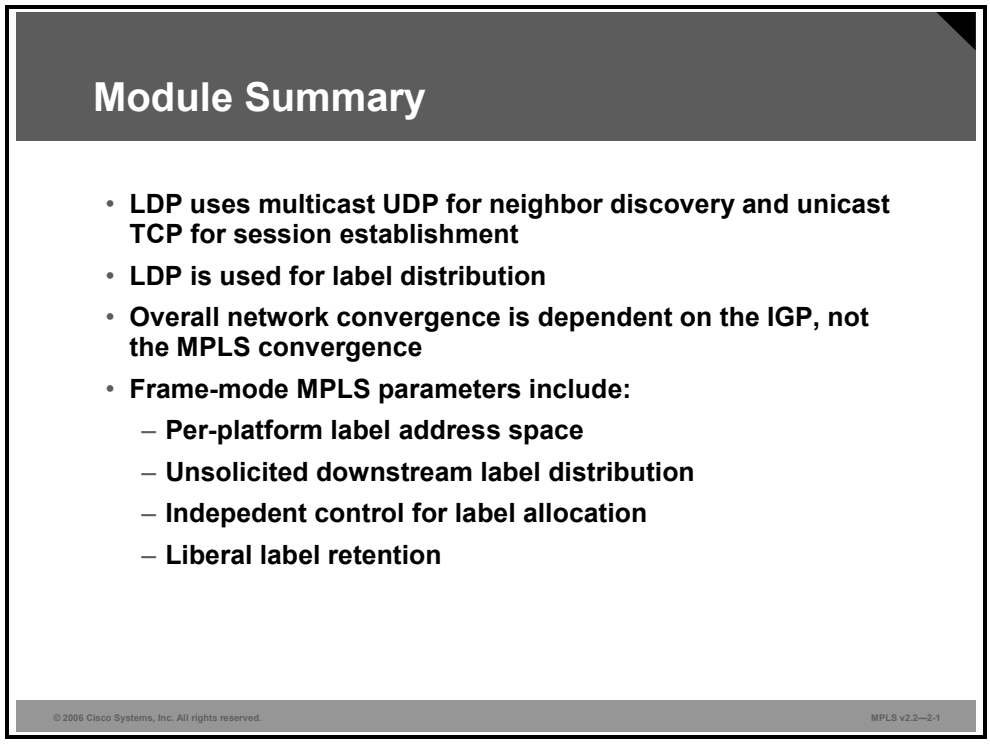

In a Multiprotocol Label Switching (MPLS) network, labels are distributed by Label Distribution Protocol (LDP) after neighbor discovery and session establishment. Label information is populated in Label Information Base (LIB), Forwarding Information Base (FIB), and label forwarding information base (LFIB) tables.

References

For additional information, refer to these resources:

- RFC 3031, *Multiprotocol Label Switching Architecture*
  http://www.ietf.org/rfc/rfc3031.txt
- RFC 3036, *LDP Specification*
  http://www.ietf.org/rfc/rfc3036.txt

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) Which multicast address does LDP use to send hello messages? (Source: Discovering LDP Neighbors)

A) 224.0.0.1
B) 224.0.0.2
C) 224.0.0.12
D) 224.0.20.0

Q2) What does per-platform label space require? (Source: Discovering LDP Neighbors)

A) only one LDP session
B) one session per interface
C) multiple sessions for parallel links
D) “Per-platform” is not a proper term in MPLS terminology.

Q3) What is the purpose of the LDP identifier in a hello message? (Source: Discovering LDP Neighbors)

A) contains the source address
B) contains the multicast address
C) contains the TCP destination port
D) uniquely identifies the neighbor and the label space

Q4) LDP sessions are initiated by using which address? (Source: Discovering LDP Neighbors)

A) the highest IP address
B) the loopback address
C) the lowest IP address
D) whichever LDP neighbor sends the fist hello message

Q5) Exchanging initialization messages is which step in the LDP session negotiation process? (Source: Discovering LDP Neighbors)

A) first step
B) second step
C) third step
D) not required in LDP session negotiation

Q6) LDP discovers nonadjacent neighbors by broadcasting _____ IP addresses. (Source: Discovering LDP Neighbors)

Q7) LDP and TDP use which two well-known port numbers? (Choose two.) (Source: Discovering LDP Neighbors)

A) LDP uses 464.
B) LDP uses 646.
C) LDP uses 711.
D) TDP uses 171.
E) TDP uses 646.
F) TDP uses 711.

Q8) In frame-mode MPLS networks, the number of LDP sessions that are required between neighbors is determined by which of these choices? (Source: Discovering LDP Neighbors)

A) the number of interfaces
B) the number of different label spaces
C) the number of LDP processes running a router
D) the information contained in the source address field of the hello message response

Q9) Which statement best describes PHP? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) PHP works only for TDP and not for LDP.
B) PHP works only for LDP and not for TDP.
C) PHP optimizes MPLS performance.
D) PHP is configurable and by default is disabled.

Q10) Which description applies to per-platform label allocation? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) default operation for frame-mode MPLS
B) an approach that results in larger LIB and LFIB tables
C) an approach that results in slower label exchange
D) a future enhancement for MPLS

Q11) Which three of the answer choices are contained in the LFIB? (Choose three.) (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) local generated label
B) outgoing label
C) local address
D) next-hop address

Q12) When an IP packet is to be label-switched as it traverses an MPLS network, which table is used to perform the label switching? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) LIB
B) FIB
C) FLIB
D) LFIB

Q13) Which statement is correct? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) An IP forwarding table resides on the data plane; LDP (or TDP) runs on the control plane; and an IP routing table resides on the data plane.
B) An IP forwarding table resides on the data plane; LDP (or TDP) runs on the control plane; and an IP routing table resides on the control plane.
C) An IP forwarding table resides on the control plane; LDP (or TDP) runs on the control plane; and an IP routing table resides on the data plane.
D) An IP forwarding table resides on the control plane; LDP (or TDP) runs on the control plane; and an IP routing table resides on the control plane.

Q14) Which two tables contain label information? (Choose two.) (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) LIB
B) main IP routing table
C) FLIB
D) LFIB

Q15) What generates a label update? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) UDP
B) OSPF
C) EIGRP
D) LDP

Q16) Which two statements are correct? (Choose two.) (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) LSPs are bidirectional.
B) LSPs are unidirectional.
C) LDP advertises labels for the entire LSP.
D) LDP advertises labels only for individual segments in the LSP.

Q17) Which statement is correct regarding TTL propagation being disabled? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

A) The label TTL is copied back into the IP TTL.
B) The IP TTL is copied back into the TTL of the label.
C) The IP TL is not copied back into the TTL of the label.
D) TTL label propagation can not be disabled.

Q18) What enables routers in a frame-mode MPLS network to store all received labels, even if they are not being used? (Source: Introducing Convergence in Frame-Mode MPLS)

A) keep-all-labels mode
B) liberal label max-all mode
C) liberal label retention mode
D) A router in a frame-mode network does not keep all labels; the router keeps only the labels that it will use.

Q19) Which table is NOT used to determine if MPLS is fully functional? (Source: Introducing Convergence in Frame-Mode MPLS)

A) LIB
B) LFIB
C) FIB
D) FLIB

Q20) Upon a link failure, which three tables are updated to reflect the failed link? (Choose three.) (Source: Introducing Convergence in Frame-Mode MPLS)

A) LIB
B) LFIB
C) FIB
D) FLIB

Q21) Which statement best describes how a link failure is handled in an MPLS network? (Source: Introducing Convergence in Frame-Mode MPLS)

A) Overall convergence depends on LDP.
B) Overall convergence depends on the IGP that is used.
C) Upon a link failure, only LDP convergence is affected.
D) Upon a link failure, only the IGP convergence is affected.

Q22) Upon a link recovery, which three tables are updated to reflect the failed link? (Choose three.) (Source: Introducing Convergence in Frame-Mode MPLS)

A) LFIB
B) FLIB
C) FIB
D) LIB

Q23) Which of the following statements best describes convergence in a frame-mode MPLS network after a link failure has occurred and been restored? (Source: Introducing Convergence in Frame-Mode MPLS)

A) MPLS convergence occurs after IGP convergence.
B) MPLS convergence occurs before IGP convergence peer to peer.
C) If a failure occurs with the IGP, MPLS convergence is not affected.
D) If a failure occurs with the IGP, MPLS will not be able to converge after the IGP failure has been corrected unless the MPLS process is bounced.

Q24) Which statement is NOT a label distribution parameter? (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) label space
B) label quality
C) label retention
D) label allocation and distribution

Q25) Frame-mode MPLS uses _____ label space. (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

Q26) Which type of label distribution is used in Cisco frame-mode MPLS networks? (Choose one.) (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) downstream-on-demand
B) unsolicited downstream
C) solicited downstream
D) unsolicited downstream-on-demand

Q27) The mode of label allocation for frame-mode MPLS is _____ control. (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

Q28) What is the label retention mode used in Cisco frame-mode MPLS networks? (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) total
B) light
C) liberal
D) conservative

Q29) Which statement is correct? (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) By default, routers with frame interfaces use per-platform label space.
B) By default, ATM switches use per-platform label space.
C) By default, routers with ATM interfaces use per-platform label space and conservative label retention mode.
D) By default, routers with frame interfaces use conservative label retention mode.

Module Self-Check Answer Key

Q1) B

Q2) A

Q3) D

Q4) A

Q5) B

Q6) unicast

Q7) B, F

Q8) B

Q9) C

Q10) A

Q11) A**,** B, D

Q12) D

Q13) B

Q14) A, D

Q15) D

Q16) B, D

Q17) C

Q18) C

Q19) D

Q20) A, B, C

Q21) B

Q22) A, C, D

Q23) A

Q24) B

Q25) per-platform

Q26) B

Q27) independent

Q28) C

Q29) A

Module 3

Frame-Mode MPLS Implementation on Cisco IOS Platforms

Overview

This module provides a review of switching implementations, focusing on Cisco Express Forwarding (CEF). The module also covers the details of implementing frame-mode Multiprotocol Label Switching (MPLS) on Cisco IOS platforms, giving detailed configuration, monitoring, and debugging guidelines. In addition, this module includes the advanced topics of controlling time-to-live (TTL) propagation and label distribution.

Module Objectives

Upon completing this module, you will be able to describe the tasks and commands necessary to implement MPLS on frame-mode Cisco IOS platforms. This ability includes being able to meet these objectives:

- Explain the features of CEF switching
- Configure frame-mode MPLS on Cisco IOS platforms
- Monitor frame-mode MPLS on Cisco IOS platforms
- Troubleshoot frame-mode MPLS problems on Cisco IOS platforms

Lesson 1

Introducing CEF Switching

Overview

This lesson explains the Cisco IOS platform-switching mechanisms by reviewing standard IP switching and Cisco Express Forwarding (CEF) switching, including configuration and monitoring commands.

It is important to understand what part CEF switching plays in a Multiprotocol Label Switching (MPLS) network. CEF must be running as a prerequisite to implementing MPLS on a Cisco router; therefore, an understanding of the purpose of CEF and how it functions will provide an awareness of how the network uses CEF information when forwarding packets.

Objectives

Upon completing this lesson, you will be able to describe the features of CEF switching. This ability includes being able to meet these objectives:

- Describe the various switching mechanisms used by Cisco IOS platforms
- Describe the function of standard IP switching on Cisco IOS platforms
- Describe the architecture of CEF switching
- Configure IP CEF switching
- Monitor IP CEF switching

What Are Cisco IOS Platform-Switching Mechanisms?

This topic describes the various switching mechanisms used by Cisco IOS platforms.

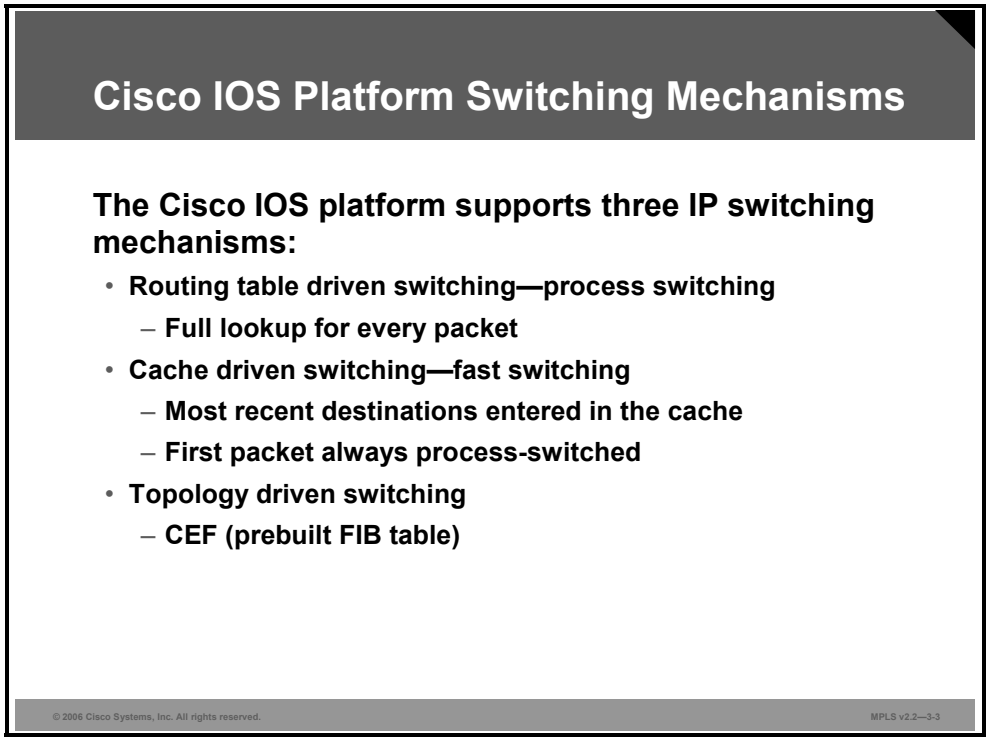

The first and the oldest switching mechanism available in Cisco routers is process switching. Because process switching must find a destination in the routing table (possibly a recursive lookup) and construct a new Layer 2 frame header for every packet, it is very slow and is normally not used.

To overcome the slow performance of process switching, Cisco IOS platforms support several switching mechanisms that use a cache to store the most recently used destinations. The cache uses a faster searching mechanism, and it stores the entire Layer 2 frame header to improve the encapsulation performance. The first packet whose destination is not found in the fast-switching cache is process-switched, and an entry is created in the cache. The subsequent packets are switched in the interrupt code using the cache to improve performance.

The latest and preferred Cisco IOS platform-switching mechanism is CEF, which incorporates the best of the previous switching mechanisms. CEF supports per-packet load balancing (previously supported only by process switching), per-source or per-destination load balancing, fast destination lookup, and many other features not supported by other switching mechanisms.

The CEF cache, or Forwarding Information Base (FIB) table, is essentially a replacement for the standard routing table.

Using Standard IP Switching

This topic describes the function of standard IP switching on Cisco IOS platforms.

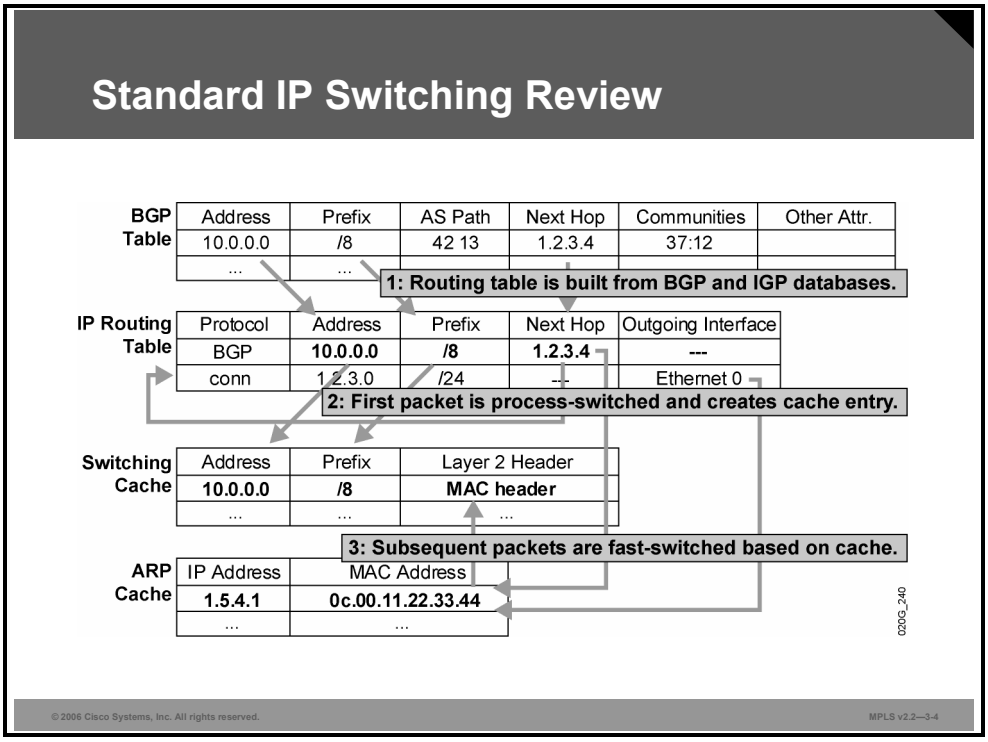

There is a specific sequence of events that occurs when process switching and fast switching are used for destinations learned through Border Gateway Protocol (BGP).

Example: Standard IP Switching

The figure illustrates this process. Here is a description of the sequence of events:

- When a BGP update is received and processed, an entry is created in the routing table.
- When the first packet arrives for this destination, the router tries to find the destination in the fast-switching cache. Because the destination is not in the fast-switching cache, process switching has to switch the packet when the process is run. The process performs a recursive lookup to find the outgoing interface. The process switching may possibly trigger an Address Resolution Protocol (ARP) request or find the Layer 2 address in the ARP cache. Finally, it creates an entry in the fast-switching cache.
- All subsequent packets for the same destination are fast-switched, as follows:
  - The switching occurs in the interrupt code (the packet is processed immediately).
  - Fast destination lookup is performed (no recursion).
  - The encapsulation uses a pregenerated Layer 2 header that contains the destination and Layer 2 source (MAC) address. (No ARP request or ARP cache lookup is necessary.)

Whenever a router receives a packet that should be fast-switched but the destination is not in the switching cache, the packet is process-switched. A full routing table lookup is performed, and an entry in the fast-switching cache is created to ensure that the subsequent packets for the same destination prefix will be fast-switched.

What Is the CEF Switching Architecture?

This topic describes the architecture of CEF switching.

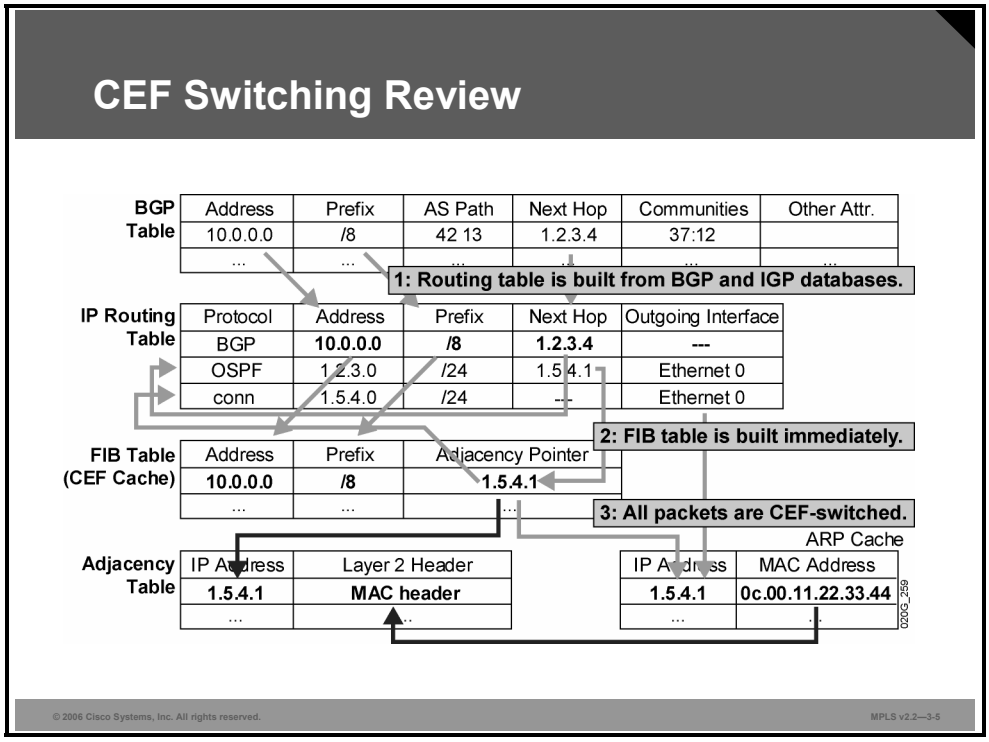

CEF uses a different architecture from process switching or any other cache-based switching mechanism. CEF uses a complete IP switching table, the FIB table, which holds the same information as the IP routing table. The generation of entries in the FIB table is not packet-triggered but change-triggered. When something changes in the IP routing table, the change is also reflected in the FIB table.

Because the FIB table contains the complete IP switching table, the router can make definitive decisions based on the information in it. Whenever a router receives a packet that should be CEF-switched, but the destination is not in the FIB, the packet is dropped.

The FIB table is also different from other fast-switching caches in that it does not contain information about the outgoing interface and the corresponding Layer 2 header. That information is stored in a separate table, the adjacency table. The adjacency table is more or less a copy of the ARP cache, but instead of holding only the destination MAC address, it holds the Layer 2 header.

Note	If the router carries full Internet routing, enabling CEF may consume additional memory. Enabling distributed CEF will also affect memory utilization on Versatile Interface Processor (VIP) modules (Cisco 7500 Series Routers) or line cards (Cisco 12000 Series Internet Routers), because the entire FIB table will be copied to all VIP modules or line cards.

Configuring IP CEF

This topic describes how to configure CEF on Cisco IOS platforms.

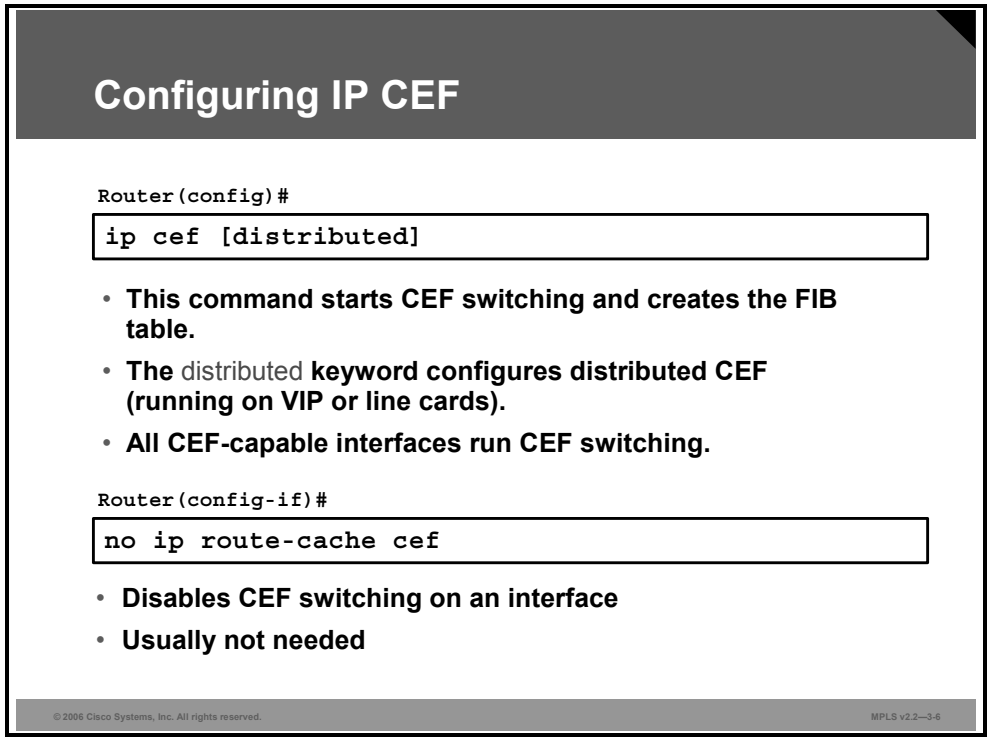

ip cef

To enable CEF on the route processor card, use the **ip cef** global command in global configuration mode. To disable CEF, use the **no** form of this command. Use the appropriate form of the command:

- **ip cef** [**distributed**]
- **no ip cef** [**distributed**]

Syntax Description

distributed (optional): Enables the distributed CEF operation. This option distributes the CEF information to the line cards. The line cards perform express forwarding.

CEF is disabled by default, excluding these platforms:

- CEF is enabled on Cisco 7100 Series Routers.
- CEF is enabled on Cisco 7200 Series Routers.
- CEF is enabled on Cisco 7500 Series Routers.
- CEF is enabled on Cisco 7600 Series Routers, and distributed CEF is enabled on some Cisco 7600 Series Line Cards.
- Distributed CEF is enabled on Cisco 12000 Series Internet Routers.

ip route-cache cef

To enable CEF operation on an interface after the CEF operation has been disabled, use the **ip route-cache cef** command in interface configuration mode. To disable CEF operation on an interface, use the **no** form of this command. Use the form that follows of the two commands:

- **ip route-cache cef**
- **no ip route-cache cef**

Syntax Description

This command has no arguments or keywords.

Defaults

When standard CEF or distributed CEF operations are enabled globally, all interfaces that support CEF are enabled by default.

Monitoring IP CEF

This topic describes how to monitor CEF on Cisco IOS platforms.

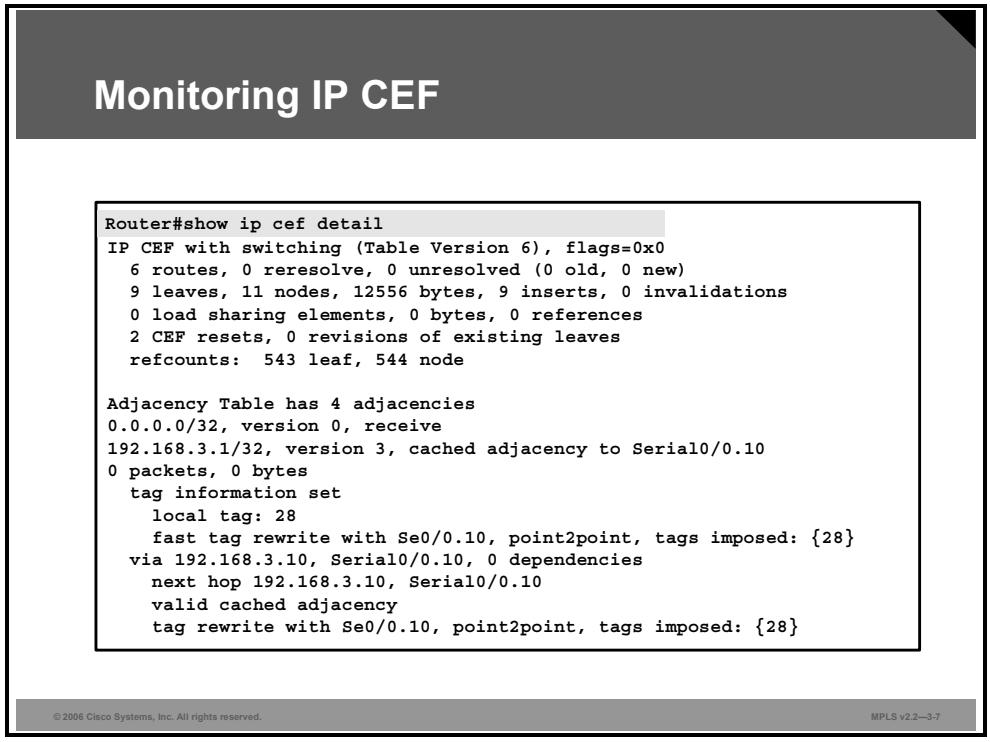

show ip cef

To display unresolved entries in the FIB table or to display a summary of the FIB, use this form of the **show ip cef** EXEC command: **show ip cef** [**unresolved** | **summary**].

To display specific entries in the FIB table based on IP address information, use this form of the **show ip cef** command in EXEC mode: **show ip cef** [*network* [*mask* [**longer-prefix**]]] [**detail**].

To display specific entries in the FIB table based on interface information, use this form of the **show ip cef** command in EXEC mode: **show ip cef** [*type number*] [**detail**].

The table describes the parameters for the **show ip cef** command.

show ip cef Syntax Description

Parameter	Description
unresolved	(Optional) Displays unresolved FIB entries
summary	(Optional) Displays a summary of the FIB
network	(Optional) Displays the FIB entry for the specified destination network
mask	(Optional) Displays the FIB entry for the specified destination network and mask
longer-prefix	(Optional) Displays the FIB entries for all the specific destinations
detail	(Optional) Displays detailed FIB entry information
type number	(Optional) Interface type and number for which to display FIB entries

Summary

This topic summarizes the key points that were discussed in this lesson.

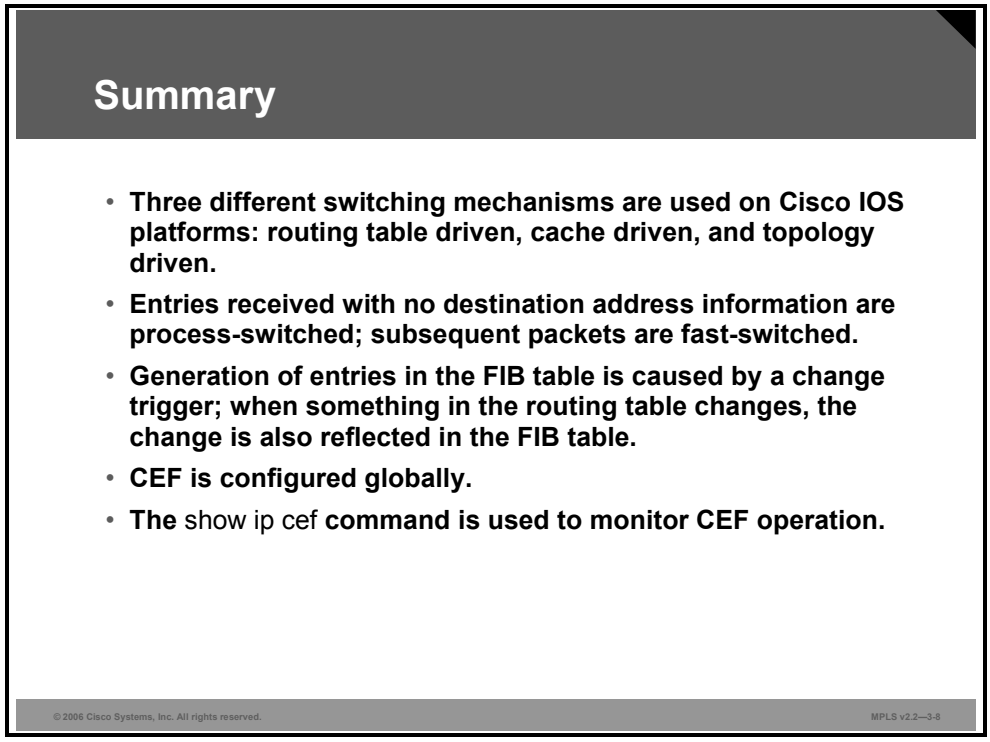

Lesson 2

Configuring Frame-Mode MPLS on Cisco IOS Platforms

Overview

This lesson describes how to configure frame-mode Multiprotocol Label Switching (MPLS) on Cisco IOS platforms. The mandatory configuration tasks, and commands and their correct syntax usage, are discussed in this lesson. The lesson also covers some advanced configurations such as label-switching maximum transmission unit (MTU), IP time-to-live (TTL) propagation, and conditional label distribution. Also discussed in this lesson is the operation of frame-mode MPLS over switched WAN media.

It is important to understand how to enable and configure MPLS to successfully complete the Lab 3-1: Establishing the Core MPLS Environment.

Objectives

Upon completing this lesson, you will be able to describe how to configure frame-mode MPLS on Cisco IOS platforms. This ability includes being able to meet these objectives:

- Describe the MPLS configuration tasks
- Configure the MPLS ID on a router
- Configure MPLS on a frame-mode interface
- Configure a label-switching MTU
- Configure IP TTL propagation
- Configure conditional label distribution
- Configure frame-mode MPLS on switched WAN media

What Are MPLS Configuration Tasks?

This topic describes the MPLS configuration tasks.

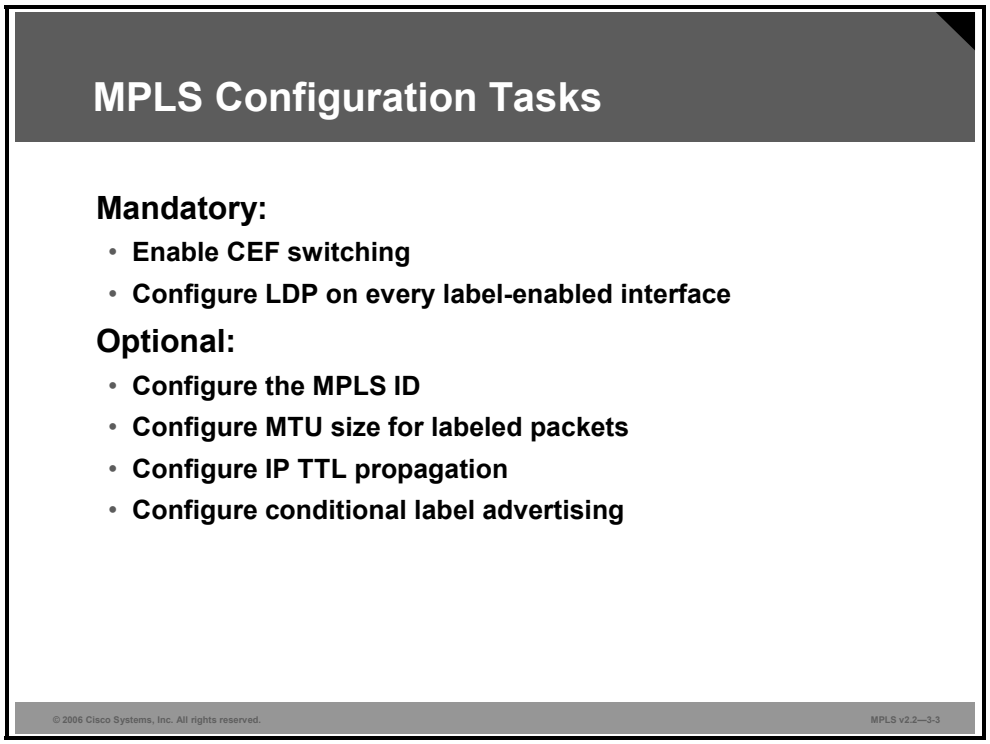

To enable MPLS, you must first enable Cisco Express Forwarding (CEF) switching. Depending on the Cisco IOS software release, you may need to establish the range for the label pool.

You must enable Label Distribution Protocol (LDP) on the interface by using label switching.

Optionally, the maximum size of labeled packets may be changed.

By default, the TTL field is copied from the IP header and placed in the MPLS label when a packet enters an MPLS network. To prevent core routers from responding with (Internet Control Message Protocol [ICMP]) TTL exceeded messages, disable TTL propagation. If TTL propagation is disabled, the value in the TTL field of the label is 255.

Note	Ensure that all routers have TTL propagation either enabled or disabled. If TTL is enabled in some routers and disabled in others, the result may be that a packet leaving the MPLS domain will have a larger TTL value than when it entered.

By default, a router will generate and propagate labels for all networks that it has in the routing table. If label switching is required for only a limited number of networks (for example, only for router loopback addresses), configure conditional label advertising.

Configuring the MPLS ID on a Router

This topic describes how to configure the MPLS identifier (MPLS ID) on a router.

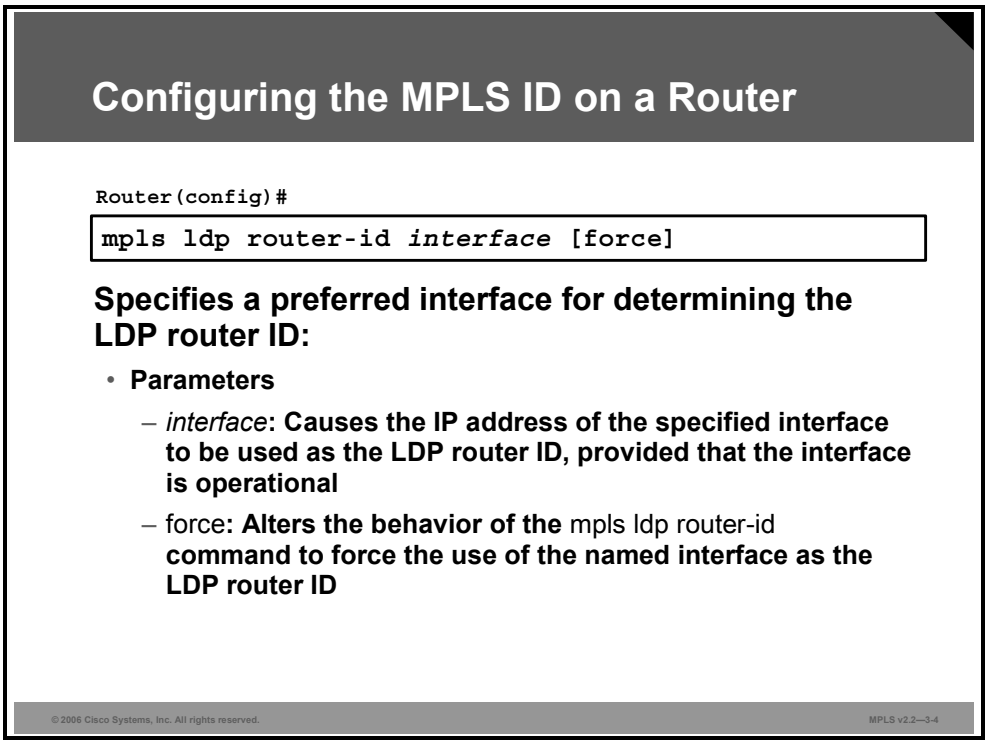

mpls ldp router-id

To specify a preferred interface for determining the LDP router ID, use the **mpls ldp router-id** command in global configuration mode. To remove the preferred interface for determining the LDP router ID, use the **no** form of this command. This illustrates the two commands:

- **mpls ldp router-id** *interface* [**force**]
- **no mpls ldp router-id**

This table describes the parameters for the **mpls idp router-id** command.

mpls idp router-id Syntax Description

Parameter	Description
interface	Causes the IP address of the specified interface to be used as the LDP router ID, provided that the interface is operational
force	(Optional) Alters the behavior of the **mpls ldp router-id** command to force the use of the named interface as the LDP router ID

Defaults

The **mpls ldp router-id** command is disabled.

Configuring MPLS on a Frame-Mode Interface

This topic describes how to configure MPLS on a frame-mode interface.

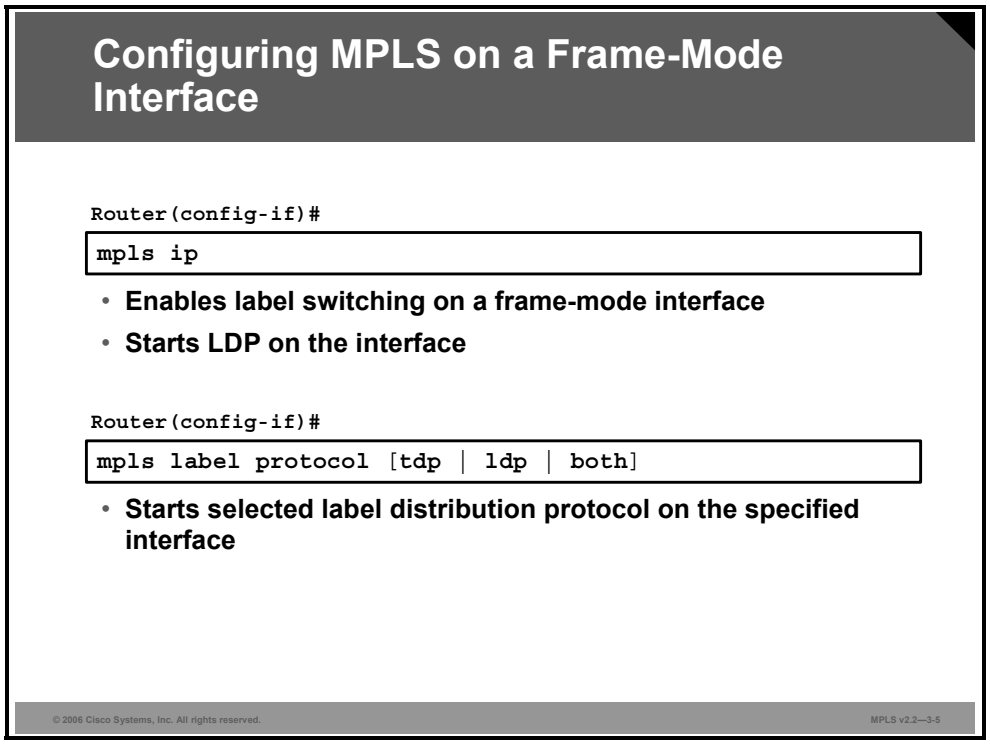

mpls ip

To enable label switching of IP version 4 (IPv4) packets on an interface, use the **mpls ip** command in interface configuration mode. To disable IP label switching on this interface, use the **no** form of this command. This illustrates the two commands:

- **mpls ip**
- **no mpls ip**

Syntax Description

This command has no arguments or keywords.

Defaults

Label switching of IPv4 packets is disabled on this interface.

mpls label protocol [tdp | ldp | both]

To select the label distribution protocol to be used on an interface, use the **mpls label protocol** command in interface configuration mode. To revert to the default label distribution protocol, use the **no** form of this command. This illustrates the two commands:

- **mpls label protocol** *protocol*
- **no mpls label protocol** *protocol*

This table describes the parameters for the **mpls label protocol** [**tdp** | **ldp** | **both**] command.

mpls label protocol [tdp | ldp | both] Syntax Description

Parameter	Description
`tdp`	Enables Tag Distribution Protocol (TDP) on an interface
`ldp`	Enables LDP on an interface
`both`	Enables TDP and LDP on an interface

Defaults

TDP has been the default label distribution protocol. Starting in Cisco IOS Release 12.4(3), the default MPLS label distribution protocol is LDP.

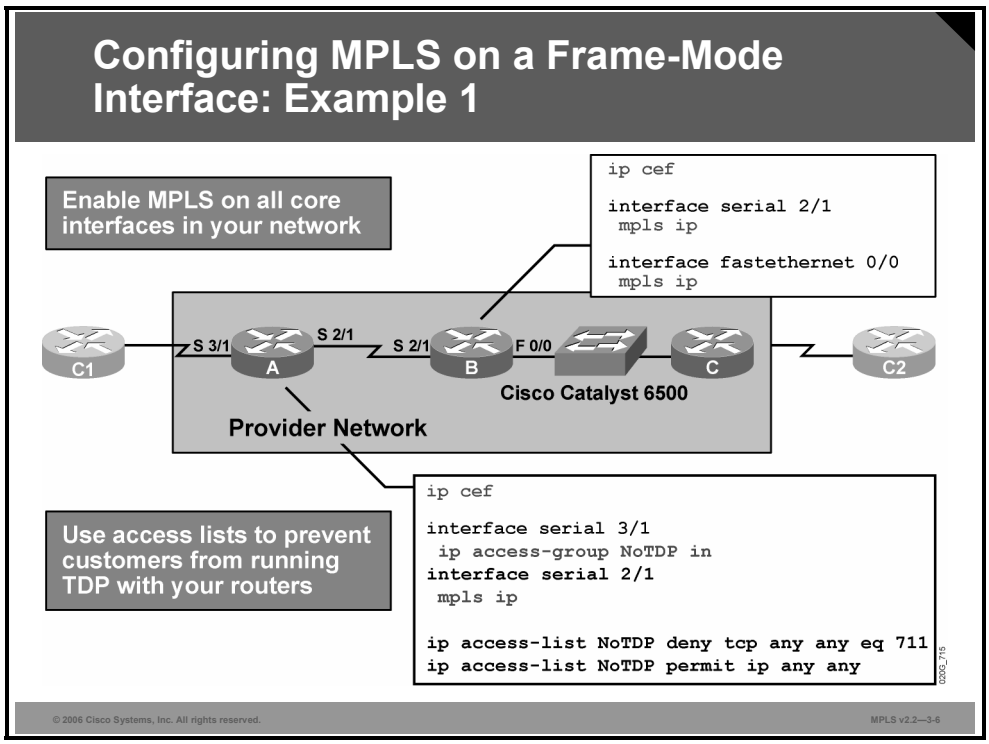

Example: Configuring MPLS on a Frame-Mode Interface

The figure shows the configuration steps needed to enable MPLS on an edge label switch router (LSR). The configuration includes an access control list (ACL) that denies any attempt to establish a TDP session from an interface that is not enabled for MPLS. In the example in the figure, router A has the NoTDP access list applied to serial 3/1, which is not enabled for MPLS.

You must globally enable CEF switching, which automatically enables CEF on all interfaces that support it. (CEF is not supported on logical interfaces, such as loopback interfaces.)

Nonbackbone interfaces have an input ACL that denies TCP sessions on the well-known port number 711 (TDP).

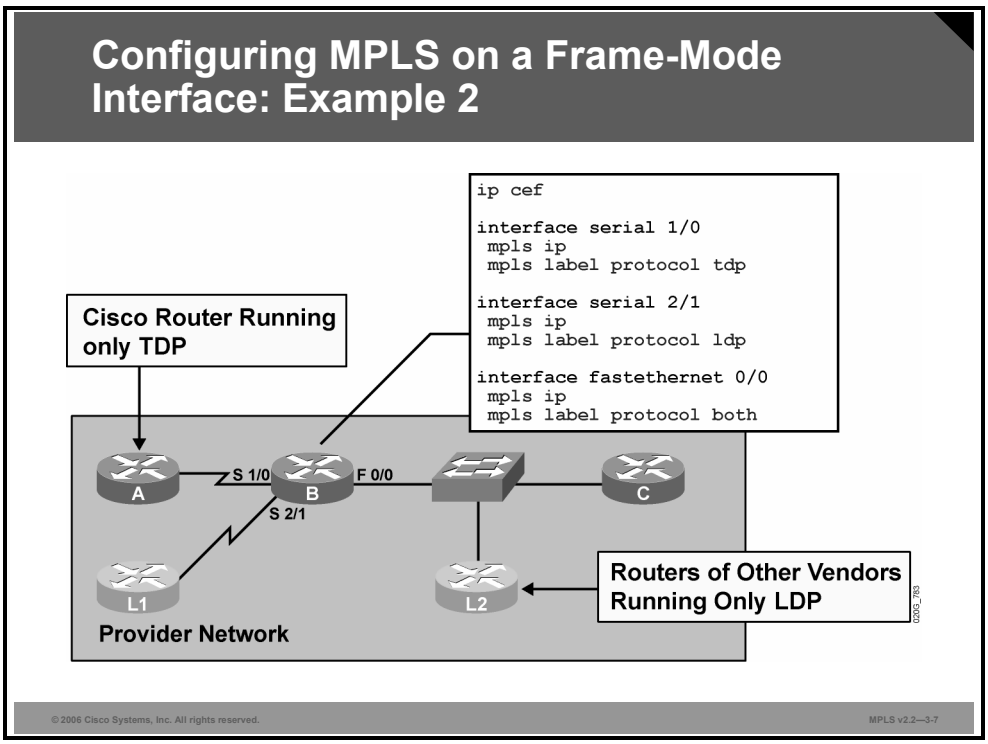

When combining Cisco routers with equipment of other vendors, you may need to use standard LDP (MPLS). TDP (tag switching) can be replaced by LDP on point-to-point interfaces. However, you can also use both protocols on shared media if some devices do not support TDP.

Label switching is more or less independent of the distribution protocol, so there should be no problem in mixing the two protocols. TDP and LDP are functionally very similar, and both populate the Label Information Base (LIB) table.

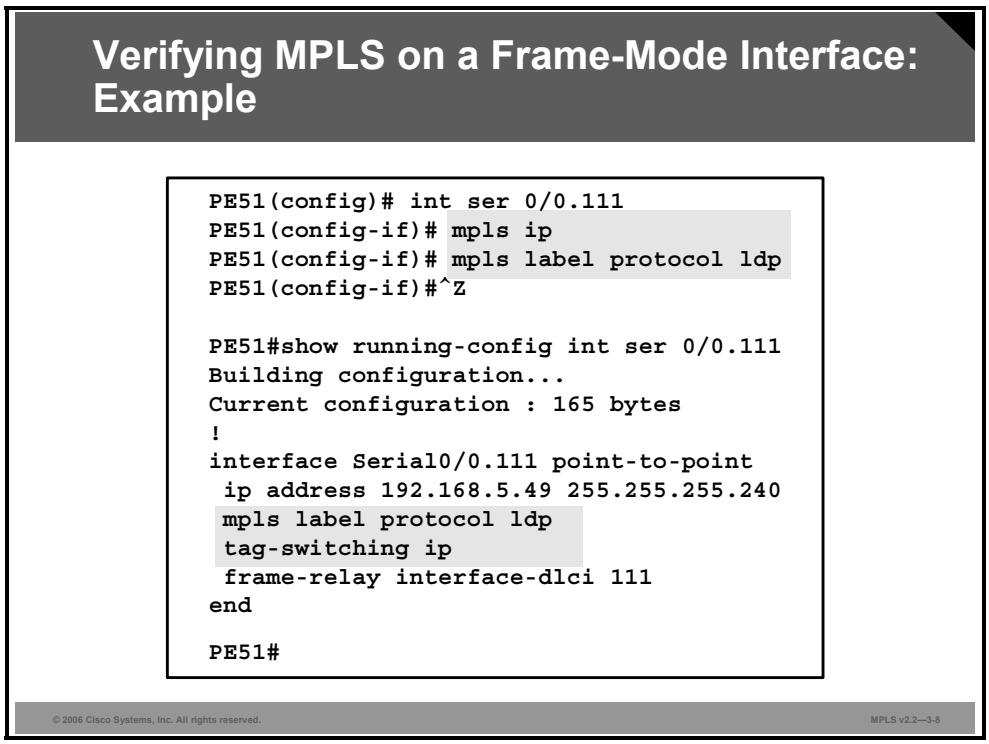

Example: Verifying MPLS on a Frame-Mode Interface

When verifying the MPLS configuration, you will find that depending on the Cisco IOS release, the **show running-config** command will display some of the **ldp** commands as **tag-switching** commands.

Note Starting with Cisco IOS Release 12.4(3), the default MPLS label distribution protocol has changed from TDP to LDP. If no protocol is explicitly configured by the **mpls label protocol** command, LDP is the default label distribution protocol. LDP configuration commands will be saved by using the MPLS form of the command rather than the tag-switching form. Before Cisco IOS Release 12.4(3), commands were saved using the tag-switching form of the command for backward compatibility.

Use caution when upgrading the image on a router that uses TDP. Ensure that the TDP sessions are established when the new image is loaded. You can accomplish this by issuing the global configuration command **mpls label protocol tdp**. Issue this command and save it to the startup configuration before loading the new image. Alternatively, you can enter the command and save the running configuration immediately after loading the new image.

Configuring a Label-Switching MTU

This topic describes how to configure a label-switching MTU.

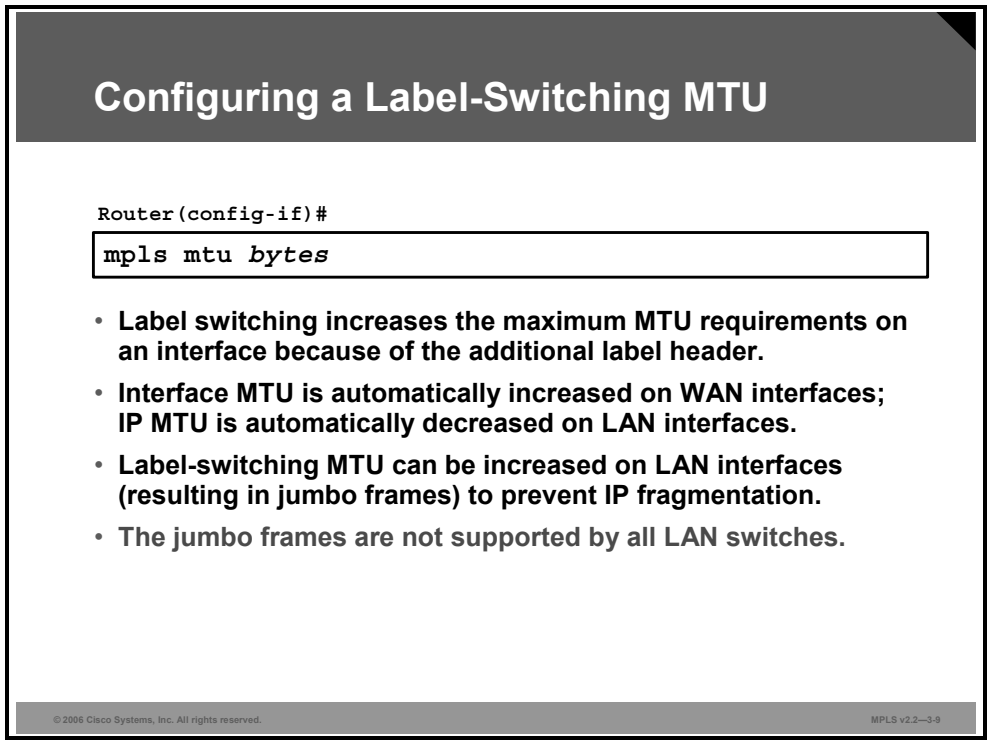

mpls mtu

To set the per-interface MTU for labeled packets, use the **mpls mtu** interface configuration command. This shows these commands:

- **mpls mtu** *bytes*
- **no mpls mtu**

This table describes the parameters for the **mpls mtu** command.

mpls mtu Syntax Description

Parameter	Description
bytes	MTU in bytes

Defaults

The minimum MTU is 64 bytes. The maximum depends on the type of interface medium.

Note	The **show mpls interface** *type number* **detail** command can be used to check the MPLS MTU setting.

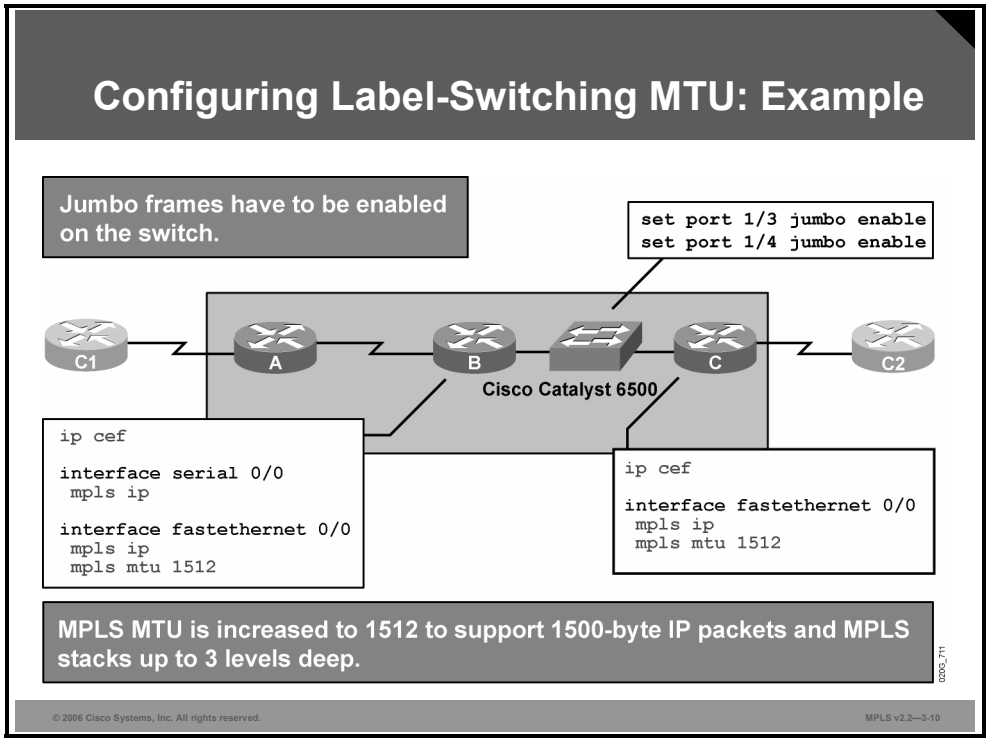

One way of preventing labeled packets from exceeding the maximum size (and being fragmented as a result) is to increase the MTU size of labeled packets for all segments in the label-switched path (LSP) tunnel. The problem will typically occur on LAN switches, where it is more likely that a device does not support oversized packets (also called jumbo frames or, sometimes, giants or baby giants). Some devices support jumbo frames, and some need to be configured to support them.

The MPLS MTU size is increased automatically on WAN interfaces and needs to be increased manually on LAN interfaces.

The MPLS MTU size has to be increased on all LSRs attached to a LAN segment. Additionally, the LAN switches used to implement switched LAN segments need to be configured to support jumbo frames. No additional configuration is necessary for shared LAN segments implemented with hubs.

A different approach is needed if a LAN switch does not support jumbo frames. The problem may be even worse for networks that do not allow ICMP MTU discovery messages to be forwarded to sources of packets and if the Don't Fragment bit (DF bit) is strictly used. This situation can be encountered where firewalls are used.

Configuring IP TTL Propagation

This topic describes how to configure IP TTL propagation.

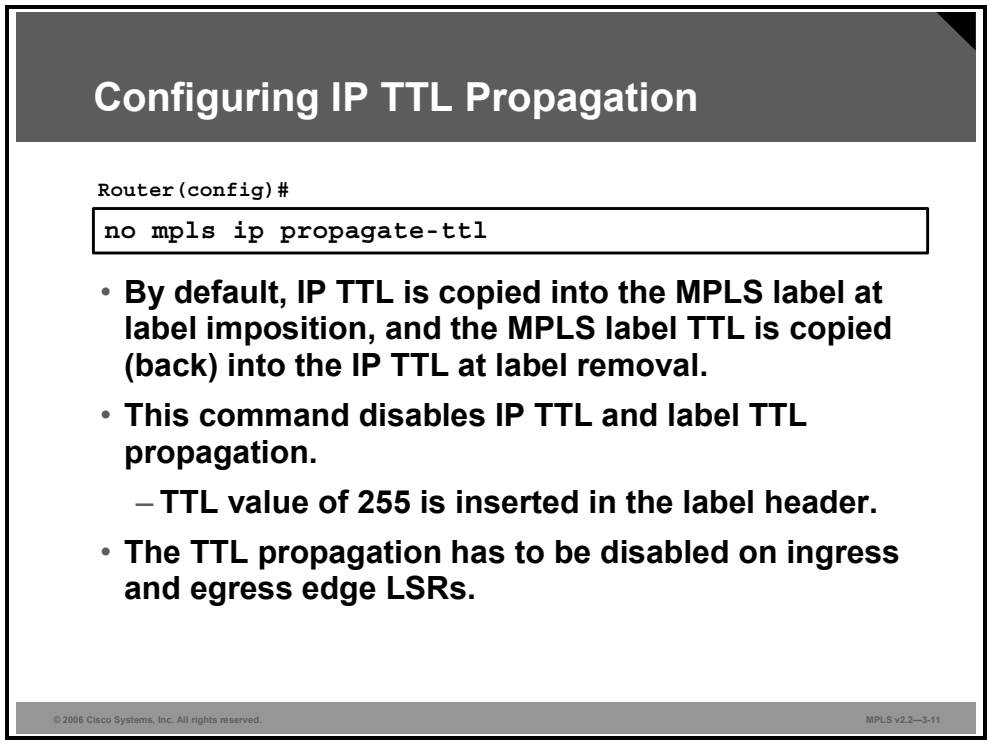

mpls ip propagate-ttl

To set the TTL value on output when the IP packets are being encapsulated in MPLS, use the **mpls ip propagate-ttl** command in privileged EXEC mode. To disable this feature, use the **no** form of this command. This illustrates these two commands:

- **mpls ip propagate-ttl**
- **no mpls ip propagate-ttl**

Syntax Description

This command has no optional keywords or arguments.

Defaults

The MPLS TTL value on packet output is set based on the IP TTL value on packet input.

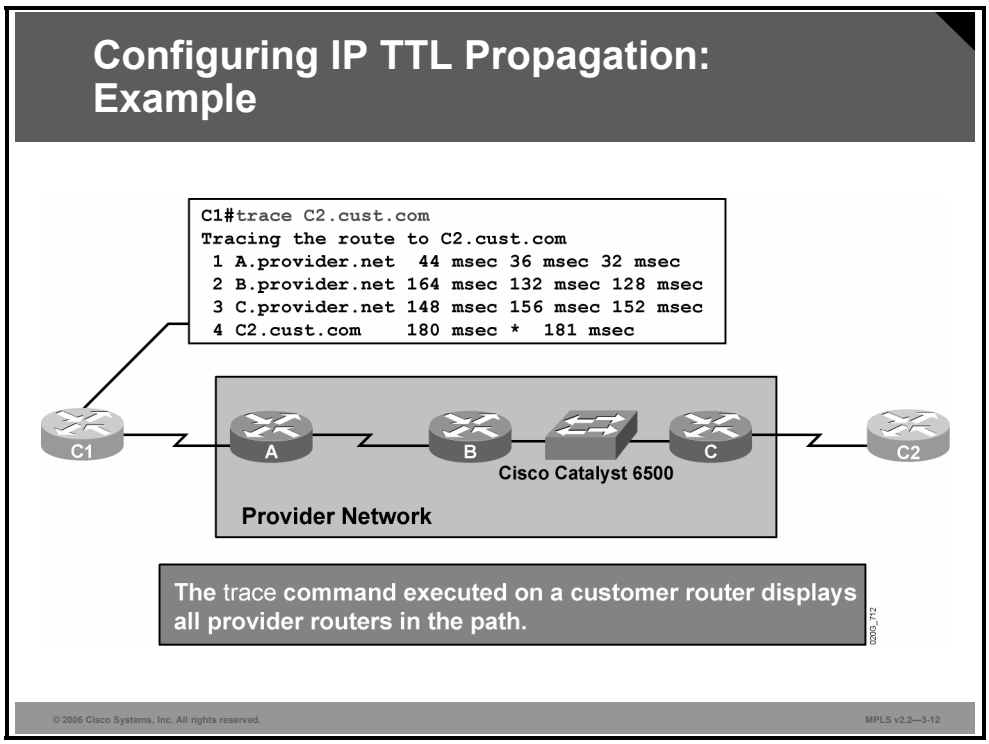

Example: Configuring IP TTL Propagation

The figure illustrates typical traceroute behavior in an MPLS network. Because the label header of a labeled packet carries the TTL value from the original IP packet, the routers in the path can drop packets when the TTL is exceeded. Traceroute will therefore show all the routers in the path. This is the default behavior.

In the example, router C1 is executing a **trace** command that results in this behavior. The steps for this process are as follows:

Step 1 The first packet is an IP packet with TTL=1. Router A decreases the TTL and drops the packet because it reaches 0. An ICMP TTL exceeded message is sent to the source.

Step 2 The second packet sent is an IP packet with TTL=2. Router A decreases the TTL, labels the packet (the TTL from the IP header is copied into the label), and forwards the packet to router B.

Step 3 Router B decreases the TTL value, drops the packet, and sends an ICMP TTL exceeded message to the source.

Step 4 The third packet (TTL=3) experiences a similar processing to the previous packets, except that router C is not the one dropping the packet based on the TTL in the IP header. Router B, because of penultimate hop popping (PHP), previously removed the label, and the TTL was copied back to the IP header (or second label).

The fourth packet (TTL=4) reaches the final destination, where the TTL of the IP packet is examined.

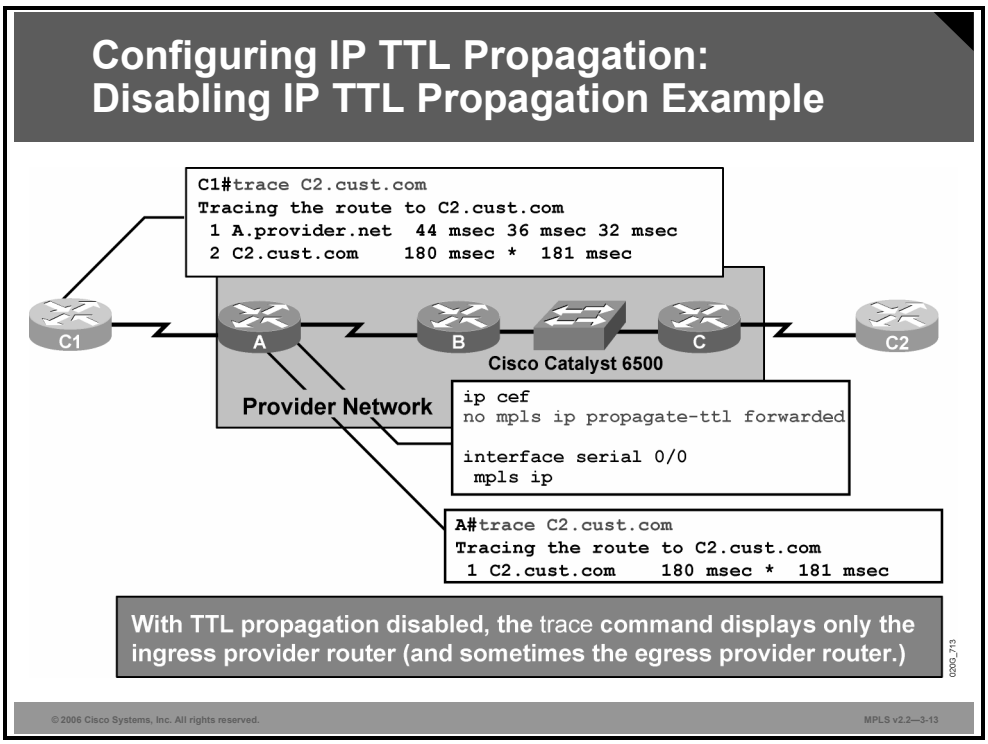

If TTL propagation is disabled, the TTL value is not copied into the label header. Instead, the label TTL field is set to 255. The probable result is that the TTL field in the label header will not decrease to 0 for any router inside the MPLS domain (unless there is a forwarding loop inside the MPLS network).

If the **traceroute** command is used, ICMP replies are received only from those routers that see the real TTL stored in the IP header.

Example: Disabling IP TTL Propagation

In the figure, router C1 is executing the **traceroute** command, but the core routers do not copy the TTL to and from the label. This situation results in this behavior:

Step 1 The first packet is an IP packet with TTL=1. Router A decreases the TTL, drops the packet, and sends an ICMP TTL exceeded message to the source.

Step 2 The second packet is an IP packet with TTL=2. Router A decreases the TTL, labels the packet, and sets the TTL to 255.

Step 3 Router B decreases the TTL in the label to 254 and forwards a labeled packet with TTL set to 254.

Step 4 Router C removes the label, decreases the IP TTL, and sends the packet to the next-hop router (C2). The packet has reached the final destination.

Note The egress MPLS router may, in some cases, be seen in the trace printout, for example, if the route toward C2 is carried in Border Gateway Protocol (BGP), not in the Interior Gateway Protocol (IGP).

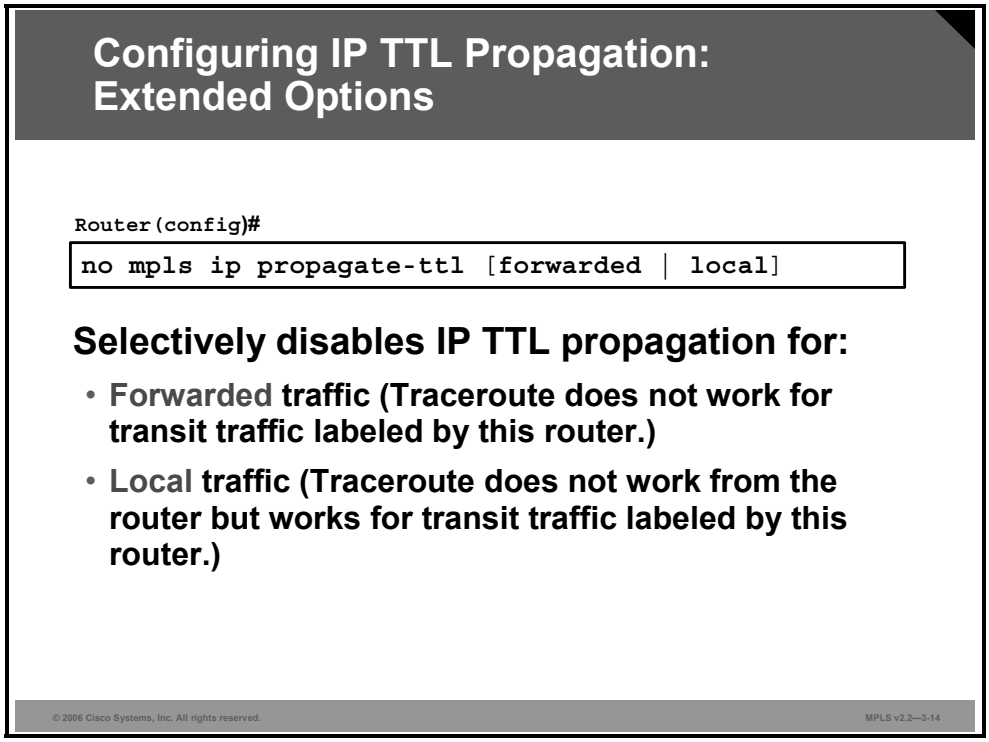

mpls ip propagate-ttl

Use the **mpls ip propagate-ttl** global configuration command to control generation of the TTL field in the label when the label is first added to the IP packet. By default, this command is enabled, which means that the TTL field is copied from the IP header and inserted into the MPLS label. This aspect allows a **trace** command to show all of the hops in the network.

To use a fixed TTL value (255) for the first label of the IP packet, use the **no** form of the **mpls ip propagate-ttl** command. This action hides the structure of the MPLS network from a **trace** command. Specify the types of packets to be hidden by using the **forwarded** and **local** arguments. Specifying **no mpls ip propagate-ttl forwarded** allows the structure of the MPLS network to be hidden from customers but not from the provider. Here are the most common applications of this command:

- **mpls ip propagate-ttl** [**forwarded** | **local**]
- **no mpls ip propagate-ttl** [**forwarded** | **local**]

This table describes the parameters for the **mpls ip propagate-ttl** command.

mpls ip propagate-ttl Syntax Description

Parameter	Description
`forwarded`	(Optional) Hides the structure of the MPLS network from a **trace** command only for forwarded packets; prevents the **trace** command from showing the hops for forwarded packets
`local`	(Optional) Hides the structure of the MPLS network from a **trace** command only for local packets; prevents the **trace** command from showing the hops only for local packets

Defaults

By default, this command is enabled. The TTL field is copied from the IP header. A **trace** command shows all of the hops in the network.

Usage Guidelines

By default, the **mpls ip propagate-ttl** command is enabled, and the IP TTL value is copied to the MPLS TTL field during label imposition. To disable TTL propagation for all packets, use the **no mpls ip propagate-ttl** command. To disable TTL propagation only for forwarded packets, use the **no mpls ip propagate-ttl forwarded** command. This action allows the structure of the MPLS network to be hidden from customers, but not from the provider.

This feature supports the Internet Engineering Task Force (IETF) document "ICMP Extensions for Multiprotocol Label Switching."

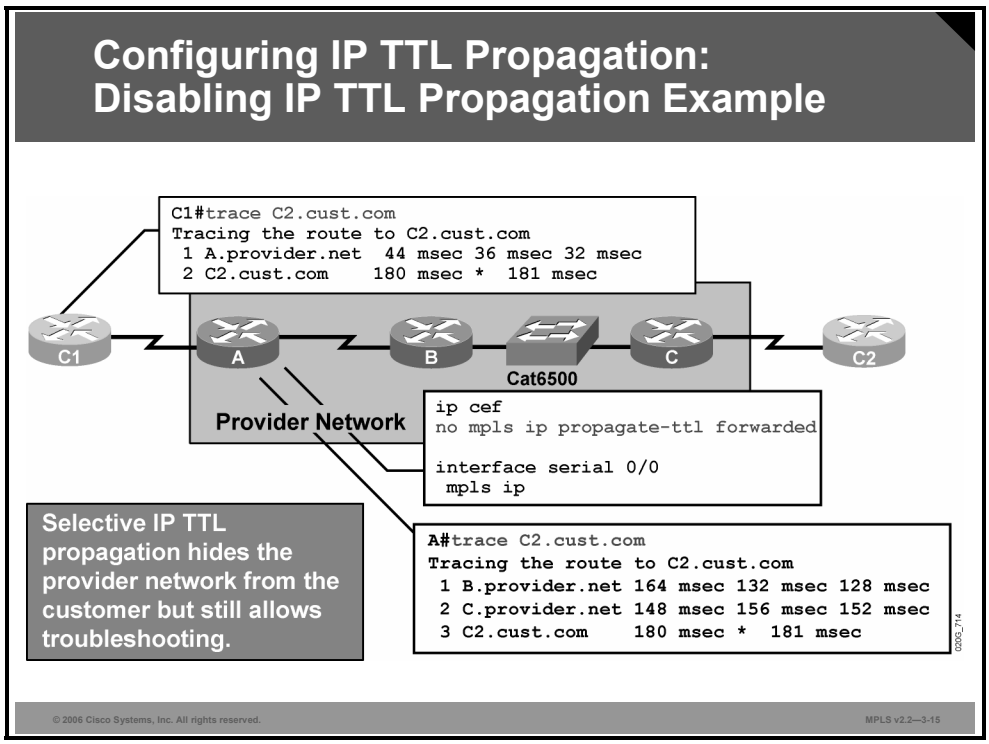

Typically, a service provider likes to hide the backbone network from outside users but allow inside traceroute to work for easier troubleshooting of the network.

This goal can be achieved by disabling TTL propagation for forwarded packets only, as described here:

- If a packet originates in the router, the real TTL value is copied into the label TTL.
- If the packet is received through an interface, the TTL field in a label is assigned a value of 255.

The result is that someone using traceroute on a provider router will see all of the backbone routers. Customers will see only edge routers.

The opposite behavior can be achieved by using the **no mpls ip propagate-ttl local** command, although this is not usually desired.

Configuring Conditional Label Distribution

This topic describes how to configure conditional label distribution.

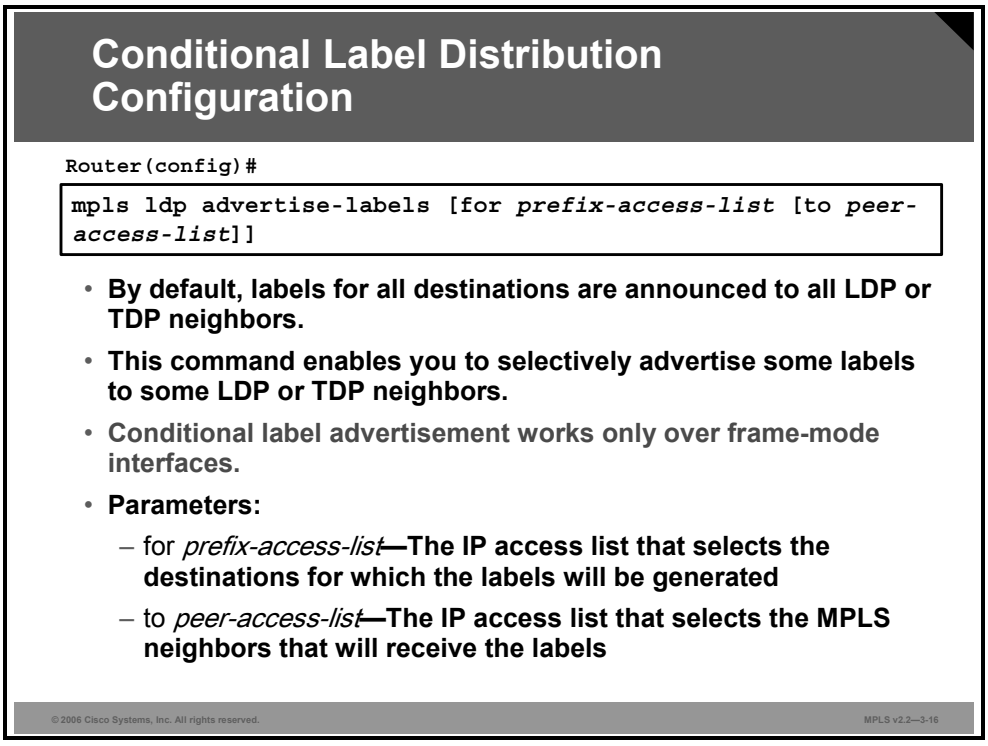

mpls ldp advertise-labels

To control the distribution of locally assigned (incoming) labels by means of LDP, use the **mpls ldp advertise-labels** command in global configuration mode. This command is used to control which labels are advertised to which LDP neighbors. To prevent the distribution of locally assigned labels, use the **no** form of this command, as shown here:

- **mpls ldp advertise-labels** [**for** *prefix-access-list* [**to** *peer-access-list*]]
- **no mpls ldp advertise-labels** [**for** *prefix-access-list* [**to** *peer-access-list*]]

This table describes the parameters for the **mpls idp advertise-labels** command.

mpls idp advertise-labels Syntax Description

Parameter	Description
for *prefix-access-list*	(Optional) This parameter specifies which destinations should have their labels advertised.
to *peer-access-list*	(Optional) This parameter specifies which LSR neighbors should receive label advertisements. An LSR is identified by its router ID, which consists of the first 4 bytes of its 6-byte LDP identifier.

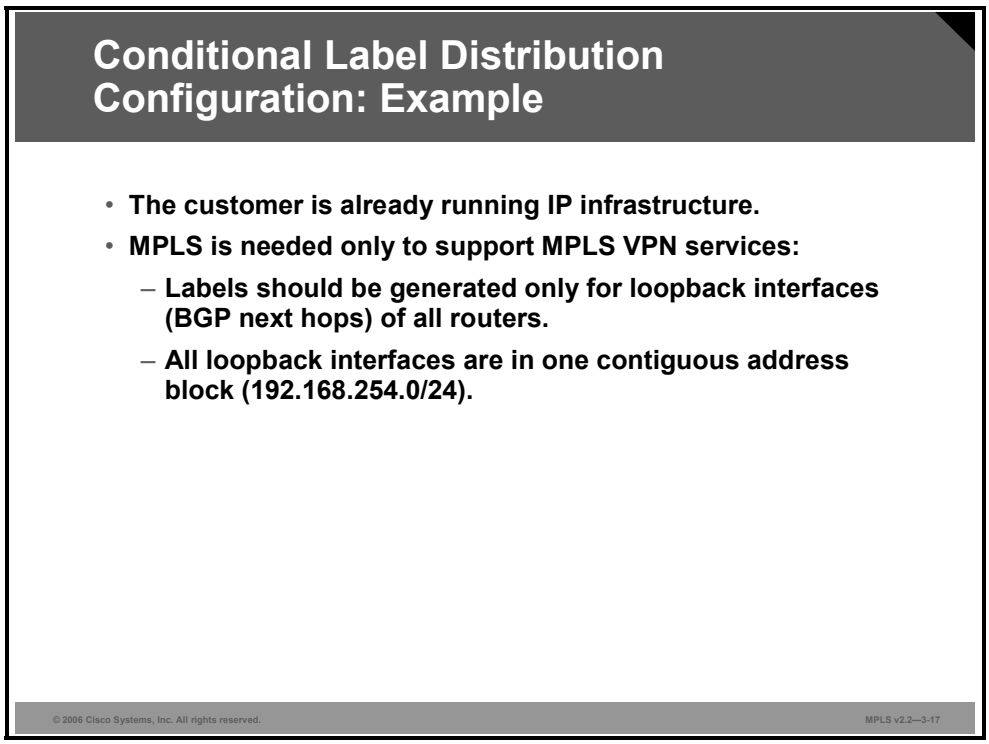

Example: Conditional Label Distribution Configuration

The example here describes where conditional label advertising can be used. The existing network still performs normal IP routing, but the MPLS LSP tunnel between the loopback interfaces of the LSR routers is needed to enable MPLS Virtual Private Network (VPN) functionality.

Using one contiguous block of IP addresses for loopbacks on provider edge (PE) routers can simplify the configuration of conditional advertising.

Conditional Label Distribution Configuration Steps

Step 1: Enable CEF and label switching.

```
ip cef
!
interface serial 0/0
  mpls ip
!
interface serial 0/1
  mpls ip
!
interface ethernet 1/0
  mpls ip
```


In the first step, CEF switching and MPLS have to be enabled on all core interfaces. The MPLS MTU size may be adjusted on the LAN interfaces.

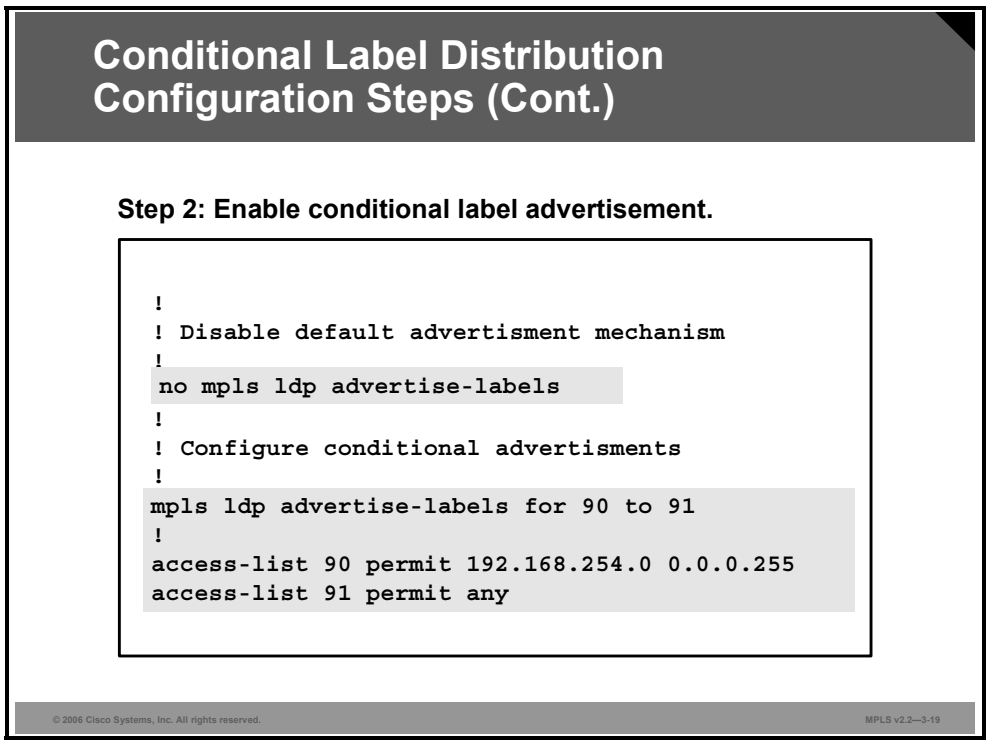

In the second step, disable label propagation and enable conditional label advertising. Within the **mpls ldp advertise-labels** command, specify the neighbors to which the labels are to be sent and the networks for which the labels are to be advertised.

Example: Enabling Conditional Label Advertisement

In the figure, the labels for all networks permitted by ACL 90 are sent to all neighbors matched by ACL 91 (in this example, this would be all TDP or LDP neighbors).

Configuring Frame-Mode MPLS on Switched WAN Media

This topic describes how to configure frame-mode MPLS on switched WAN media.

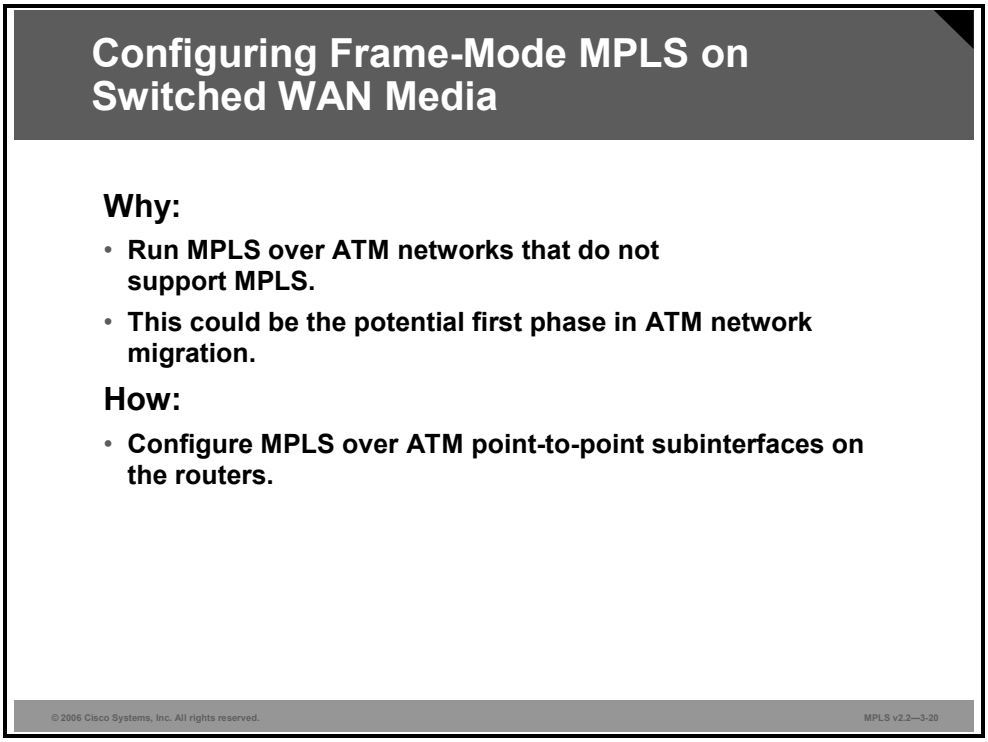

When an underlying ATM infrastructure that does not support cell-mode MPLS is used, MPLS can still be used across point-to-point permanent virtual circuits (PVCs). The MPLS configuration is equal to that on any other Layer 2 media.

This activity could be the first phase of an ATM network migration.

Configuring Frame-Mode MPLS on Switched WAN Media: MPLS over ATM Forum PVCs

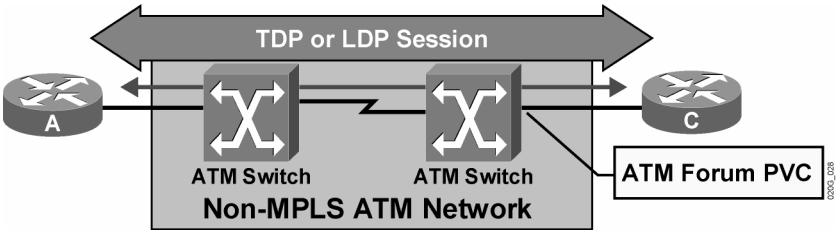

- **Routers view the ATM PVC as a frame-mode MPLS interface.**
- **TDP or LDP is run between the adjacent routers.**
- **Many LSPs can be established over one ATM PVC.**
- **The ATM network is not aware of MPLS between the routers.**

If frame-mode MPLS on an ATM interface is enabled, TDP or LDP neighbor relationships are established between the two PVC endpoint routers and not with the attached ATM switch.

Labeling of packets happens at the process level (in software), while segmentation and reassembly happen on the interface (in hardware), regardless of the type of packet.

Switching is performed based on the virtual path identifier/virtual channel identifier (VPI/VCI) value in the ATM header that is used for this particular PVC, and is not related to Layer 3 IP information.

Configuring Frame-Mode MPLS on Switched WAN Media: MPLS over ATM Forum PVCs (Cont.)

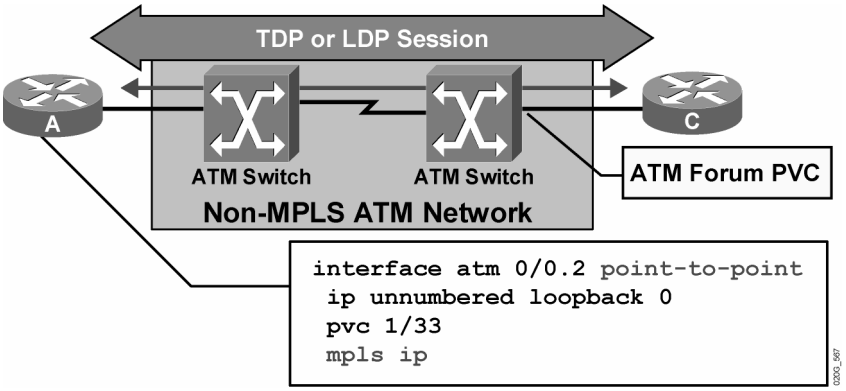

- **Create a point-to-point ATM subinterface.**
- **Configure ATM PVC on the subinterface.**
- **Start label switching and LDP or TDP on the interface.**

Configuring frame-mode MPLS on an ATM interface involves using the same interface command **(mpls ip**). Because this implementation is frame-mode MPLS (versus cell-mode) over ATM, the interface is defined as a point-to-point subinterface.

The ATM parameters are not related to MPLS, because the labeled traffic is using a standard ATM Forum point-to-point PVC.

Configuring Frame-Mode MPLS on Switched WAN Media: MPLS over Frame Relay Networks

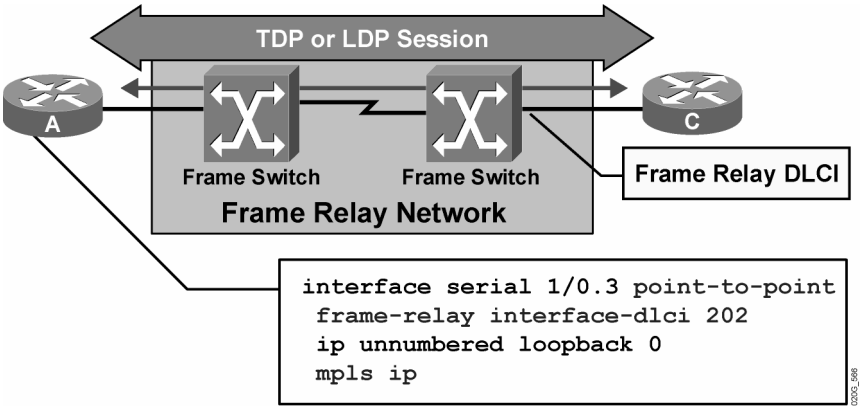

- **Create a point-to-point or multipoint Frame Relay subinterface.**
- **Configure Frame Relay DLCI on the subinterface.**
- **Start label switching and LDP or TDP on the interface.**

Enabling MPLS on a Frame Relay PVC, also called a data-link connection identifier (DLCI), is no different from doing so on any other point-to-point media.

Routers insert a label between the frame and the IP header. The TDP or LDP session is established between the two IP endpoints connected through a Frame Relay network.

Summary

This topic summarizes the key points that were discussed in this lesson.

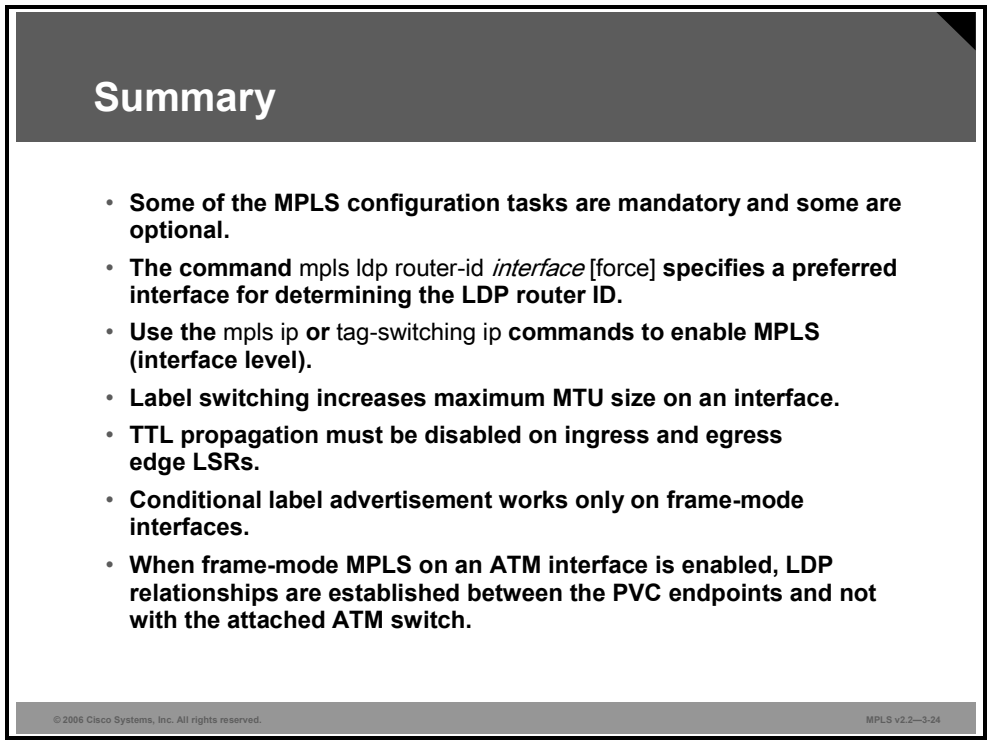

Lesson 3

Monitoring Frame-Mode MPLS on Cisco IOS Platforms

Overview

This lesson describes the procedures for monitoring Multiprotocol Label Switching (MPLS) on Cisco IOS platforms by listing the syntax and parameter descriptions; looking at interfaces, neighbor nodes, and Label Information Base (LIB) and label forwarding information base (LFIB) tables; and outlining the usage guidelines for the commands.

It is very important to know what commands you can use to verify correct operation of MPLS in the network. The information here will help you when you encounter problems with frame-mode interfaces that have MPLS running in the network.

Objectives

Upon completing this lesson, you will be able to describe how to use monitoring commands in frame-mode MPLS on Cisco IOS platforms. This ability includes being able to meet these objectives:

- Monitor MPLS
- Monitor LDP
- Monitor label switching
- Debug MPLS and LDP

Monitoring MPLS

This topic describes how to monitor MPLS.

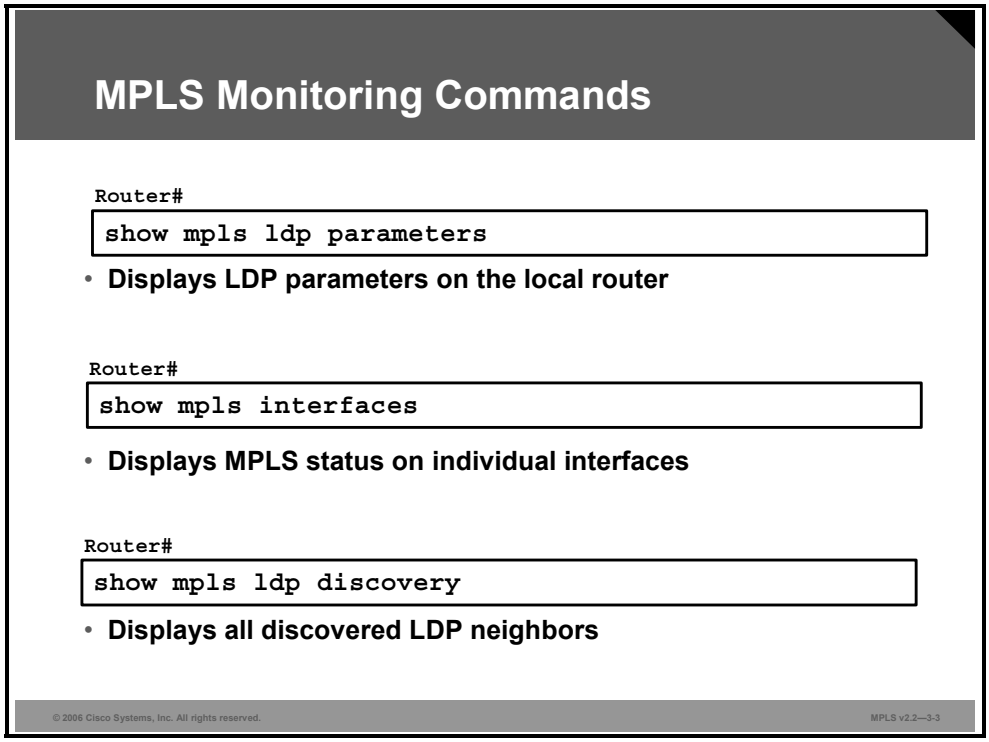

show mpls ldp parameters

To display available Label Distribution Protocol (LDP) parameters, use the **show mpls ldp parameters** command in privileged EXEC mode.

Syntax Description

This command has no arguments or keywords.

show mpls interfaces

To display information about one or more interfaces that have the MPLS feature enabled, use the **show mpls interfaces** [*interface*] [**detail**] command in EXEC mode.

The table describes the parameters for the **show mpls interfaces** command.

show mpls interfaces Syntax Description

Parameter	Description
interface	(Optional) Defines the interface about which to display label-switching information
detail	(Optional) Displays detailed label-switching information for the specified interface

show mpls ldp discovery

To display the status of the LDP discovery process (Hello protocol), use these commands in privileged EXEC mode:

- **show mpls ldp discovery** [**vrf** *vpn-name*]
- **show mpls ldp discovery** [**all**]

The **show mpls ldp discovery** command displays all MPLS-enabled interfaces and the neighbors that are present on the interfaces.

show mpls ldp discovery Syntax Description

Parameter	Description
vrf *vpn-name*	(Optional) Displays the neighbor discovery information for the specified Virtual Private Network (VPN) routing or forwarding instance (vpn-name)
all	(Optional) Displays LDP discovery information for all VPNs when the **all** keyword is specified alone in this command, including those in the default routing domain

MPLS Monitoring Commands: show mpls ldp parameters

```
Router#show mpls ldp parameters
Protocol version: 1
Downstream label pool: min label: 16; max label:
   100000
   [Configured: min label: 1000; max label: 1999]
Session hold time: 180 sec; keep alive interval: 60
   sec
Discovery hello: holdtime: 15 sec; interval: 5 sec
Discovery targeted hello: holdtime: 180 sec; interval:
   5 sec
Downstream on Demand max hop count: 255
LDP for targeted sessions
LDP initial/maximum backoff: 15/120 sec
LDP loop detection: off
```


The table describes the significant fields in the display.

show mpls ldp parameters Field Description

Field	Description
Protocol version	This field indicates the version of LDP running on the platform.
Downstream label pool	This field describes the range of labels available for the platform to assign for label-switching purposes. The available labels range from the smallest value (min label) to the largest label value (max label), with a modest number of labels at the low end of the range (reserved labels), reserved for diagnostic purposes.
Session hold time	This field indicates the time that an LDP session is to be maintained with an LDP peer without receiving LDP traffic or an LDP keepalive message from the peer.
Keepalive interval	This field indicates the interval of time between consecutive transmissions of LDP keepalive messages to an LDP peer.
Discovery hello	This field indicates the amount of time to remember that a neighbor platform wants an LDP session without receiving an LDP hello message from the neighbor (hold time), and the time interval between the transmissions of consecutive LDP hello messages to neighbors (interval).
Discovery targeted hello	This field indicates the amount of time to remember that a neighbor platform wants an LDP session when one of the these situations occurs: ■ The neighbor platform is not directly connected to the router. ■ The neighbor platform has not sent an LDP hello message. This intervening interval is known as hold time. This field also indicates the time interval between the transmissions of consecutive hello messages to a neighbor not directly connected to the router.
LDP for targeted sessions	This field reports the parameters that have been set by the **mpls ldp neighbor targeted** command.
LDP initial/maximum backoff	This field reports the parameters that have been set by the **mpls ldp backoff** command.

MPLS Monitoring Commands: show mpls interfaces

```
Router#show mpls interfaces
Interface Serial0/0:
        IP labeling enabled (ldp)
        LSP Tunnel labeling enabled
        Tag Frame Relay Transport tagging not enabled
        Tagging operational
        Fast Switching Vectors:
          IP to MPLS Fast Switching Vector
          MPLS Turbo Vector
        MTU = 1500
Interface Serial0/3:
        IP labeling enabled (ldp)
        LSP Tunnel labeling not enabled
        Tag Frame Relay Transport tagging not enabled
        Tagging operational
        Fast Switching Vectors:
          IP to MPLS Fast Feature Switching Vector
          MPLS Feature Vector
        MTU = 1500
```


The **show mpls interfaces** command will show only those interfaces on which MPLS has been configured.

The table describes the significant fields in the display.

show mpls interfaces Field Description

Field	Description
Interface	Interface name
IP	“Yes” if IP label switching (sometimes called hop-by-hop label switching) has been enabled on this interface
Tunnel	“Yes” if label-switched path (LSP) tunnel labeling has been enabled on this interface
Tagging operational	Operational state; “Yes” if labeled packets can be sent over this interface Labeled packets can be sent over an interface if an MPLS protocol is configured on the interface and the required Layer 2 negotiations have occurred.

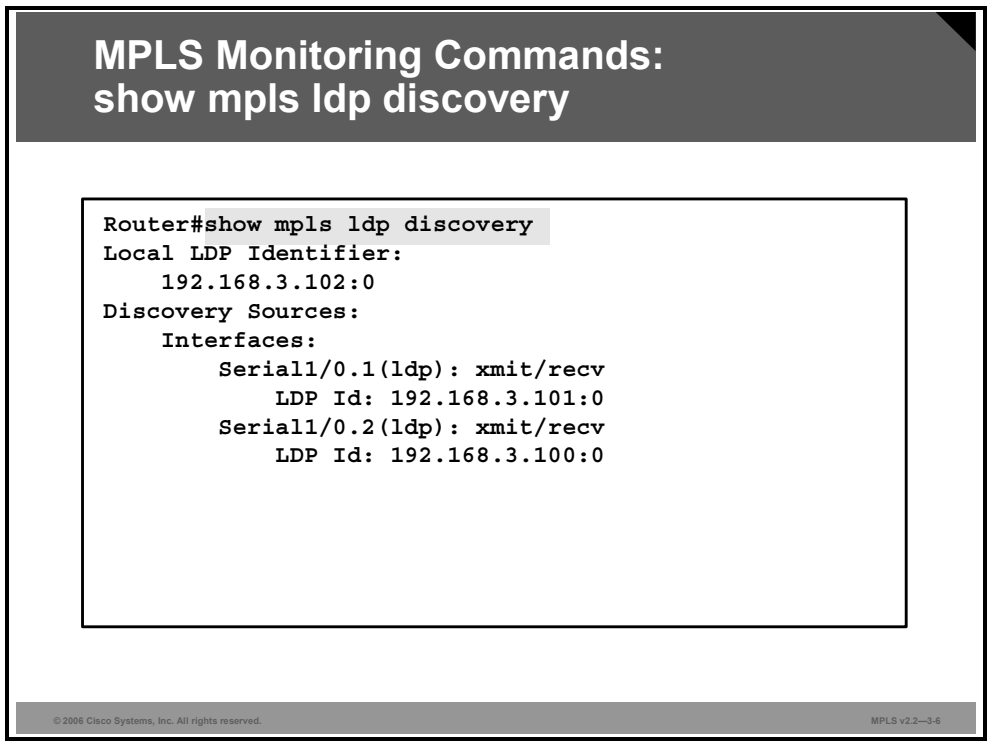

show mpls ldp discovery

The table describes the significant fields in the display.

show mpls ldp discovery Field Description

Field	Description
Local LDP Identifier	This field indicates the LDP identifier (LDP ID) for the local router. An LDP ID is a 6-byte construct displayed in the form “IP address:number”. By convention, the first 4 bytes of the LDP ID constitute the router ID; integers, starting with 0, constitute the final 2 bytes of the IP address:number construct.
Interfaces	This field lists the interfaces that are engaging in LDP discovery activity, as described here: ■ The xmit field indicates that the interface is transmitting LDP discovery hello packets. ■ The recv field indicates that the interface is receiving LDP discovery hello packets. ■ The (ldp) or (tdp) field indicates the label distribution protocol configured for the interface. The LDP (or Tag Distribution Protocol [TDP]) identifiers indicate LDP (or TDP) neighbors discovered on the interface.
Targeted Hellos	This field lists the platforms to which targeted hello messages are being sent, as described here: ■ The xmit, recv, and (ldp) or (tdp) fields are as described for the Interfaces field. ■ The active field indicates that this label switch router (LSR) has initiated targeted hello messages. ■ The passive field indicates that the neighbor LSR has initiated targeted hello messages and that this LSR is configured to respond to the targeted hello messages from the neighbor.

Monitoring LDP

This topic describes how to monitor LDP.

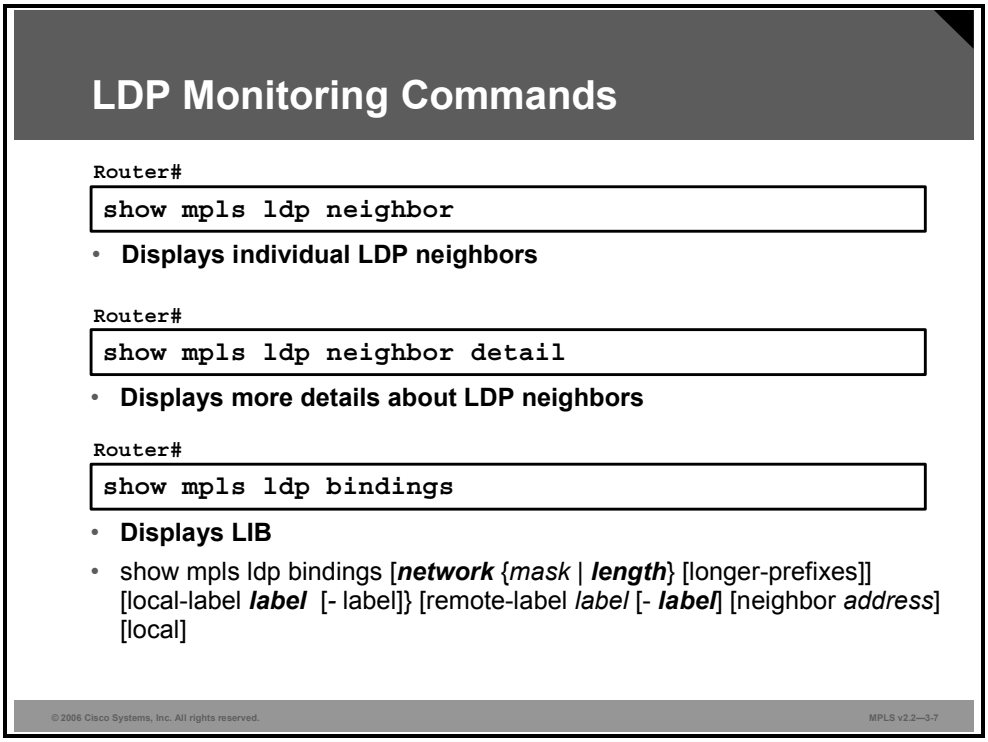

show mpls ldp neighbor

To display the status of LDP sessions, use these **show mpls ldp neighbor** commands in privileged EXEC mode:

- **show mpls ldp neighbor** [**vrf** *vpn-name*] [*address*] [*interface*] [**detail**]
- **show mpls ldp neighbor** [**all**]

show mpls ldp neighbor Syntax Description

Parameter	Description
vrf *vpn-name*	(Optional) Displays the LDP neighbors for the specified VPN routing or forwarding instance (*vpn-name*)
address	(Optional) Identifies the neighbor with this IP address
interface	(Optional) Defines the LDP neighbors accessible over this interface
detail	(Optional) Displays information in long form
all	(Optional) Displays LDP neighbor information for all VPNs when the **all** keyword is specified alone in this command, including those in the default routing domain

show mpls ldp bindings

To display the contents of the LIB, use this **show mpls ldp bindings** command in privileged EXEC mode: **show mpls ldp bindings** [*network* {*mask* | *length*} [**longer-prefixes**]] [**local-label** *label* [-*label*]] [**remote-label** *label* [-*label*]] [**neighbor** *address*] [**local**].

show mpls ldp bindings Syntax Description

Parameter	Description
vrf *vpn-name*	(Optional) This parameter displays the label bindings for the specified VPN routing or forwarding instance (*vpn-name*).
network	(Optional) This parameter defines the destination network number.
mask	(Optional) This parameter specifies the network mask, written as A.B.C.D.
length	(Optional) This parameter specifies the mask length (1 to 32 characters).
longer-prefixes	(Optional) This parameter selects any prefix that matches *mask* with a length from 1 to 32 characters.
local-label *label-label*	(Optional) This parameter displays entries matching local label values. Use the *label-label* argument to indicate the label range.
remote-label *label-label*	(Optional) This parameter displays entries matching the label values assigned by a neighbor router. Use the *label-label* argument to indicate the label range.
neighbor *address*	(Optional) This parameter displays the label bindings assigned by the selected neighbor.
local	(Optional) This parameter displays the local label bindings.

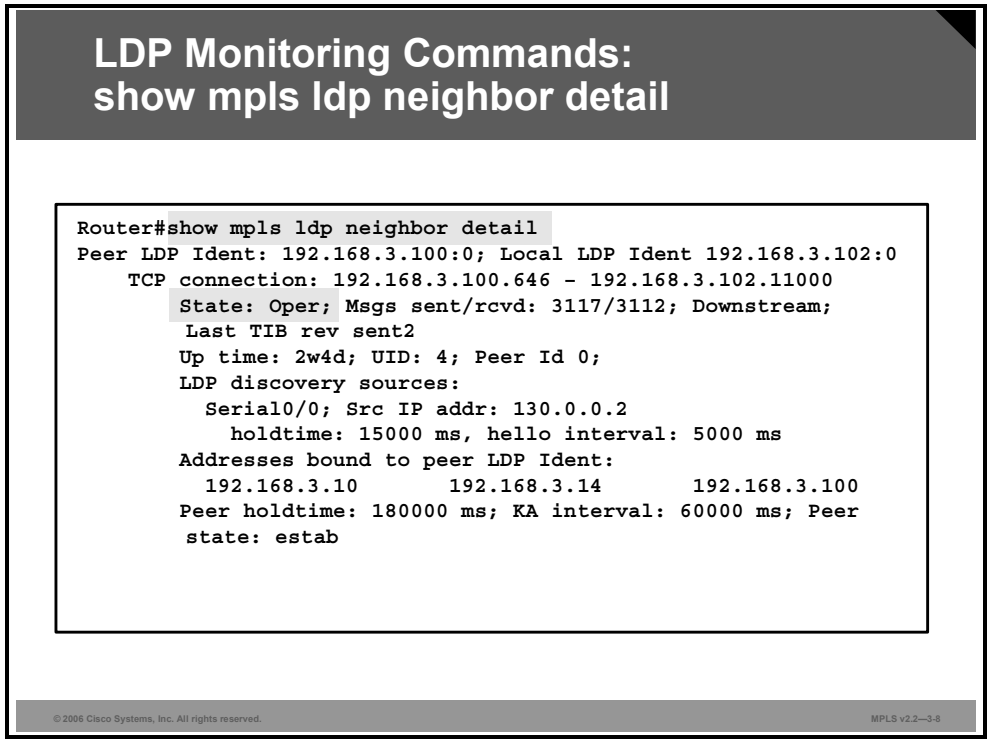

The status of the LDP session is indicated by "State: Oper" (operational).

show mpls ldp neighbor

To display the status of LDP sessions, issue these **show mpls ldp neighbor** commands in privileged EXEC mode:

- **show mpls ldp neighbor** [**vrf** *vpn-name*] [*address*] [**interface**] [**detail**]
- **show mpls ldp neighbor** [**all**]

Usage Guidelines

The **show mpls ldp neighbor** command can provide information about all LDP neighbors, or the information can be limited to the following:

- Neighbor with specific IP address
- LDP neighbors known to be accessible over a specific interface

This table describes the significant fields in the display.

show mpls ldp neighbor Field Description

Field	Description
Peer LDP Ident	This field displays the LDP ID of the neighbor (peer) for this session.
Local LDP Ident	This field displays the LDP ID for the local LSR for this session.
TCP connection	This field displays the TCP connection used to support the LDP session, shown in the format that follows: ■ peer IP address.peer port ■ local IP address.local port
State	This field displays the state of the LDP session. Generally, this is “Oper” (operational), but “transient” is another possible state.
Msgs sent/rcvd	This field displays number of LDP messages sent to and received from the session peer. The count includes the transmission and receipt of periodic keepalive messages, which are required for maintenance of the LDP session.
Downstream on demand	This field indicates that the downstream-on-demand method of label distribution is being used for this LDP session. When the downstream-on-demand method is used, an LSR advertises its locally assigned (incoming) labels to its LDP peer only when the peer requests them.
Downstream	This field Indicates that the downstream method of label distribution is being used for this LDP session. When the downstream method is used, an LSR advertises all of its locally assigned (incoming) labels to its LDP peer (subject to any configured access list restrictions).
Up time	This field displays the length of time that the LDP session has existed.
LDP discovery sources	This field displays the source (or sources) of LDP discovery activity that led to the establishment of this LDP session.
Addresses bound to peer LDP Ident	This field displays the known interface addresses of the LDP session peer. These are addresses that might appear as next-hop addresses in the local routing table. They are used to maintain the LFIB.
Peer holdtime	This field displays the time that it takes to remove the relationship if no keepalives are received within this period.
KA interval	This field displays the keepalive interval.
Peer state	This field shows the status of the neighbor relationship.

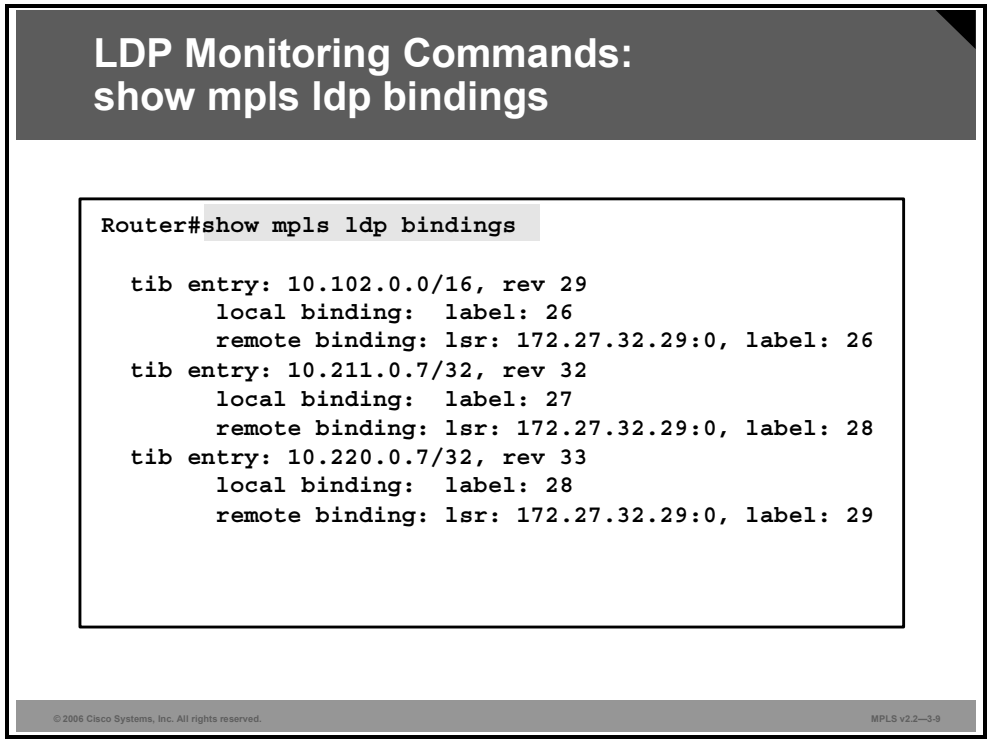

show mpls ldp bindings

To display the contents of the LIB, use this **show mpls ldp bindings** command in privileged EXEC mode: **show mpls ldp bindings** [**vrf** *vpn-name*] [*network* {*mask* | *length*} [**longer-prefixes**]] [**local-label** *label* [-*label*]] [**remote-label** *label* [-*label*]] [**neighbor** *address*] [**local**].

show mpls ldp bindings Syntax Description

Parameter	Description
vrf *vpn-name*	(Optional) This parameter displays the label bindings for the specified VPN routing or forwarding instance (vpn-name).
network	(Optional) This parameter defines the destination network number.
mask	(Optional) This parameter specifies the network mask, written as A.B.C.D.
length	(Optional) This parameter specifies the mask length (1 to 32 characters).
longer-prefixes	(Optional) This parameter selects any prefix that matches *mask* with a length from 1 to 32 characters.
local-label *label-label*	(Optional) This parameter displays entries matching local label values. Use the *label-label* argument to indicate the label range.
remote-label *label-label*	(Optional) This parameter displays entries matching the label values assigned by a neighbor router. Use the *label-label* argument to indicate the label range.
neighbor *address*	(Optional) This parameter displays the label bindings assigned by the selected neighbor.
local	(Optional) This parameter displays the local label bindings.

Usage Guidelines

The **show mpls ldp bindings** command displays label bindings learned by the LDP or TDP.

Examples

This sample output from the **show mpls ldp bindings** command displays the contents of the entire LIB.

```
Router1#show mpls ldp bindings
  10.92.0.0/16, rev 28
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  10.102.0.0/16, rev 29
        local binding:  label: 26
        remote binding: lsr: 172.27.32.29:0, label: 26
  10.105.0.0/16, rev 30
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  10.205.0.0/16, rev 31
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  10.211.0.7/32, rev 32
        local binding:  label: 27
        remote binding: lsr: 172.27.32.29:0, label: 28
  10.220.0.7/32, rev 33
        local binding:  label: 28
        remote binding: lsr: 172.27.32.29:0, label: 29
  99.101.0.0/16, rev 35
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  100.101.0.0/16, rev 36
        local binding:  label: 29
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  171.69.204.0/24, rev 37
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  172.27.32.0/22, rev 38
        local binding:  label: imp-null
        remote binding: lsr: 172.27.32.29:0, label: imp-null
  210.10.0.0/16, rev 39
        local binding:  label: imp-null
```

Monitoring Label Switching

This topic describes how to monitor label switching.

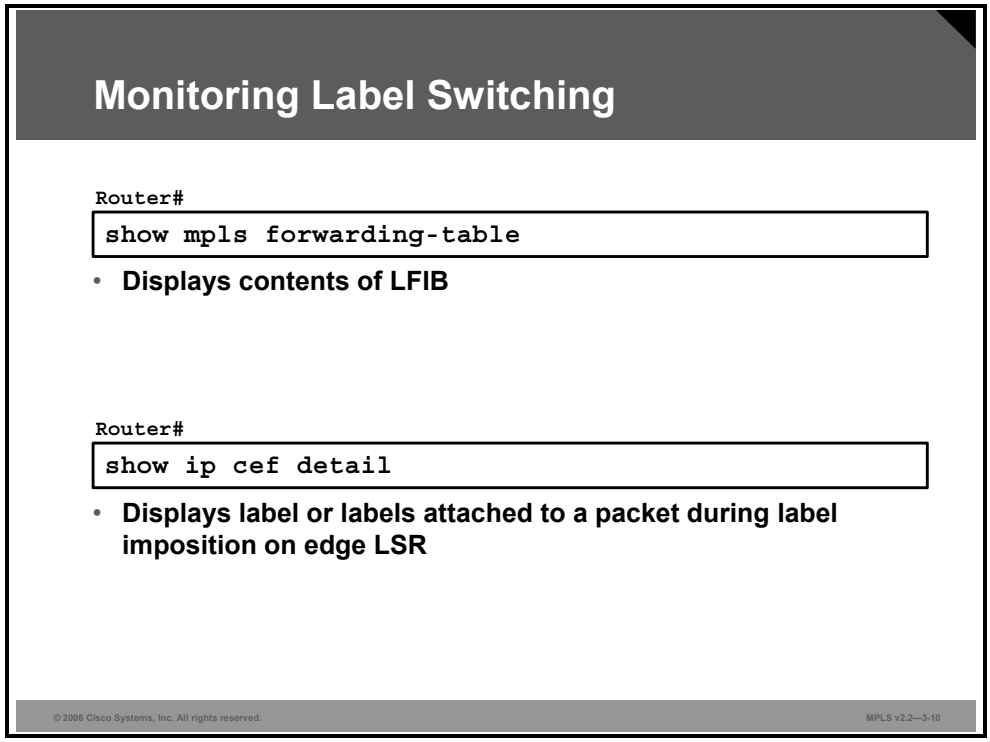

show mpls forwarding-table

To display the contents of the MPLS LFIB, use this **show mpls forwarding-table** command in privileged EXEC mode: **show mpls forwarding-table** [{*network* {*mask* | *length*} | **labels** *label* [-*label*]| **interface** *interface* | **next-hop** *address* | **lsp-tunnel** [*tunnel-id*]}] [**detail**].

show ip cef

To display entries in the FIB that are unresolved or to display a summary of the FIB, use the this form of the **show ip cef** command in privileged EXEC mode: **show ip cef** [**unresolved** | **summary**].

To display specific entries in the FIB based on IP address information, use this form of the **show ip cef** command in privileged EXEC mode: **show ip cef** [*network* [*mask* [**longer-prefix**]]] [**detail**].

To display specific entries in the FIB based on interface information, use this form of the **show ip cef** command in privileged EXEC mode: **show ip cef** [*type number*] [**detail**].

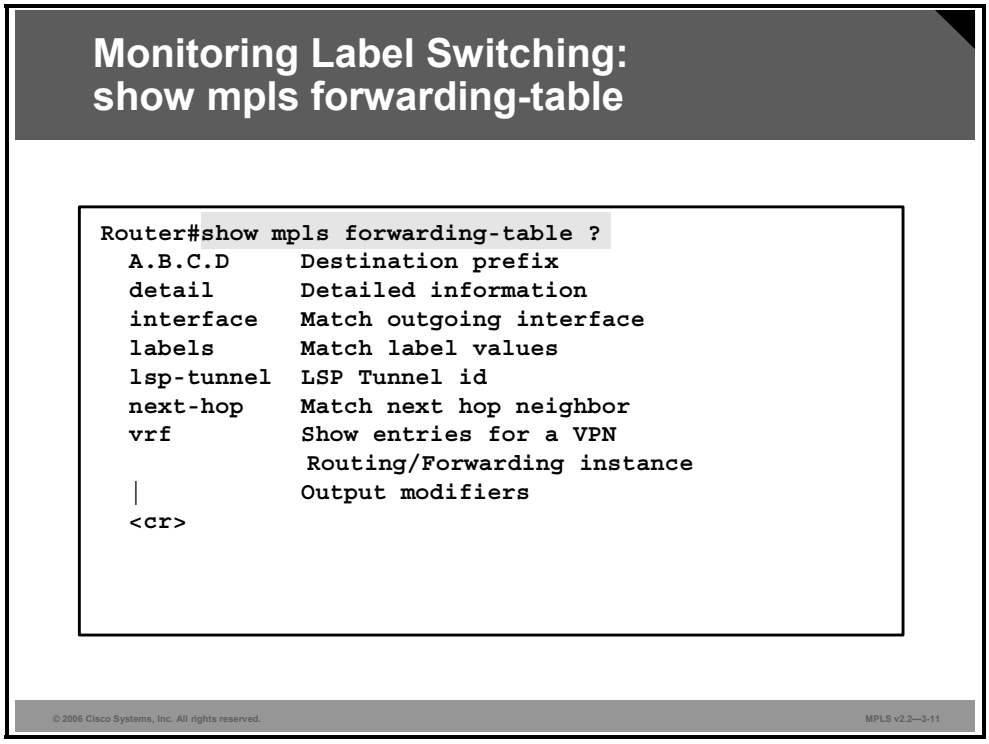

show mpls forwarding-table

To display the contents of the MPLS LFIB, use this **show mpls forwarding-table** command in privileged EXEC mode: **show mpls forwarding-table** [{*network* {*mask* | *length*} | **labels** *label* [-*label*]| **interface** *interface* | **next-hop** *address* | **lsp-tunnel** [*tunnel-id*]}] [**detail**].

show mpls forwarding-table Syntax Description

Parameter	Description
network	(Optional) Displays destination network number
mask	Displays IP address of destination mask whose entry is to be shown
length	Displays number of bits in mask of destination
labels *label-label*	(Optional) Shows only entries with specified local labels
interface *interface*	(Optional) Shows only entries with specified outgoing interface
next-hop *address*	(Optional) Shows only entries with specified neighbor as next hop
lsp-tunnel *tunnel-id*	(Optional) Shows only entries with specified LSP tunnel, or all LSP tunnel entries
detail	(Optional) Displays information in long form (includes length of encapsulation, length of MAC string, maximum transmission unit (MTU), and all labels)

Examples: show mpls forwarding table Command Output

This is a sample output from the **show mpls forwarding table** command.

```
Router#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing     Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
26     Untagged    10.253.0.0/16     0          Et4/0/0      172.27.232.6
28     1/33        10.15.0.0/16      0          AT0/0.1      point2point
29     Pop tag     10.91.0.0/16      0          Hs5/0        point2point
       1/36        10.91.0.0/16      0          AT0/0.1      point2point
30     32          10.250.0.97/32    0          Et4/0/2      10.92.0.7
       32          10.250.0.97/32    0          Hs5/0        point2point
34     26          10.77.0.0/24      0          Et4/0/2      point2point
       26          10.77.0.0/24      0          Hs5/0        point2point
35     Untagged[T] 10.100.100.101/32 0          Tu1          point2point
36     Pop tag     168.1.0.0/16      0          Hs5/0        point2point
       1/37        168.1.0.0/16      0          AT0/0.1      point2point
```

[T] = Forwarding through an LSP tunnel.

Note View additional tagging information with the **detail** option.

Monitoring Label Switching: show mpls forwarding-table detail

```
Router#show mpls forwarding-table detail
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
70     Pop tag     192.168.3.3/32       0        Se0/0.111  point2point
        MAC/Encaps=4/4, MTU=1504, Tag Stack{}
        18F18847
        No output feature configured
    Per-packet load-sharing
71     Pop tag     192.168.3.4/32    0          Se0/0.111   point2point
        MAC/Encaps=4/4, MTU=1504, Tag Stack{}
        18F18847
        No output feature configured
    Per-packet load-sharing
72     16          192.168.1.97/32   0          Se0/0.111  point2point
        MAC/Encaps=4/8, MTU=1500, Tag Stack{16}
        18F18847 00010000
        No output feature configured
    Per-packet load-sharing
```


The table describes the significant fields in the display.

show mpls forwarding table Field Description

Field	Description
Local tag	This field displays the label assigned by this router.
Outgoing tag or VC	This field displays the label assigned by the next hop or virtual path identifier/virtual channel identifier (VPI/VCI) used to get to next hop. Some of the entries that you can specify in this column are as follows: **[T]:** Forwarding is through an LSP tunnel. **untagged:** There is no label for the destination from the next hop, or label switching is not enabled on the outgoing interface. **Pop tag:** The next hop advertised an implicit null label for the destination, and this router popped the top label.
Prefix or Tunnel ID	This field displays the address or tunnel to which packets with this label are going.
Bytes tag switched	This field displays the number of bytes switched with this incoming label.
Outgoing interface	This field displays the interface through which packets with this label are sent.
Next Hop	This field displays the IP address of the neighbor that assigned the outgoing label.
MAC/Encaps	This field displays the length in bytes of Layer 2 header, and length in bytes of packet encapsulation, including Layer 2 header and label header.
MTU	This field displays the MTU of the labeled packet.
Tag Stack	This field displays all the outgoing labels. If the outgoing interface is transmission convergence-ATM (TC-ATM), the virtual circuit descriptor (VCD) is also shown.
18F18847 00010000	This field displays the actual encapsulation in hexadecimal form. There is a space shown between Layer 2 and the label header.

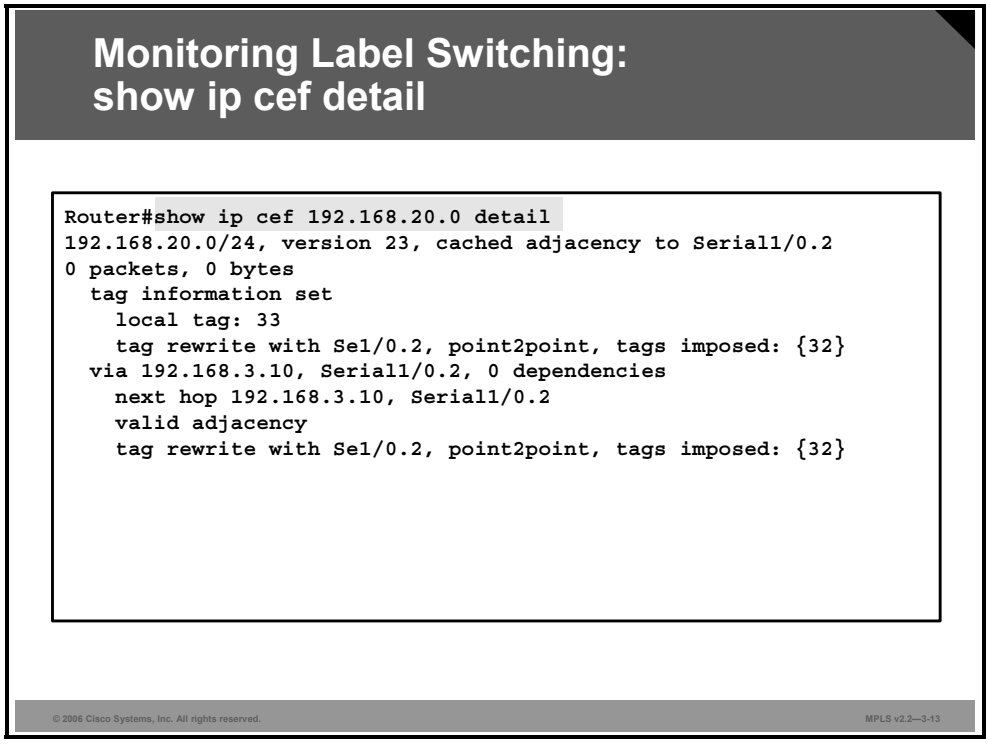

show ip cef detail

To display detailed FIB entry information for all FIB entries, use this **show ip cef detail** command in privileged EXEC mode: **show ip cef** [*type number*] [**detail**].

show ip cef detail Syntax Description

Parameter	Description
unresolved	(Optional) Displays unresolved FIB entries
summary	(Optional) Displays a summary of the FIB
network	(Optional) Displays the FIB entry for the specified destination network
mask	(Optional) Displays the FIB entry for the specified destination network and mask
longer-prefix	(Optional) Displays FIB entries for all more specific destinations
detail	(Optional) Displays detailed FIB entry information
type number	(Optional) Displays interface type and number for which to display FIB entries

Usage Guidelines

The **show ip cef** command without any keywords or arguments shows a brief display of all FIB entries.

The **show ip cef detail** command shows detailed FIB entry information for all FIB entries.

Debugging MPLS and LDP

This topic describes how to debug MPLS and LDP.

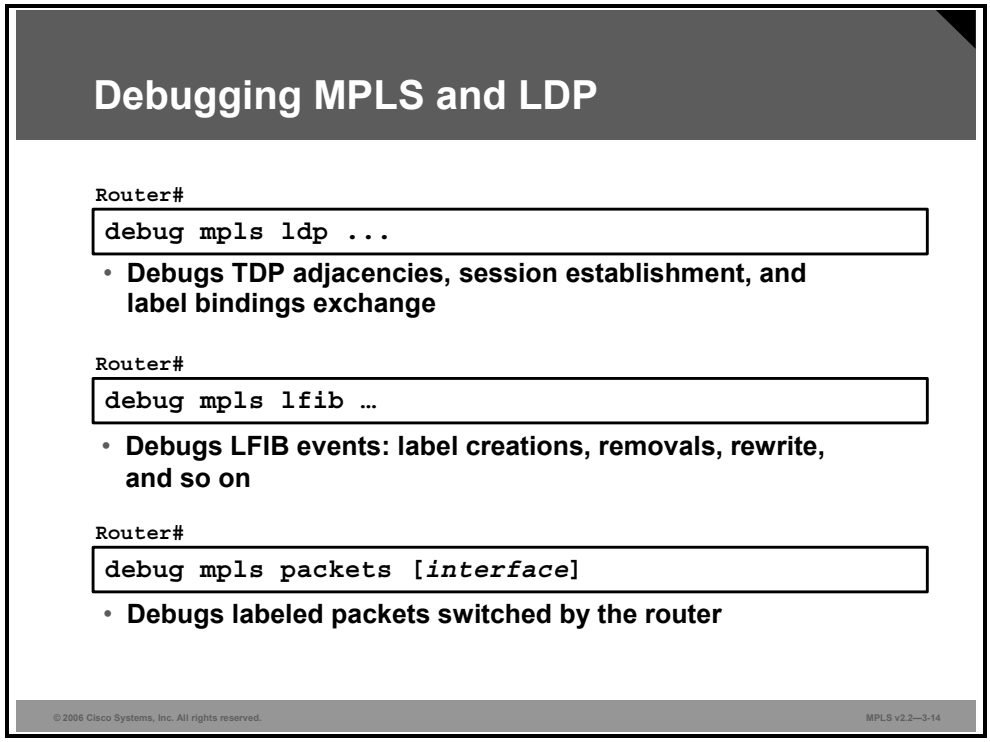

A large number of **debug** commands are associated with MPLS on Cisco IOS platforms. The **debug mpls ldp** commands debug various aspects of LDP protocol, from label distribution to exchange of the application layer data between adjacent LDP-speaking routers.

Note	Use **debug** commands with caution. Enabling debugging can disrupt operation of the router under high load conditions. Before you start a **debug** command, always consider the output that the command may generate and the amount of time this may take. You should also look at your CPU load before debugging by using the **show processes cpu** command. Verify that you have ample CPU available before beginning the debugging process.

The **debug mpls lfib** commands display LFIB-related events (allocation of new labels, removal of labels, and so on).

The **debug mpls packets** command displays all labeled packets switched by the router (through the specified interface).

Caution	Use the **debug mpls packets** command with care, because it generates output for every packet processed. Furthermore, enabling the **debug mpls packets** command causes fast and distributed label switching to be disabled for the selected interfaces. To avoid adversely affecting other system activity, use this command only when traffic on the network is at a minimum.

debug mpls packets

To display labeled packets switched by the host router, use the **debug mpls packets** command in privileged EXEC mode. To disable debugging output, use the **no** form of this command. This illustrates these two commands:

- **debug mpls packets** [*interface*]
- **no debug mpls packets** [*interface*]

debug mpls packets Syntax Description

Field	Description
Hs0/0	Displays the identifier for the interface on which the packet was received or transmitted
Recvd	Displays the packet received
Xmit	Displays the packet transmitted
CoS	Displays the class of service (CoS) field from the packet label header
TTL	Displays the time-to-live (TTL) field from the packet label header
(no tag)	Displays the last label popped off the packet and transmitted unlabeled
Tag(s)	Displays a list of labels on the packet, ordered from the top of the stack to the bottom

Summary

This topic summarizes the key points that were discussed in this lesson.

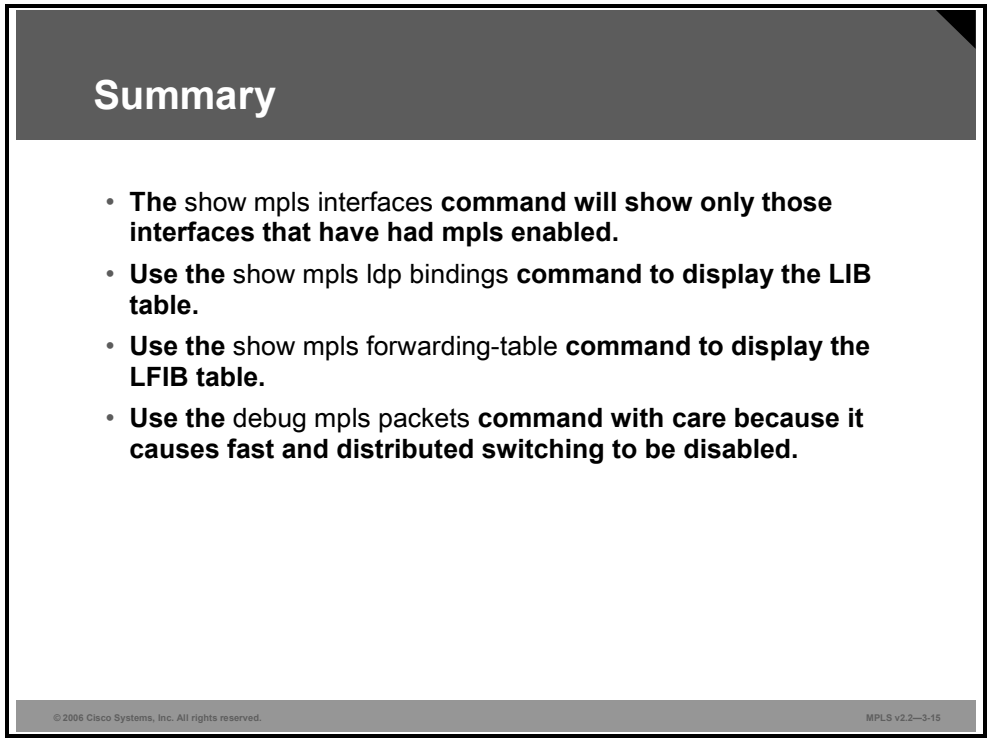

Lesson 4

Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms

Overview

This lesson looks at some of the common issues that arise in Multiprotocol Label Switching (MPLS) networks. For each issue discussed, there is a recommended troubleshooting procedure to resolve the issue.

It is very important to know what commands you can use to verify correct operation of MPLS in the network. The information here will help you when you encounter problems with frame-mode interfaces that have MPLS running in the network.

Objectives

Upon completing this lesson, you will be able to describe how to troubleshoot frame-mode MPLS problems on Cisco IOS platforms. This ability includes being able to meet these objectives:

- Identify the common issues that arise in MPLS networks
- Solve LDP session startup issues
- Solve label allocation issues that can arise in MPLS networks
- Solve label distribution issues that can arise in MPLS networks
- Solve packet-labeling issues that can arise in MPLS networks
- Solve intermittent MPLS failures
- Solve packet propagation issues in MPLS networks

What Are Common Frame-Mode MPLS Issues?

This topic identifies some of the common frame-mode issues that arise in MPLS networks.

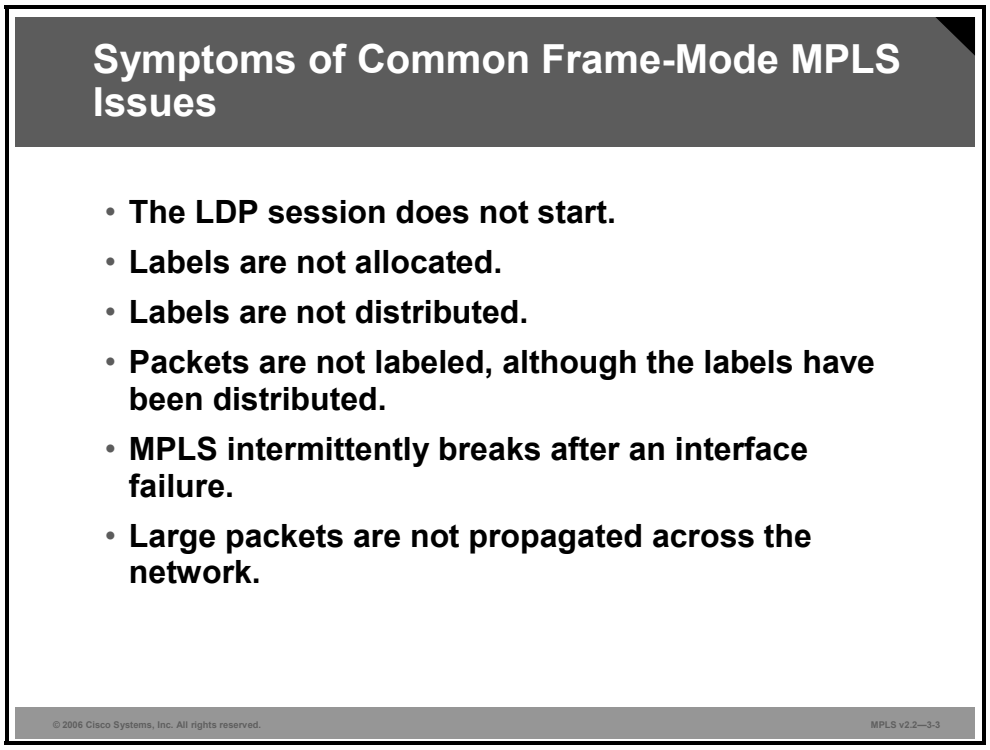

Here are the common issues that can be encountered while you are troubleshooting a frame-mode MPLS network:

- The Label Distribution Protocol (LDP) session does not start.
- The LDP session starts, but the labels are not allocated or distributed.
- Labels are allocated and distributed, but the forwarded packets are not labeled.
- MPLS stops working intermittently after an interface failure, even on interfaces totally unrelated to the failed interface.
- Large IP packets are not propagated across the MPLS backbone, even though the packets were successfully propagated across the pure IP backbone.

This discussion will cover each of these issues and provide recommended steps for troubleshooting them.

Solving LDP Session Startup Issues

This topic describes how to solve LDP session startup issues found in MPLS networks.

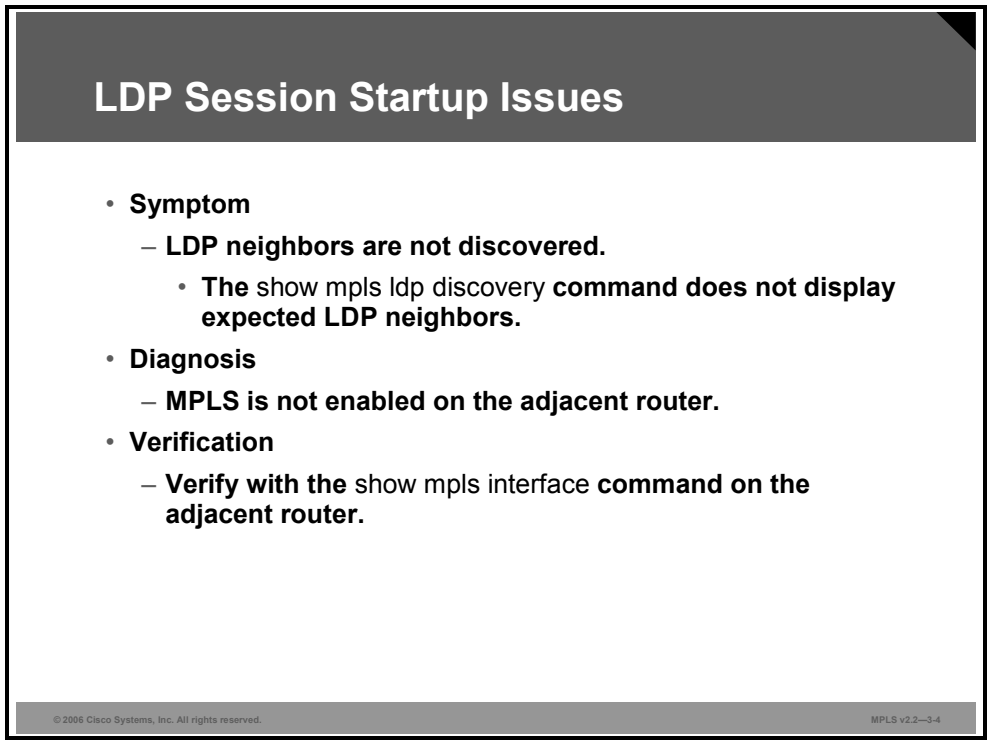

Diagnosis: If MPLS is enabled on an interface, but no neighbors are discovered, it is likely that MPLS is not enabled on the neighbor.

The router is sending discovery messages, but the neighbor is not replying because it does not have LDP enabled.

Solution: Enable MPLS on the neighboring router.

LDP Session Startup Issues (Cont.)

- **Symptom**
  - **LDP neighbors are not discovered.**
- **Diagnosis**
  - **There is a label distribution protocol mismatch—TDP on one end, LDP on the other end.**
- **Verification**
  - **Verify with the** show mpls interface detail **command on both routers.**

Diagnosis: Another possibility is that the neighbor has a different label distribution protocol enabled on the interface, such as when LDP is enabled on one end and Tag Distribution Protocol (TDP) is enabled on the other end.

Solution: Use one of these solutions:

- Change the label distribution protocol on this end.
- Change the label distribution protocol on the other end.
- Enable both label distribution protocols on this end.
- Enable both label distribution protocols on the other end.

LDP Session Startup Issues (Cont.)

- **Symptom**
  - **LDP neighbors are not discovered.**
- **Diagnosis**
  - **Packet filter drops LDP neighbor discovery packets.**
- **Verification**
  - **Verify access list presence with the** show ip interface **command.**
  - **Verify access list contents with the** show access-list **command.**

Diagnosis: MPLS configurations match on both ends, but the session still does not get established. Check whether there are any input access lists that deny discovery messages.

Solution: Remove or change the access list to allow User Datagram Protocol (UDP) packets with source and destination port number 646 (711 for TDP).

Make sure that the access list also allows TCP to and from port 646 (711 for TDP).

LDP Session Startup Issues (Cont.)

- **Symptom**
  - **LDP neighbors are discovered; the LDP session is not established.**
    - **The** show mpls ldp neighbor **command does not display a neighbor in operational state.**
- **Diagnosis**
  - **The connectivity between loopback interfaces is broken; the LDP session is usually established between loopback interfaces of adjacent LSRs.**
- **Verification**
  - **Verify connectivity with the extended** ping **command.**

Diagnosis: LDP neighbors are exchanging hello packets, but the LDP session is never established.

Solution: Check the reachability of the loopback interfaces, because they are typically used to establish the LDP session. Make sure that the loopback addresses are exchanged via the Interior Gateway Protocol (IGP) used in the network.

Solving Label Allocation Issues

This topic describes how to solve label allocation issues that could arise in MPLS networks.

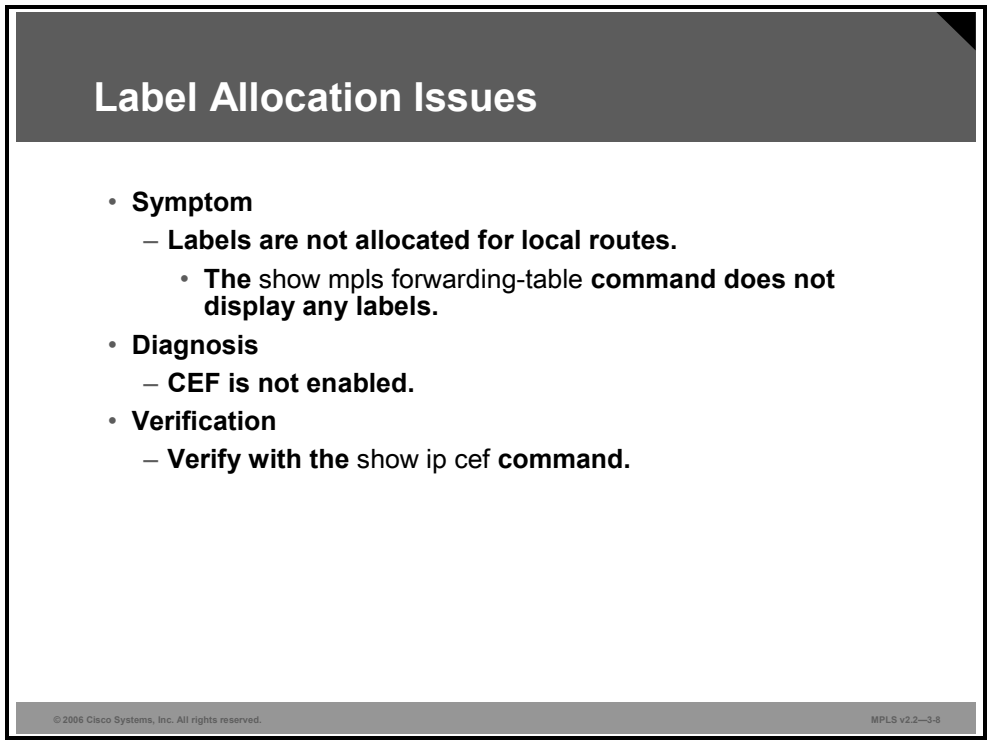

Diagnosis: Labels are not allocated for any or some of the local routes. Use the **show ip cef** command to verify whether Cisco Express Forwarding (CEF) switching is enabled on all MPLS-enabled interfaces.

Solution: Enable CEF switching by using the **ip cef** command in global configuration mode or the **ip route-cache cef** command in interface mode.

Solving Label Distribution Issues

This topic describes how to solve label distribution issues that can arise in MPLS networks.

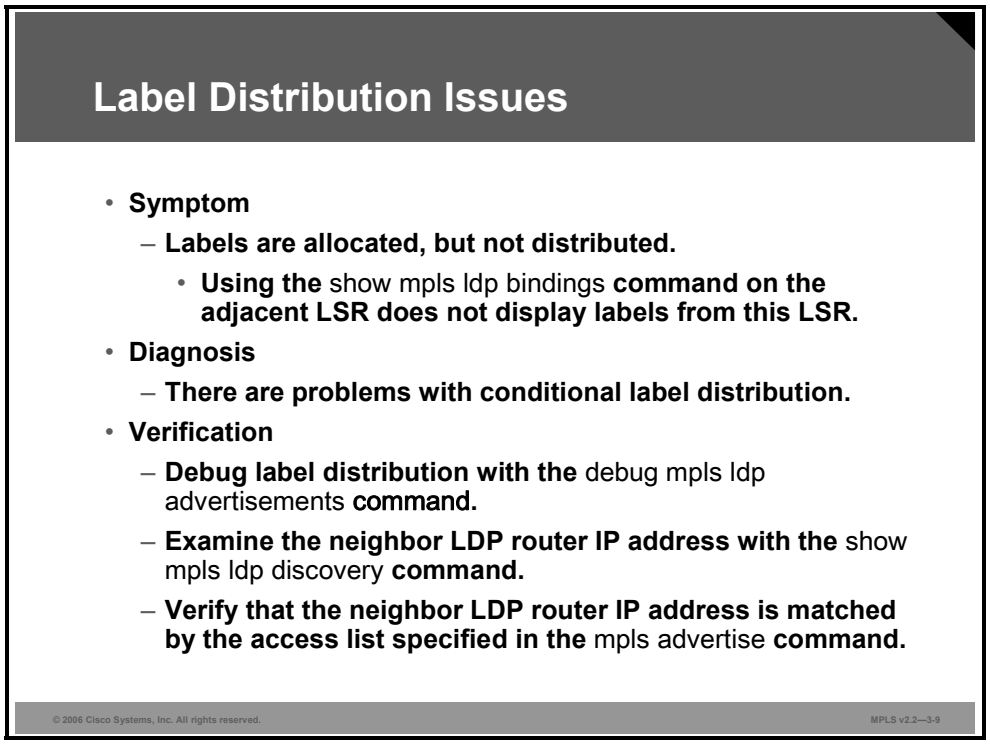

Symptom: Labels are generated for local routes on one label switch router (LSR) but are not received on neighboring LSRs.

Solution: Check whether conditional label advertising is enabled and verify both access lists that are used with the command.

Solving Packet-Labeling Issues

This topic describes how to solve packet-labeling issues that can arise in MPLS networks.

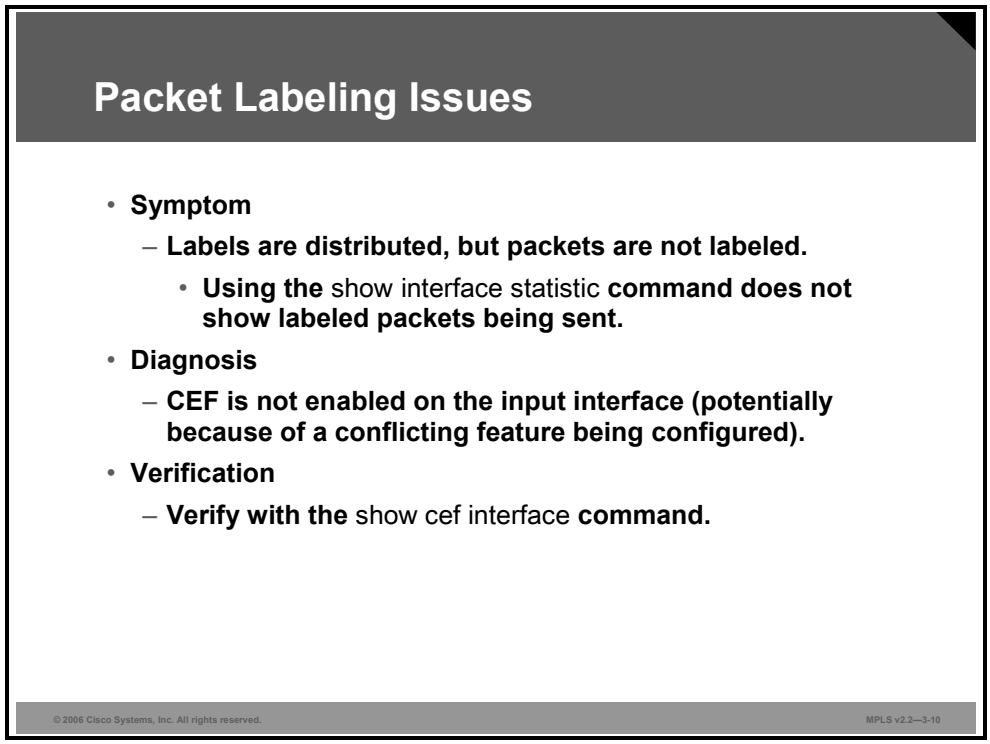

Symptom: Labels exist, but packets are not labeled.

Solution: Enable CEF switching by using the **ip route-cache cef** interface command and make sure that there is no feature enabled on the interface that is not supported in combination with CEF switching. Verify whether CEF is enabled on an individual interface with the **show cef interface** command.

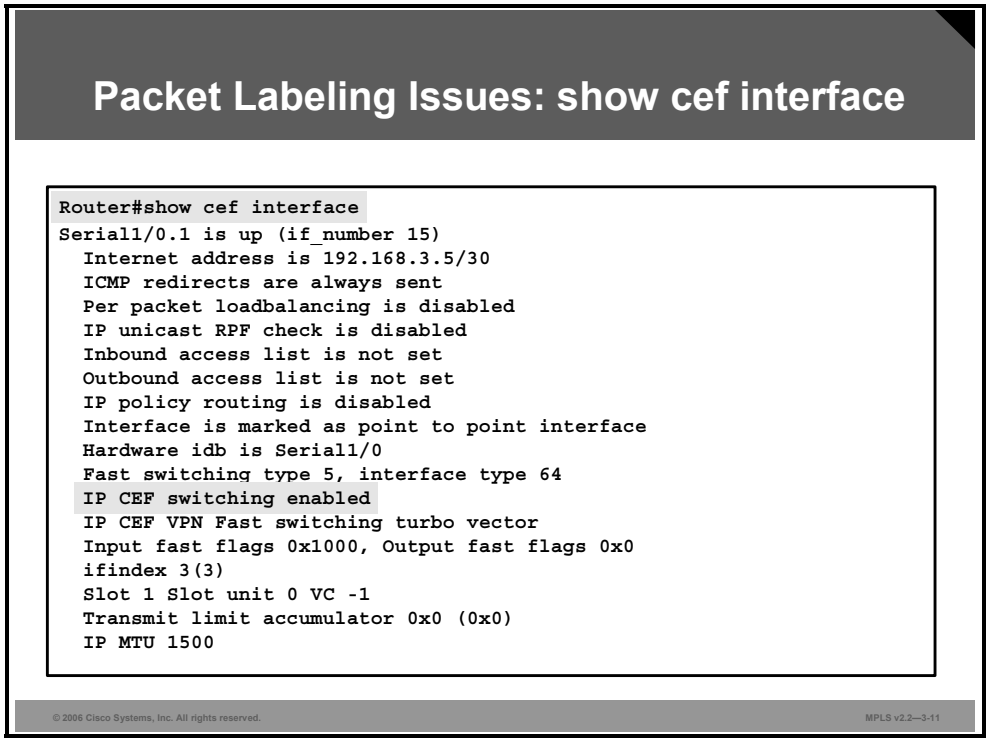

show cef interface

The **show cef interface** command is used in privileged EXEC mode to display CEF interface information.

The table describes the parameters for the **show cef interface** *type number* [**detail**] command.

show cef interface Syntax Descriptions

Parameter	Description
type number	Displays interface number and the number about which to display CEF-related information
detail	(Optional) Displays detailed CEF information for the specified interface port number

Usage Guidelines

The **show cef interface** command is available on routers that have route processor (RP) cards and line cards.

You can use this command to show the CEF state on an individual interface.

The table describes the significant fields in the display.

show cef interface Field Descriptions

Field	Description
interface type number is {up | down}	Indicates status of the interface
Internet address	Displays Internet address of the interface
ICMP redirects are {always sent | never sent}	Indicates how packet forwarding is configured
Per-packet load balancing	Displays status of load balancing in use on the interface (enabled or disabled)
Inbound access list {# | Not set}	Displays number of access lists defined for the interface
Outbound access list	Displays number of access lists defined for the interface
Hardware idb is *type number*	Displays interface type and number configured
Fast switching type	Indicates switching mode in use—used for troubleshooting
IP Distributed CEF switching {enabled | disabled}	Indicates the switching path used
Slot *n* Slot unit *n*	Displays the slot number
Transmit line accumulator	Indicates the maximum number of packets allowed in the transmit queue
IP MTU	Displays the value of the maximum transmission unit (MTU) size set on the interface

Solving Intermittent MPLS Failures

This topic describes how to solve intermittent MPLS failures.

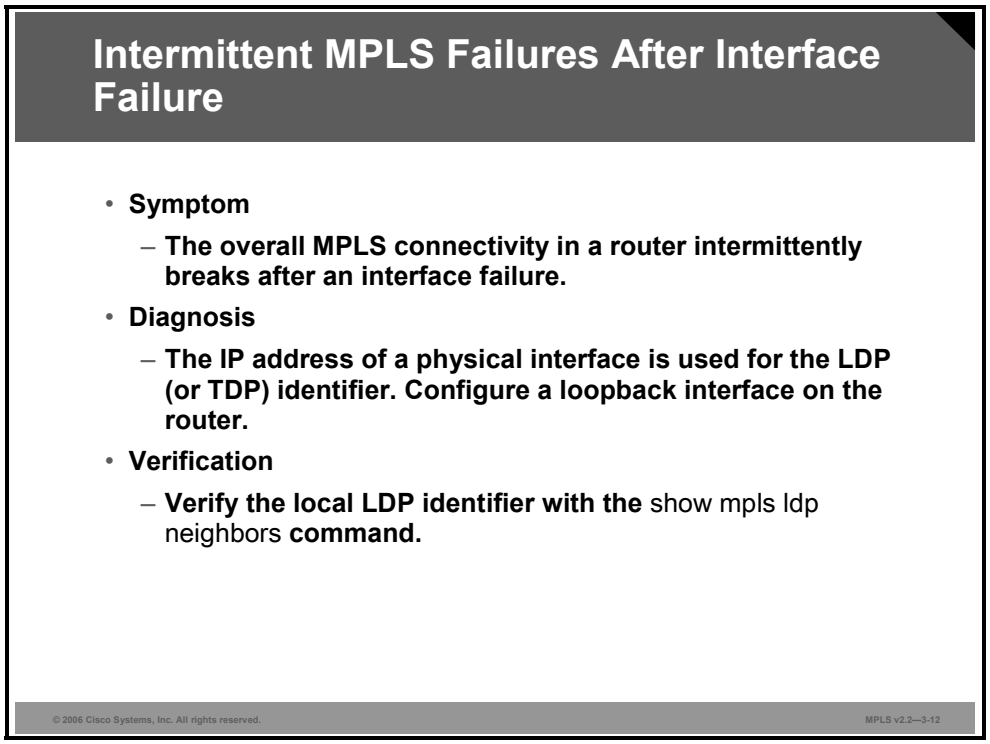

Symptom: MPLS connectivity is established, labels are exchanged, and packets are labeled and forwarded. However, an interface failure can sporadically stop an MPLS operation on unrelated interfaces in the same router.

Details: LDP sessions are established between IP addresses that correspond to the LDP LSR identifier (ID). The LDP LSR ID is assigned using the algorithm that is also used to assign an Open Shortest Path First (OSPF) or a Border Gateway Protocol (BGP) router ID.

This algorithm selects the highest IP address of an active interface if there are no loopback interfaces configured on the router. If that interface fails, the LDP LSR ID is lost and the TCP session carrying the LDP data is torn down, resulting in loss of all neighbor-assigned label information.

The symptom can be easily verified with the **show mpls ldp neighbors** command, which displays the local and remote LSR ID. Verify that both of these IP addresses are associated with a loopback interface.

Solution: Configure a loopback interface on the LSR.

Note	The LDP LSR ID will change only after the router is reloaded.

Solving Packet Propagation Issues

This topic describes how to solve packet propagation issues in an MPLS network.

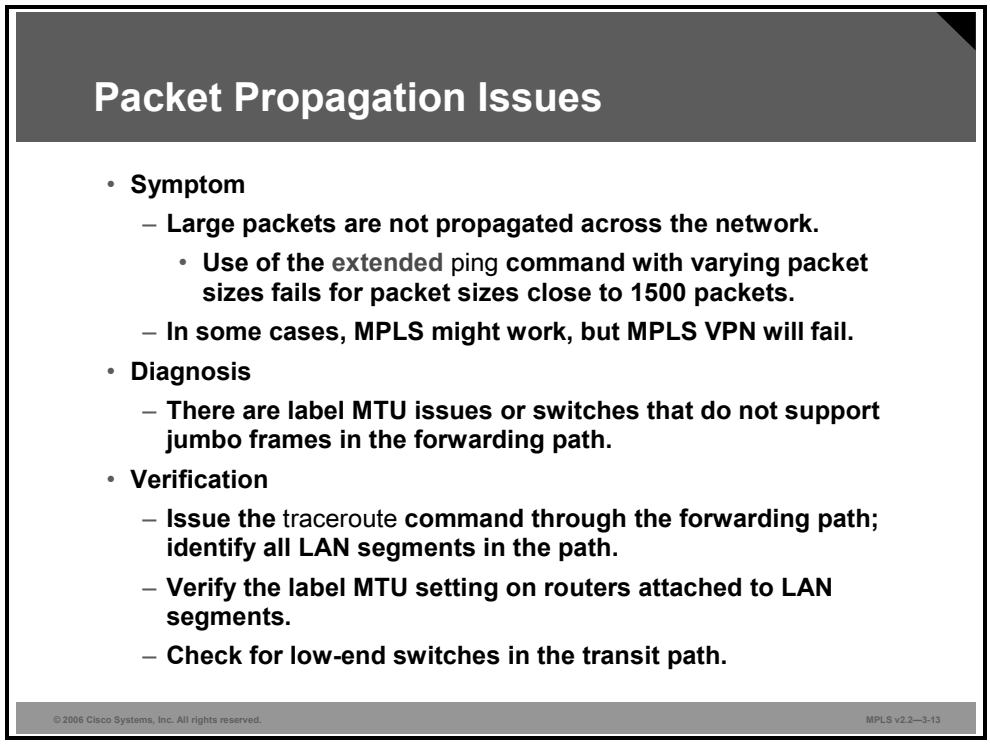

Symptom: Packets are labeled and sent, but they are not received on the neighboring router. A LAN switch between the adjacent MPLS-enabled routers may drop the packets if it does not support jumbo frames. In some cases, MPLS might work, but MPLS Virtual Private Network (VPN) will fail.

Solution: Change the MPLS MTU size, taking into account the maximum number of labels that may appear in a packet.

Summary

This topic summarizes the key points that were discussed in this lesson.

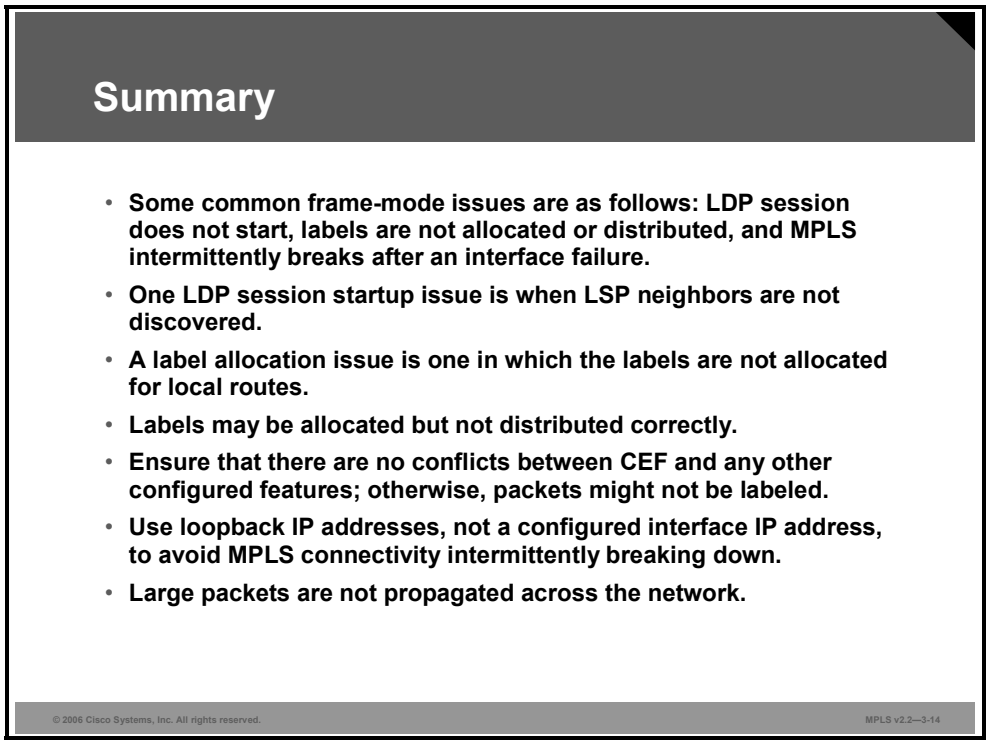

Module Summary

This topic summarizes the key points that were discussed in this module.

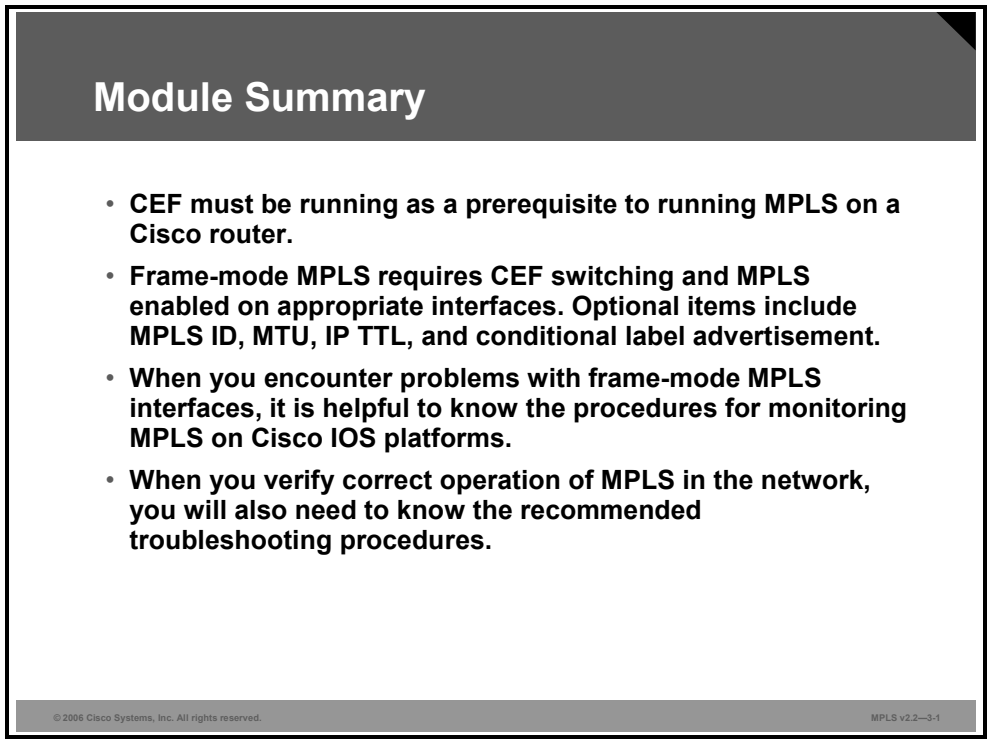

There are many detailed configuration, monitoring, and debugging guidelines when implementing frame-mode Multiprotocol Label Switching (MPLS) on Cisco IOS platforms. Advanced technologies, such as time-to-live (TTL) propagation and label distribution, are also critical when switching implementations.

References

For additional information, refer to these resources:

- Cisco Express Forwarding Overview
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca7cb.html
- Configuring Cisco Express Forwarding
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca7cc.html
- Multiprotocol Label Switching Overview
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca7cc.html

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) What is another name for topology-driven switching? (Source: Introducing CEF Switching)

A) CEF
B) fast switching
C) cache switching
D) process switching

Q2) What is the command to monitor CEF? (Source: Introducing CEF Switching)

A) Router#**show cef**
B) Router**#show mpls ip cef**
C) Router#**show ip cef**
D) Router#**show mpls cef**

Q3) What is the command to enable CEF on a Cisco router? (Source: Introducing CEF Switching)

A) Router(config)#**mpls ip cef**
B) Router(config)#**ip mpls cef**
C) Router(config)#**cef**
D) Router(config)#**ip cef**

Q4) In CEF switching, what is the difference between the adjacency table and the ARP cache? (Source: Introducing CEF Switching)

A) The adjacency table holds the Layer 2 header, and the ARP cache does not.
B) The ARP cache holds the Layer 2 header, and the adjacency table does not.
C) Both the adjacency table and the ARP cache hold the Layer 2 header.
D) Neither the adjacency table nor the ARP cache holds the Layer 2 header.

Q5) What happens to a packet that should be fast-switched but the destination is not in the switching cache? (Source: Introducing CEF Switching)

A) The packet is dropped.
B) The packet is cache-switched.
C) The packet is process-switched.
D) CEF switching is used.

Q6) If IP TTL propagation is not allowed, what is the value that is placed in the MPLS header? (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

A) 0
B) 1
C) 254
D) 255

Q7) The MPLS MTU is increased to _____ to support 1500-byte IP packets and MPLS stacks up to 3 levels deep. (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

Q8) What is the correct command to enable MPLS in Cisco IOS software? (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

A) Router(config)#**ip mpls**
B) Router(config-if)#**ip mpls**
C) Router(config)#**mpls ip**
D) Router(config-if)#**mpls ip**

Q9) Which two steps are NOT mandatory to enable MPLS? (Choose two.) (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

A) Enable CEF switching.
B) Configure the size of the label pool.
C) Configure the MTU size for labeled packets.
D) Configure LDP (or TDP) on every interface that will run MPLS.

Q10) What needs to be configured to specify which neighbors would selectively receive label advertisements? (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

A) Controlled label distribution needs to be configured.
B) Conditional label distribution needs to be configured.
C) Unsolicited label distribution needs to be configured.
D) All neighbors will receive all labels.

Q11) If frame-mode MPLS is run on ATM interfaces, LDP or LDP neighbor relationships are established between the _____ routers. (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

Q12) Which command is used to display information about the LDP Hello protocol timers? (Source: Monitoring Frame-Mode MPLS on Cisco IOS Platforms)

A) **show ip cef**
B) **show mpls ldp parameters**
C) **show ldp forwarding-table**
D) **show mpls ldp discovery**

Q13) Which command is used to display the contents of the LIB table? (Source: Monitoring Frame-Mode MPLS on Cisco IOS Platforms)

A) **show mpls ldp labels**
B) **show mpls ldp bindings**
C) **show mpls ldp neighbors**
D) **show mpls forwarding-table**

Q14) Which command is used to display the contents of the LFIB table? (Source: Monitoring Frame-Mode MPLS on Cisco IOS Platforms)

A) **show mpls ldp labels**
B) **show mpls ldp bindings**
C) **show mpls ldp neighbors**
D) **show mpls forwarding-table**

Q15) Which command would NOT be used to debug MPLS or LDP? (Source: Monitoring Frame-Mode MPLS on Cisco IOS Platforms)

A) **debug mpls ldp**
B) **debug mpls lfib**
C) **debug mpls packets**
D) **debug mpls ldp neighbors**

Q16) Which two of the answer choices would cause an LDP (or TDP) session NOT to be established between two LSRs? (Choose two.) (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) an access list that allows TCP/UDP port number 646
B) an access list that allows TCP/UDP port number 711
C) an access list that does not allow TCP/UDP port number 646
D) an access list that does not allow TCP/UDP port number 711

Q17) Which command is issued to troubleshoot label allocation issues? (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) **show cef**
B) **show lfib**
C) **show ip cef**
D) **show mpls lfib**

Q18) Which command is issued to see if labels are being distributed from the local LSR? (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) **show mpls ldp lib** (on the local router)
B) **show mpls ldp lib** (on the remote router)
C) **show mpls ldp bindings** (on the local router)
D) **show mpls ldp bindings** (on the remote router)

Q19) Which command displays CEF interface information? (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) router#**show mpls cef interface**
B) router#**show cef interface**
C) router(config)#**show cef interface**
D) router(config)#**show mpls cef interface**

Q20) To reduce the chances of having intermittent MPLS failures because of an interface failing, a _____ address should be configured. (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

Q21) A LAN switch is in the network path between two LSRs. It has been discovered that large packets are not being propagated across the network. Which of the answer choices represents the most possible cause? (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) The precedence bit has not been set in the MPLS label.
B) The TTL has not been set correctly to address this issue.
C) The MTU size has not been set correctly to address this issue.
D) This is not a legal configuration. LSRs must be directly connected.

Module Self-Check Answer Key

Q1) A

Q2) C

Q3) D

Q4) A

Q5) C

Q6) D

Q7) 1512

Q8) D

Q9) B, C

Q10) B

Q11) PVC endpoint

Q12) B

Q13) B

Q14) D

Q15) D

Q16) C, D

Q17) C

Q18) D

Q19) B

Q20) loopback

Q21) C

Module 4

MPLS VPN Technology

Overview

This module introduces Virtual Private Networks (VPNs) and two major VPN design options: the overlay VPN and the peer-to-peer VPN. The module also introduces VPN terminology and topologies, and describes Multiprotocol Label Switching (MPLS) VPN architecture and operations. This module details various customer edge-provider edge (CE-PE) routing options and Border Gateway Protocol (BGP) extensions (route targets and extended community attributes) that allow Internal Border Gateway Protocol (IBGP) to transport customer routes over a provider network. The MPLS VPN forwarding model is also covered together with how it integrates with core routing protocols.

Module Objectives

Upon completing this module, you will be able to describe the MPLS peer-to-peer architecture and explain the routing and packet-forwarding model in this architecture. This ability includes being able to meet these objectives:

- Identify the major terminology and topology of VPNs
- Describe the characteristics of the different VPN topologies
- Describe the major architectural components of an MPLS VPN
- Identify the routing requirements for MPLS VPNs
- Describe how packets are forwarded in an MPLS VPN environment

Lesson 1

Introducing VPNs

Overview

This lesson explains the concept of Virtual Private Networks (VPNs), and explains the VPN terminology that is also used by the Multiprotocol Label Switching (MPLS) VPN architecture. The lesson looks at why VPNs were first introduced, and also explains the differences between the overlay and peer-to-peer VPN models, how they are implemented, and the benefits and drawbacks of each implementation.

It is important to understand the background of VPNs, because you should be able to determine when an organization might need a VPN, and explain how MPLS VPNs can help save time and money. Understanding the different types of VPNs will allow you to recognize where the various types of VPNs would be best used in their associated networks.

Objectives

Upon completing this lesson, you will be able to identify the major terminology and topology of VPNs. This ability includes being able to meet these objectives:

- Describe the connectivity of traditional router-based networks
- Describe the advantages of VPN connectivity as compared to traditional router-based networks
- Identify the two major VPN implementation models
  - Describe the characteristics and technologies of overlay VPNs
  - Describe the characteristics and technologies of peer-to-peer VPNs
  - Describe the benefits of each type of VPN model
  - Describe the drawbacks of each VPN model

Traditional Router-Based Network Connectivity

This topic describes the connectivity of traditional router-based networks.

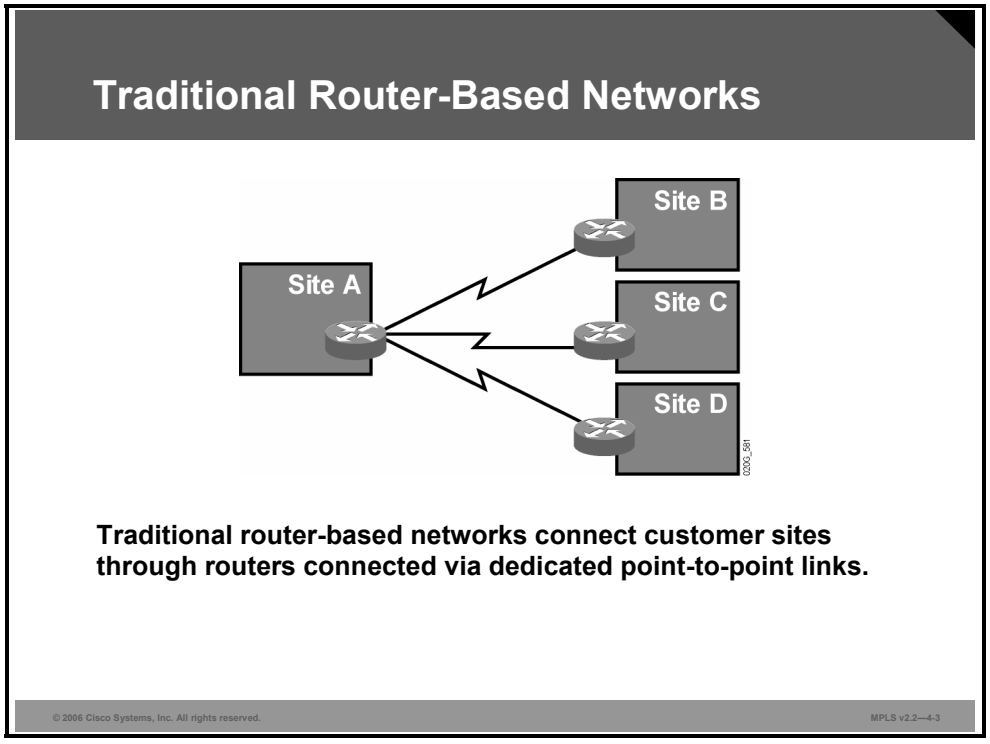

Traditional router-based networks were implemented with dedicated point-to-point links connecting customer sites. The cost of this approach was comparatively high for these reasons:

- The dedicated point-to-point links prevented any form of statistical infrastructure sharing on the service provider side, resulting in high costs for the end user.
- Every link required a dedicated port on a router, resulting in high equipment costs.

Advantages of VPNs

This topic describes the advantages of VPN connectivity as compared to traditional router-based networks.

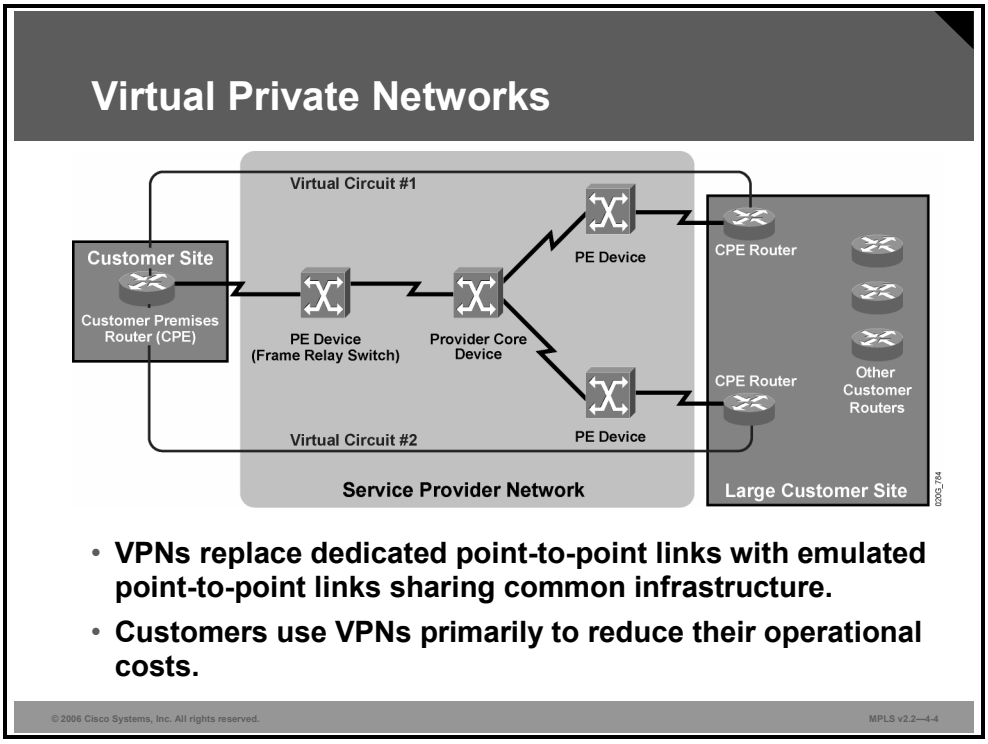

VPNs were introduced very early in the history of data communications with technologies such as X.25 and Frame Relay, which use virtual circuits to establish the end-to-end connection over a shared service provider infrastructure. Although X.25 and Frame Relay are sometimes considered legacy technologies and obsolete, they still share these basic benefits with modern VPNs:

- The dedicated links of traditional router-based networks have been replaced with a common infrastructure that emulates point-to-point links for the customer, resulting in statistical sharing of the service provider infrastructure.
- Statistical sharing of the infrastructure enables the service provider to offer connectivity for a lower price, resulting in lower operational costs for the end user.

Example: VPNs

The figure shows the statistical sharing, where the customer premises equipment (CPE) router on the left has one physical connection to the provider edge (PE) device, and two virtual circuits have been provisioned. Virtual circuit #1 provides connectivity to the top CPE router on the right. Virtual circuit #2 provides connectivity to the bottom CPE router on the right.

VPN Terminology

This topic identifies the terminology used in describing VPNs.

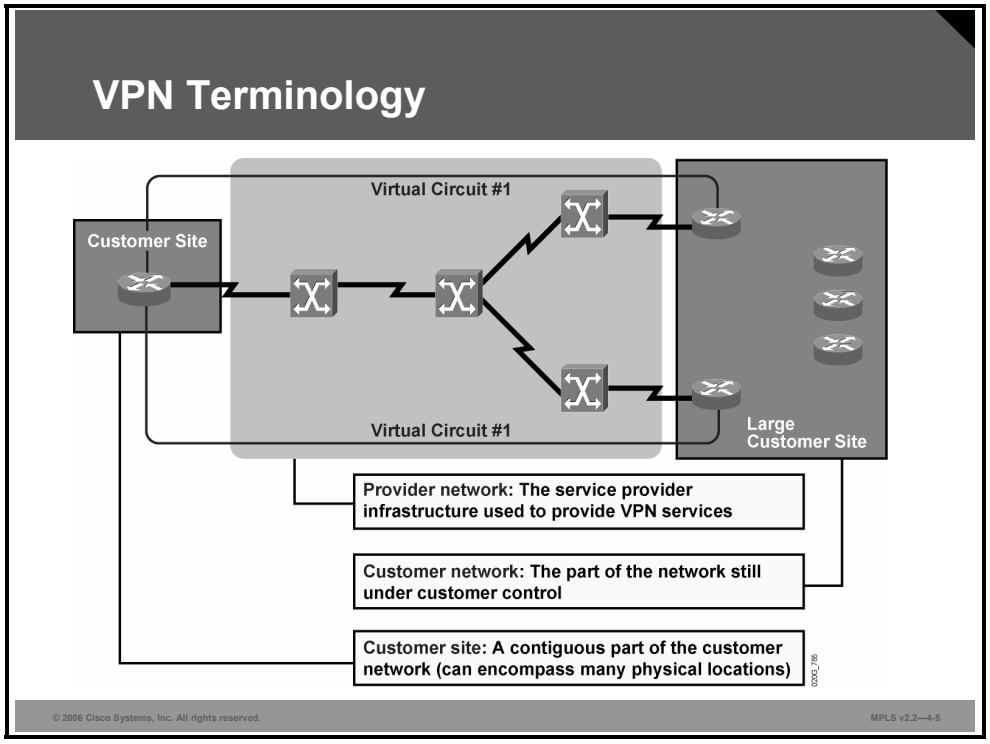

There are many conceptual models and terminologies describing various VPN technologies and implementations. The terminology is generic enough to cover nearly any VPN technology or implementation and is thus extremely versatile.

The major parts of an overall VPN solution are always those listed here:

- **Provider network (P-network):** The common infrastructure that the service provider uses to offer VPN services to customers
- **Customer network (C-network):** The part of the overall customer network that is still exclusively under customer control
- **Customer sites:** Contiguous parts of the C-network

A typical C-network implemented with any VPN technology would contain islands of connectivity under customer control (customer sites) connected together via the service provider infrastructure (P-network).

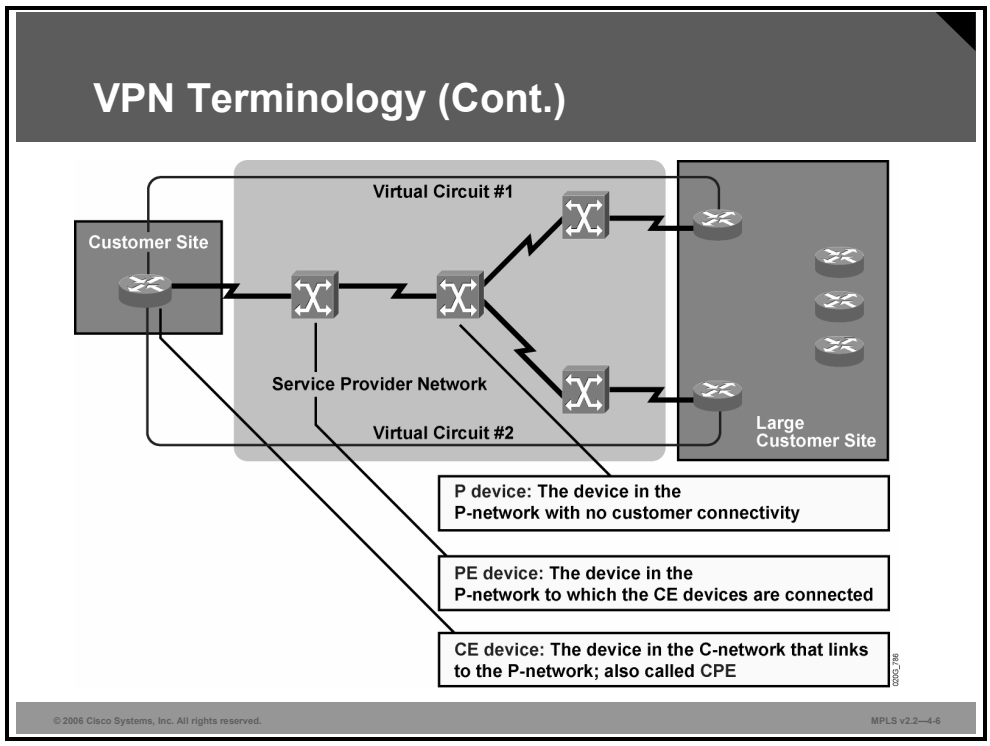

Here is a description of the devices that enable the overall VPN solution, which are named based on their position in the network:

- The customer router that connects the customer site to the service provider network is called a customer edge (CE) router, or CE device. Traditionally, this device is called CPE.
- Service provider devices to which customer devices are attached are called PE devices. In traditional switched WAN implementations, these devices would be Frame Relay or X.25 edge switches. In an MPLS implementation, these devices would be the edge label switch routers (edge LSRs).
- Service provider devices that provide only data transport across the service provider backbone, and have no customers attached to them, are called provider (P) devices. In traditional switched WAN implementations, these devices would be core (or transit) switches. In an MPLS implementation, these devices would be the LSRs.

Note The connecting device is still called a CE device even if it is not a router. For example, a packet assembler/disassembler (PAD) is a CE device.

What Are the VPN Implementation Models?

This topic describes the two major VPN implementation models.

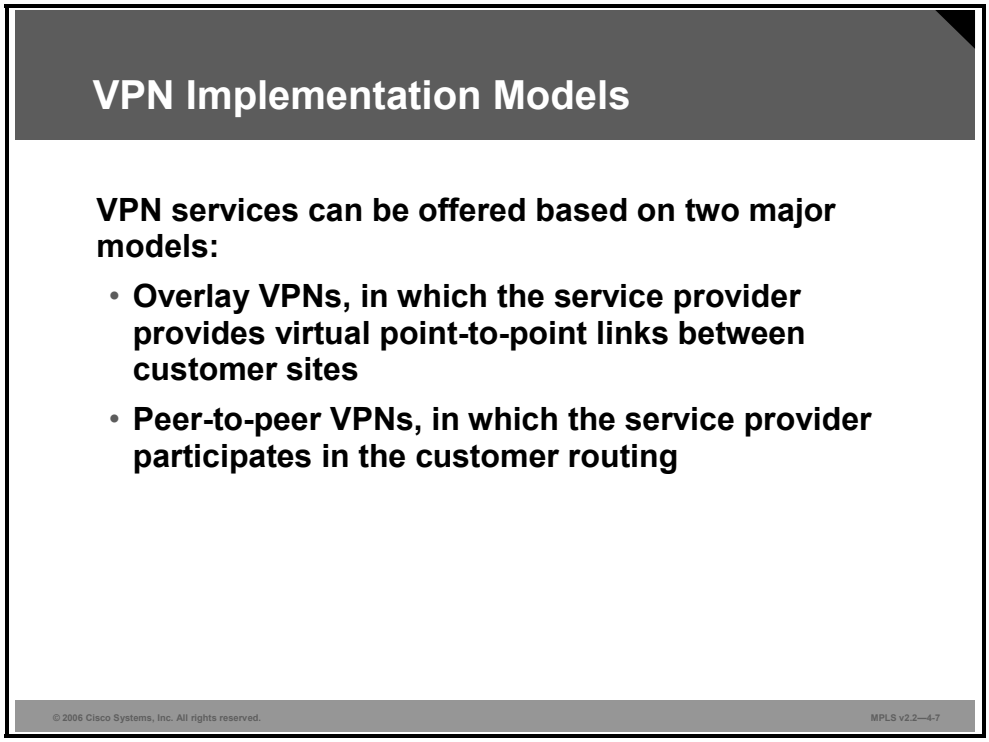

Traditional VPN implementations were all based on the overlay model, in which the service provider sold virtual circuits between customer sites as a replacement for dedicated point-to-point links. The overlay model had a number of drawbacks, which are identified in this lesson. To overcome these drawbacks (particularly in IP-based customer networks), a new model called the peer-to-peer VPN was introduced. In the peer-to-peer VPN model, the service provider actively participates in customer routing.

What Are Overlay VPN Technologies?

This topic describes the characteristics and technologies of overlay VPNs.

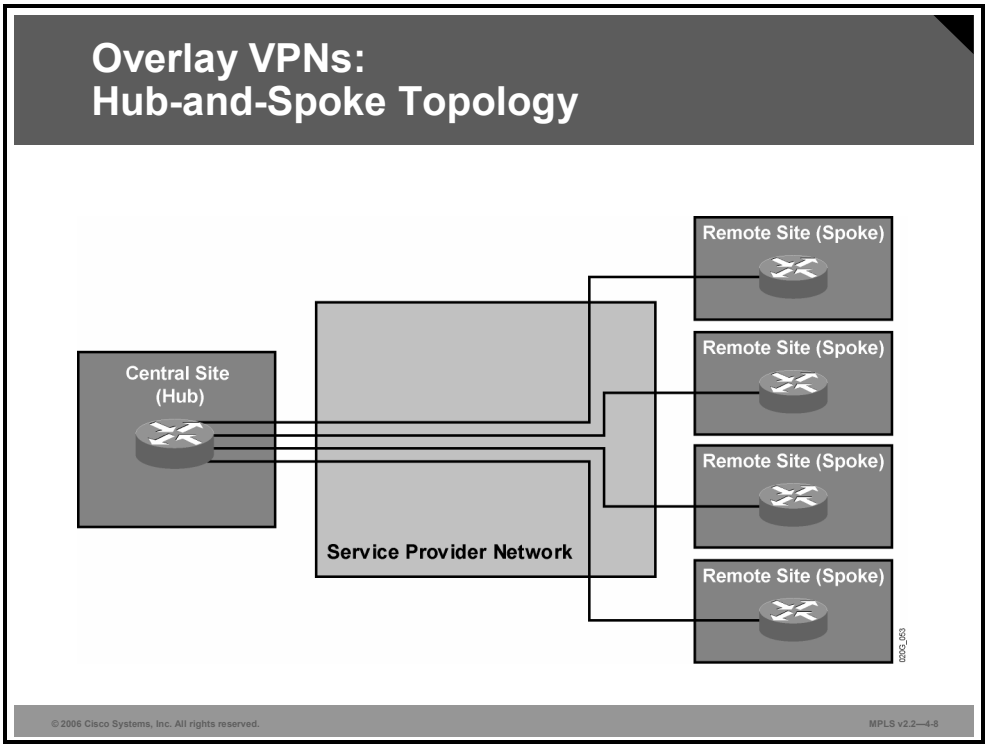

The hub-and-spoke topology is the simplest overlay VPN topology—all remote sites are linked with a single virtual circuit to a central CE router. The routing is also extremely simple—static routing or a distance vector protocol such as Routing Information Protocol (RIP) is more than adequate. If a dynamic routing protocol such as RIP is used, split-horizon updates must be disabled at the hub router or point-to-point subinterfaces must be used at the hub router to overcome the split-horizon problem.

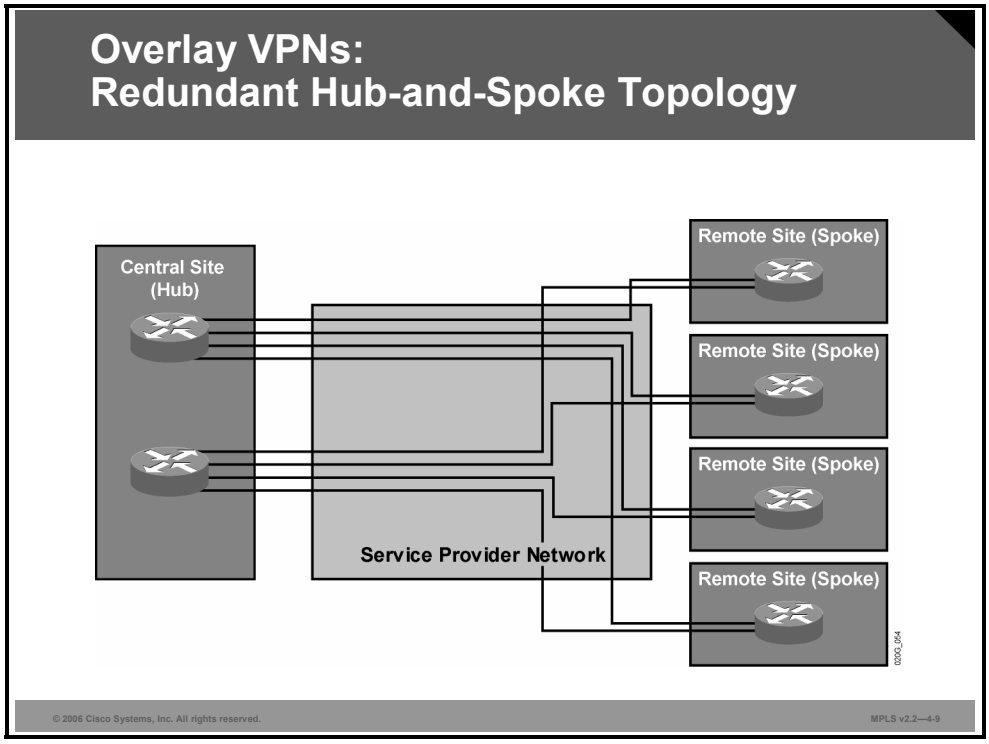

A typical redundant hub-and-spoke topology introduces central site redundancy (more complex topologies might also introduce router redundancy at spokes).

Each remote site is linked with two central routers via two virtual circuits. The two virtual circuits can be used for load sharing or in a primary circuit with backup circuit configuration.

Overlay VPNs: Layer 2 Implementation

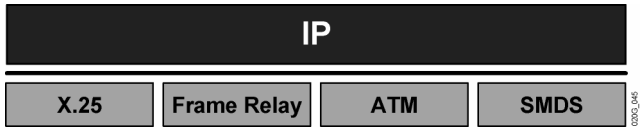

This is the traditional switched WAN solution:

- **The service provider establishes Layer 2 virtual circuits between customer sites.**
- **The customer is responsible for all higher layers.**

A Layer 2 overlay VPN implementation is the traditional switched WAN model, implemented with technologies such as X.25, Frame Relay, ATM, and Switched Multimegabit Data Service (SMDS). The service provider is responsible for transport of Layer 2 frames between customer sites, and the customer is responsible for all higher layers.

Overlay VPNs: IP Tunneling

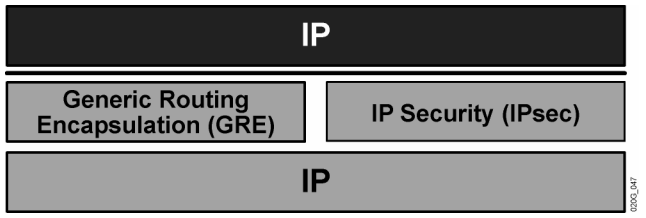

VPN is implemented with IP-over-IP tunnels:

- **Tunnels are established with GRE or IPsec.**
- **GRE is simpler (and quicker); IPsec provides authentication and security.**

With the success of IP and associated technologies, some service providers started to implement pure IP backbones to offer VPN services based on IP. In other cases, customers wanted to take advantage of the low cost and universal availability of the Internet to build low-cost private networks over it.

Whatever the business reasons behind it, Layer 3 VPN implementations over the IP backbone always involve tunneling—encapsulation of protocol units at a certain layer of the Open Systems Interconnection (OSI) reference model into protocol units at the same or higher layer of the OSI model.

Two well-known tunneling technologies are IP Security (IPsec) and generic routing encapsulation (GRE). GRE is fast and simple to implement and supports multiple routed protocols, but it provides no security and is thus unsuitable for deployment over the Internet. An alternative tunneling technology is IPsec, which provides network layer authentication and optional encryption to make data transfer over the Internet secure. IPsec supports only the IP routed protocol.

Overlay VPNs: Layer 2 Forwarding

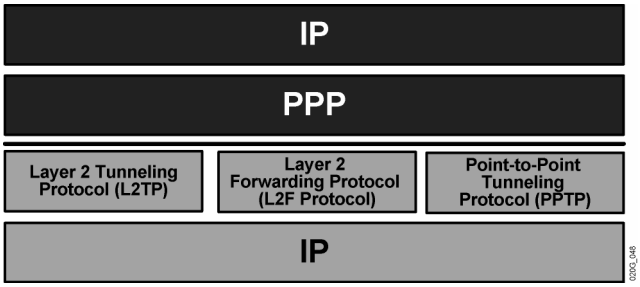

- **VPN is implemented with PPP-over-IP tunnels.**
- **VPN is usually used in access environments (dialup, digital subscriber line).**

Yet another tunneling technique was first implemented in dialup networks, where service providers wanted to tunnel customer dialup data encapsulated in PPP frames over an IP backbone to the customer central site. To make the service provider transport transparent to the customer, PPP frames are exchanged between the customer sites (usually a dialup user and a central site) and the customer is responsible for establishing Layer 3 connectivity above PPP.

Here are three well-known PPP forwarding implementations:

- Layer 2 Forwarding Protocol (L2F Protocol)
- Layer 2 Tunneling Protocol (L2TP)
- Point-to-Point Tunneling Protocol (PPTP)

Overlay VPNs: Layer 3 Routing

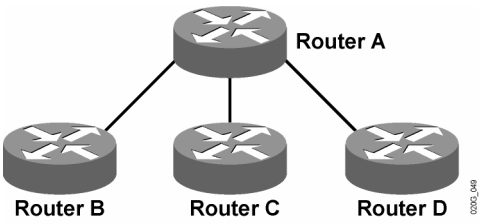

- **The service provider infrastructure appears as point-to-point links to customer routes.**
- **Routing protocols run directly between customer routers.**
- **The service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.**

From the Layer 3 perspective, the P-network is invisible to the customer routers, which are linked with emulated point-to-point links. The routing protocol runs directly between customer routers that establish routing adjacencies and exchange routing information.

The service provider is not aware of customer routing and has no information about customer routes. The responsibility of the service provider is purely the point-to-point data transport between customer sites.

The overlay VPN model has a number of drawbacks, most significantly the need for customers to establish point-to-point links or virtual circuits between sites. The formula to calculate how many point-to-point links or virtual circuits are needed in the worst case is ([*n*][*n*-1])/2, where *n* is the number of sites to be connected. For example, if you need to have full mesh connectivity between four sites, you will need a total of six point-to-point links or virtual circuits: (4 * (4-1)) / 2. This leads to scalability issues.

What Are Peer-to-Peer VPN Technologies?

This topic describes the characteristics and technologies of peer-to-peer VPNs.

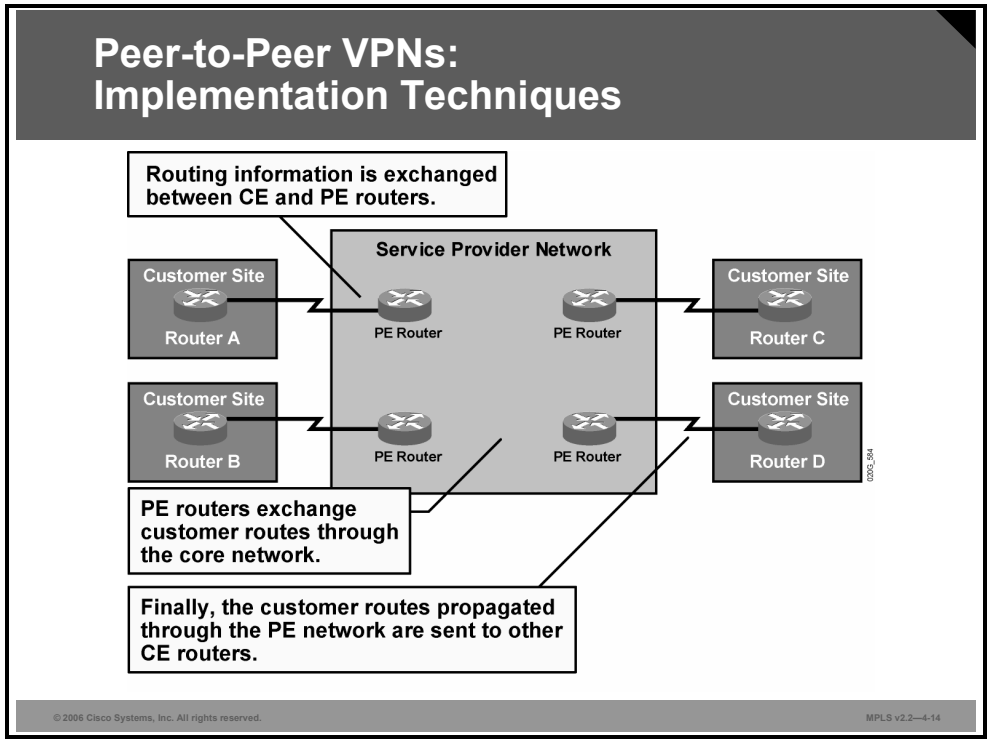

To overcome the scalability issue and provide the customer with optimum data transport across the service provider backbone, the peer-to-peer VPN concept was introduced. Here, the service provider actively participates in customer routing, accepting customer routes, transporting those customer routes across the service provider backbone, and finally propagating them to other customer sites.

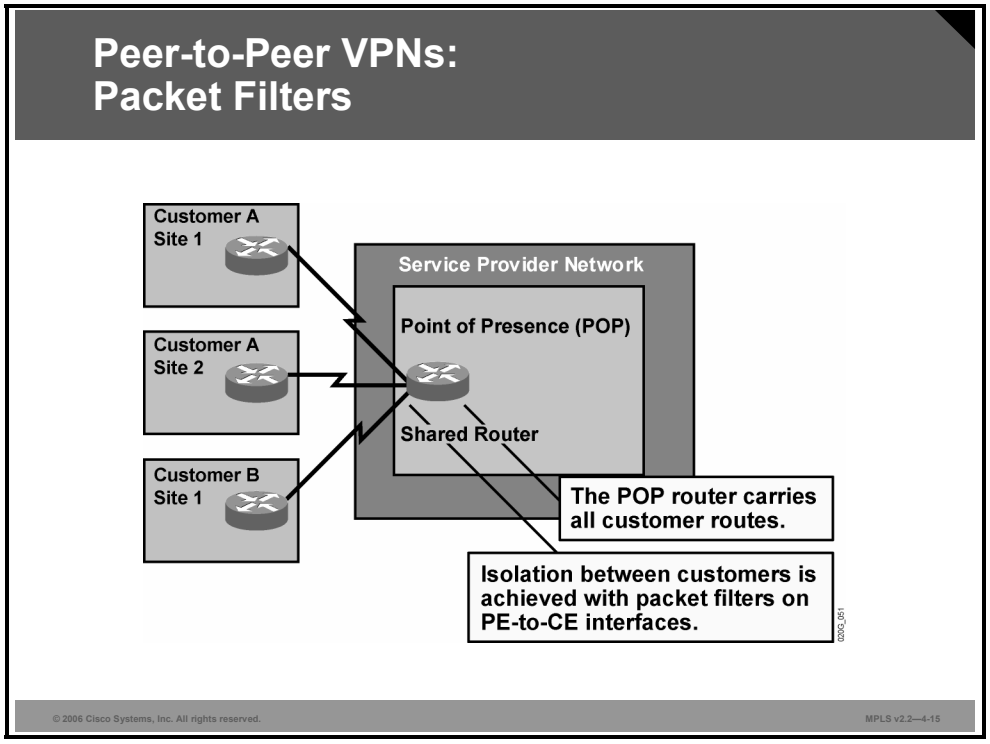

The first peer-to-peer VPN solutions appeared with the widespread deployment of IP in service provider networks. Architectures similar to that of the Internet were used to build them. Special provisions were taken into account to transform the architecture, which was targeted toward public backbones (Internet), into a solution in which customers would be totally isolated and be able to exchange corporate data securely.

The more common peer-to-peer VPN implementation allowed a PE router to be shared between two or more customers. Packet filters were used on the shared PE routers to isolate the customers. In this implementation, it was common for the service provider to allocate a portion of its address space to each customer and manage the packet filters on the PE routers to ensure full reachability between sites of a single customer and isolation between separate customers.

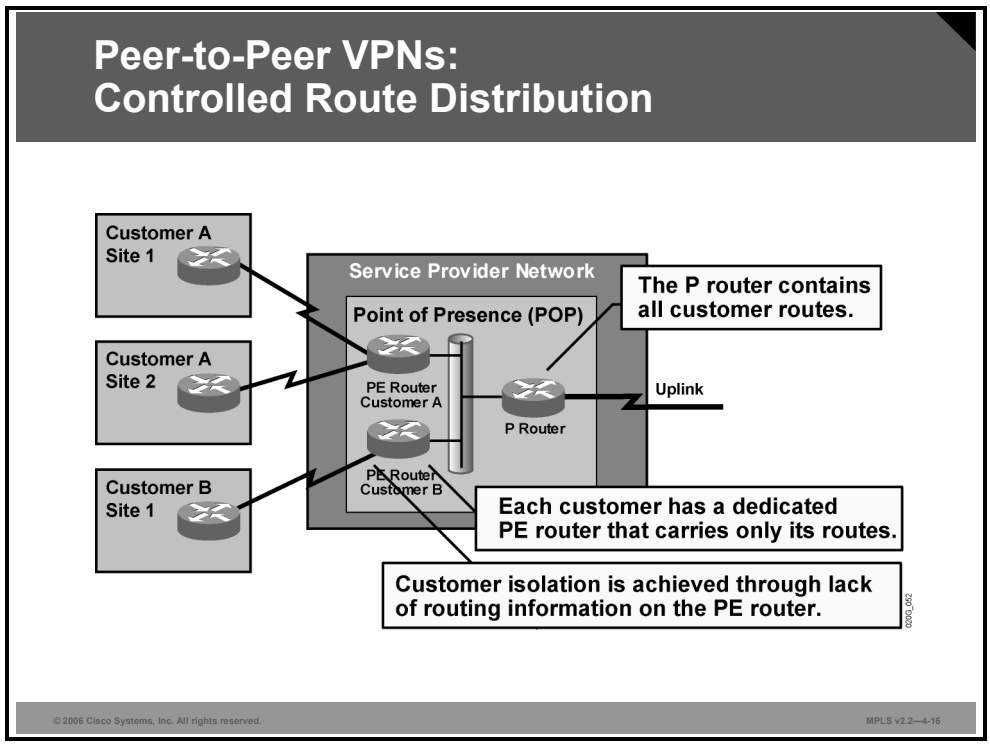

Maintaining packet filters is a mundane and error-prone task. Some service providers have thus implemented more innovative solutions based on controlled route distribution. In this approach, the customer has a dedicated PE router. The core service provider (P) routers contain all customer routes, and the dedicated PE routers contain only the routes of a single customer. This approach requires a dedicated PE router per customer per point of presence (POP). Customer isolation is achieved solely through lack of routing information on the PE router.

Example: Controlled Route Distribution

In the figure, the PE router for customer A, using route filtering between the P router and the PE routers, learns only routes belonging to customer A, and the PE router for customer B learns only routes belonging to customer B. Border Gateway Protocol (BGP) with BGP communities is usually used inside the provider backbone, because it offers the most versatile route-filtering tools.

Note	Default routes used anywhere in the C-network or P-network break isolation between customers and have to be avoided.

What Are the Benefits of VPNs?

This topic describes the benefits of the two VPN models.

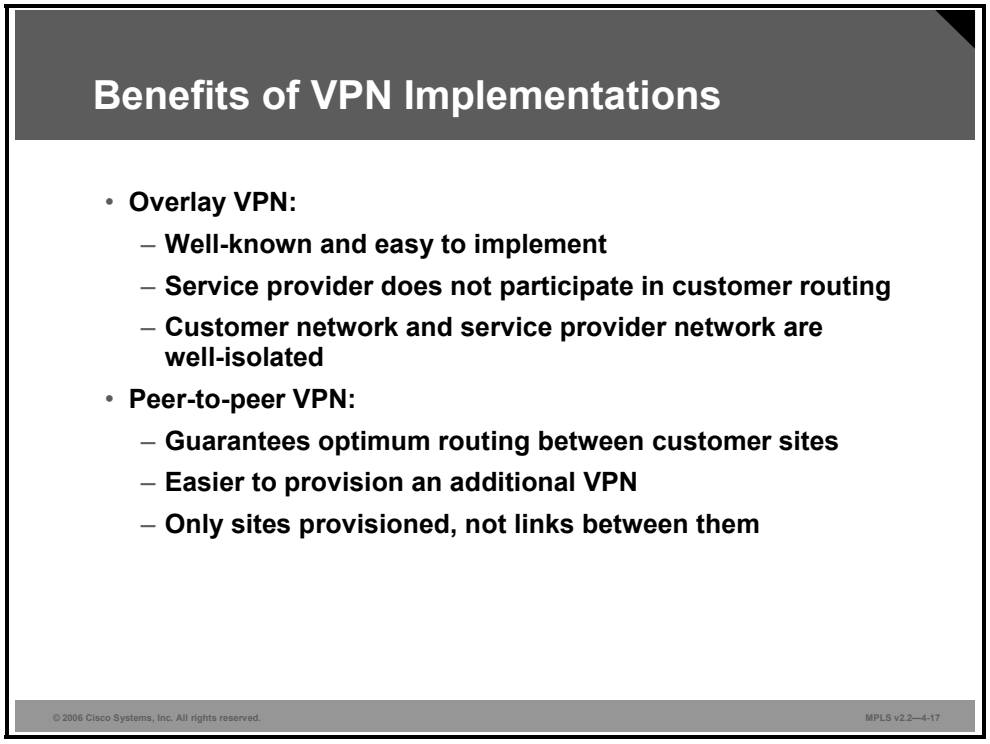

Each VPN model has a number of benefits. For example, overlay VPNs have these advantages:

- Overlay VPNs are well-known and easy to implement from both customer and service provider perspectives.
- The service provider does not participate in customer routing, making the demarcation point between service provider and customer easier to manage.

On the other hand, peer-to-peer VPNs provide these advantages:

- Optimum routing between customer sites without any special design or configuration effort
- Easy provisioning of additional VPNs or customer sites, because the service provider provisions only individual sites, not the links between individual customer sites

What Are the Drawbacks of VPNs?

This topic describes the drawbacks of the VPN models.

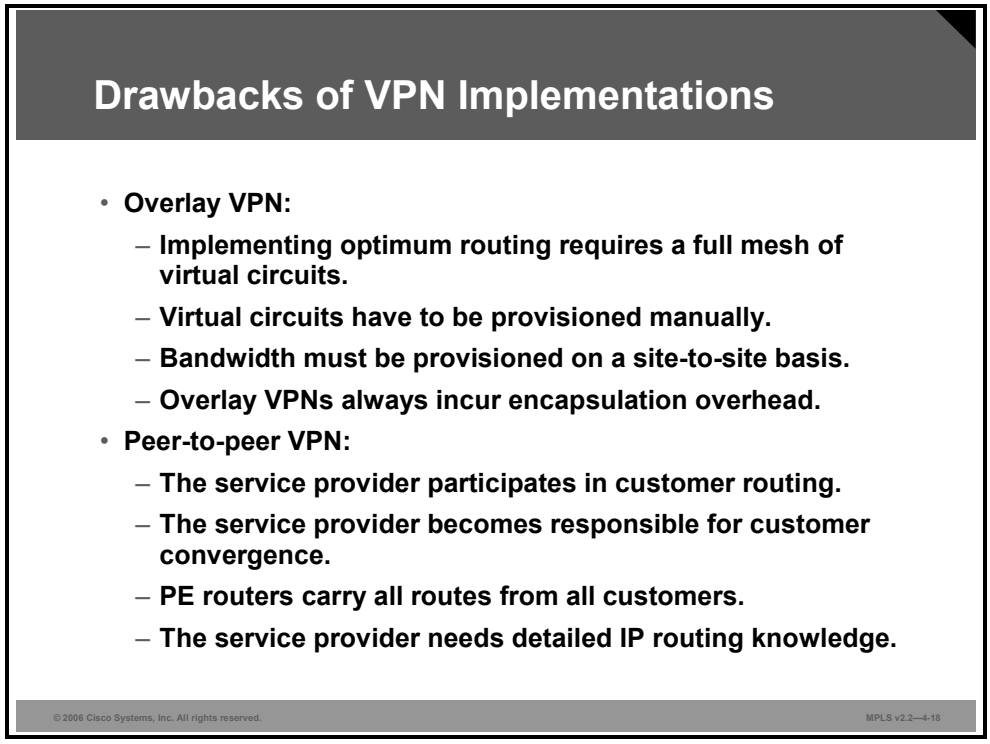

Each VPN model also has a number of drawbacks. Overlay VPNs have these disadvantages:

- Overlay VPNs require a full mesh of virtual circuits between customer sites to provide optimum intersite routing.
- All virtual circuits between customer sites have to be provisioned manually, and the bandwidth must be provisioned on a site-to-site basis (which is not always easy to achieve).
- The IP-based overlay VPN implementations (with IPsec or GRE) incur high encapsulation overhead—ranging from 20 bytes to 80 bytes per transported datagram.

The major drawbacks of peer-to-peer VPNs arise from service provider involvement in customer routing, such as these disadvantages:

- The service provider becomes responsible for correct customer routing and for fast convergence of the C-network following a link failure.
- The service provider PE routers have to carry all customer routes that were hidden from the service provider in the overlay VPN model.
- The service provider needs detailed IP routing knowledge, which is not readily available in traditional service provider teams.

Summary

This topic summarizes the key points that were discussed in this lesson.

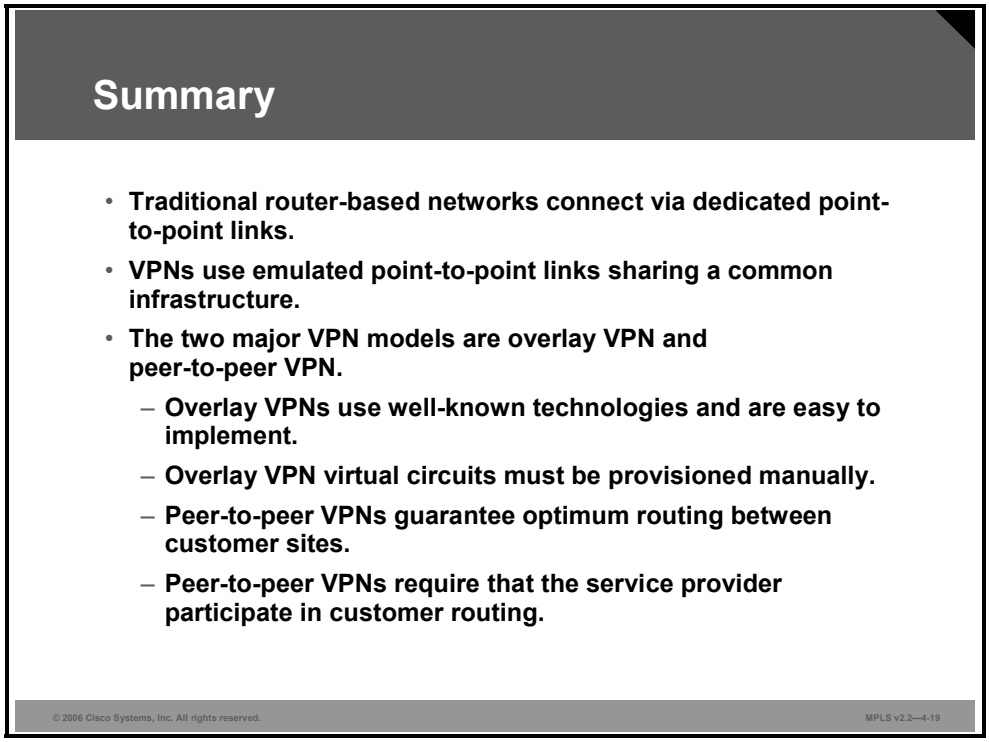

Lesson 2

Categorizing VPNs

Overview

This lesson explains how Virtual Private Networks (VPNs) can be categorized based on business needs or connectivity requirements.

It is important to understand the different categories of VPNs and to know into which environments those VPNs can be applied.

Objectives

Upon completing this lesson, you will be able to describe the characteristics of the different VPN topology categories. This ability includes being able to meet these objectives:

- Identify the major business categories for VPNs
- Describe the characteristics of extranet VPNs
- Identify the major connectivity categories for VPNs
- Describe the characteristics of central services extranet VPNs
- Describe the characteristics of managed network VPNs

What Are the Business Categories for VPNs?

This topic describes how VPNs can be categorized based on business needs.

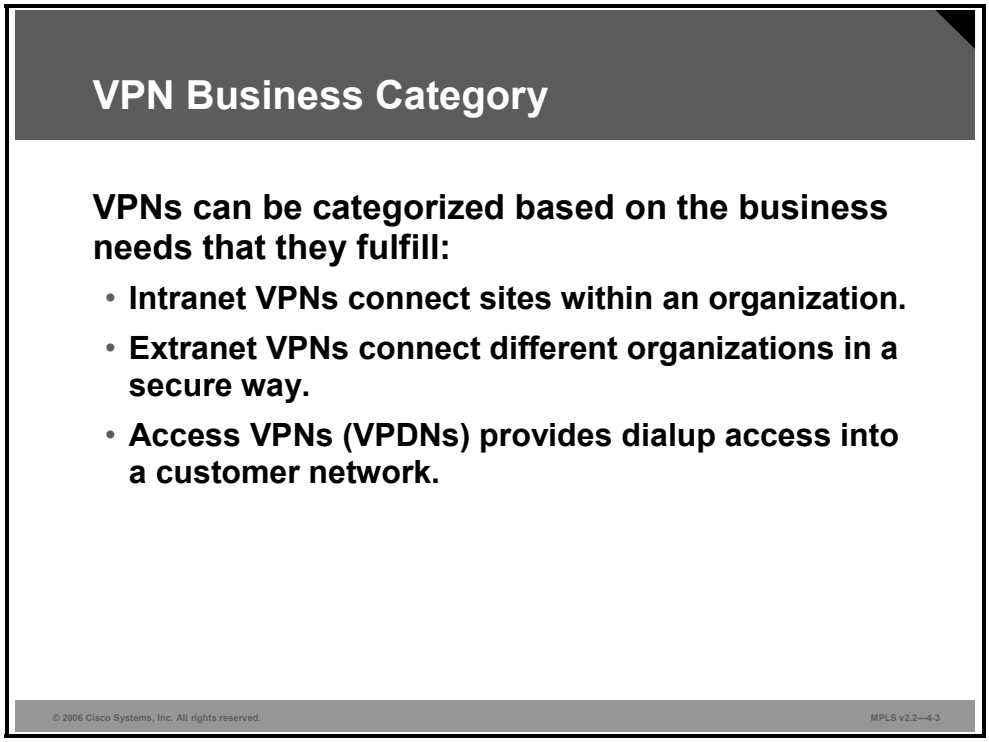

Here is a list of some very popular VPN categories that classify VPNs based on the business needs that they fulfill:

- **Intranet VPN:** Intranet VPNs connect sites within an organization. Security mechanisms are usually not deployed in an intranet, because all sites belong to the same organization.
- **Extranet VPN:** Extranet VPNs connect different organizations. Extranets usually rely on security mechanisms to ensure the protection of participating individual organizations. Security mechanisms are usually the responsibility of individual participating organizations.
- **Access VPN:** Access VPNs are virtual private dial-up networks (VPDNs) that provide dialup access into a customer network.

What Are Extranet VPNs?

This topic describes the characteristics of extranet VPNs.

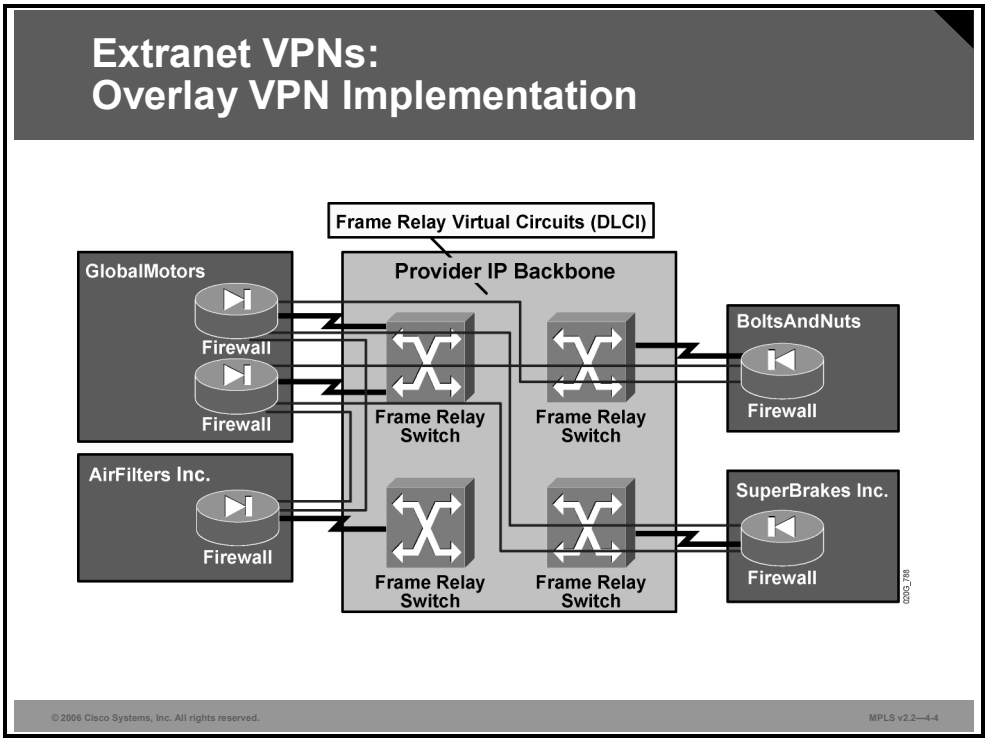

In an overlay implementation of an extranet, organizations are linked with dedicated virtual circuits.

Example: Overlay VPN—Extranet VPNs

This figure illustrates an overlay VPN implementation of an extranet. Traffic between two organizations can flow only if one of these conditions is met:

- There is a direct virtual circuit between the organizations.
- A third organization linked with both organizations is willing to provide transit traffic capability to those organizations. Because establishing virtual circuits between two organizations is always associated with costs, the transit traffic capability is almost never granted free of charge.

Example: Peer-to-Peer VPN—Extranet VPNs

This figure illustrates a peer-to-peer VPN implementation of an extranet.

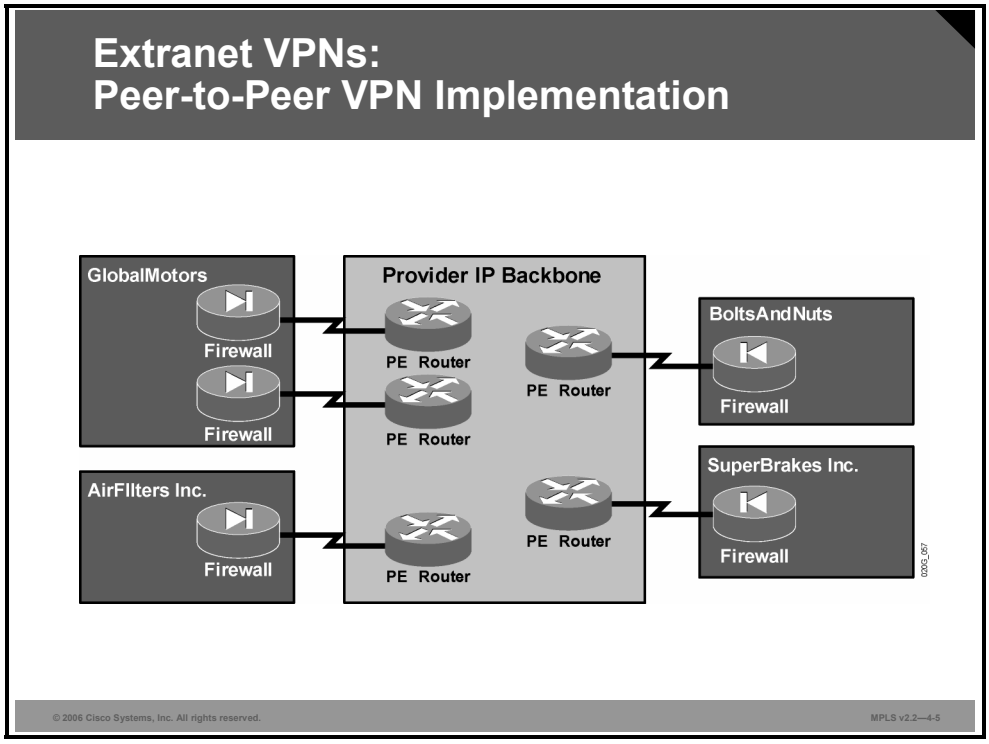

Peer-to-peer VPN implementation of an extranet VPN is very simple compared with overlay VPN implementation—all sites are connected to the provider network (P-network), and optimum routing between sites is enabled by default.

The cost model of peer-to-peer implementation is also simpler—usually every organization pays its connectivity fees for participation in the extranet and gets full connectivity to all other sites.

What Are the Connectivity Categories for VPNs?

This topic identifies the major connectivity categories for VPNs.

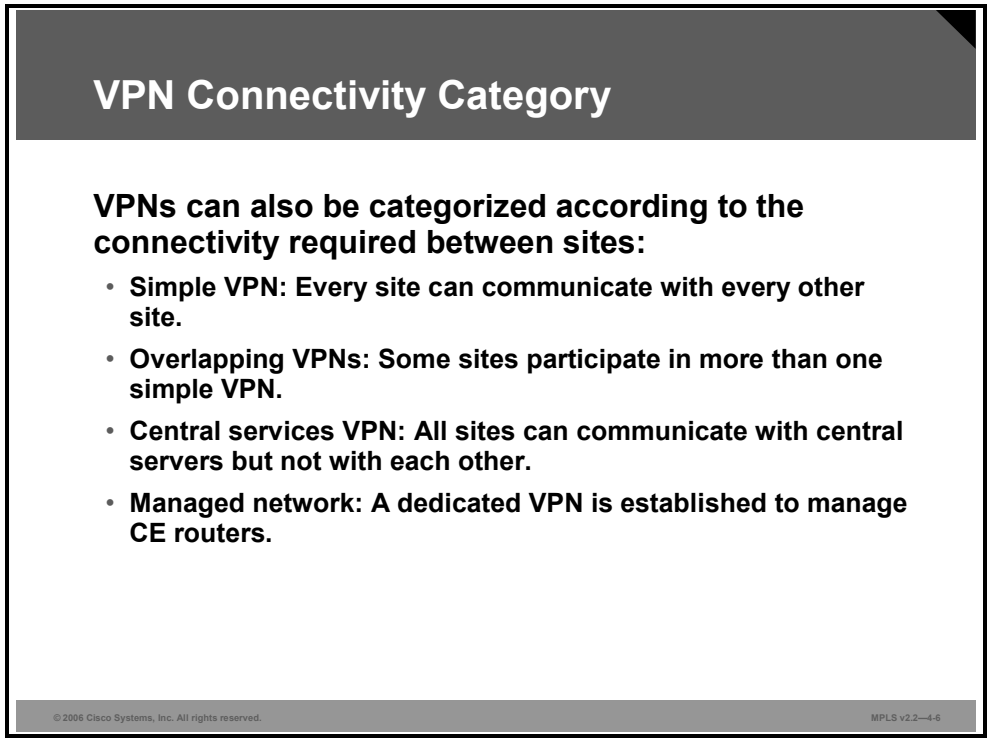

The VPNs discussed so far have usually been very simple in terms of connectivity, as described here:

- In most cases, full connectivity between sites is required. (In an overlay implementation of either an intranet or extranet VPN, this requirement usually means that a common site acts as a transit site).
- In an overlay implementation of an extranet VPN, the connectivity is limited to sites that have direct virtual circuits established between them.

Here are descriptions of a number of advanced VPN topologies with more complex connectivity requirements:

- Overlapping VPN connectivity, in which a site participates in more than one VPN
- Central services VPNs, in which the sites are split into two classes: server sites, which can communicate with all other sites, and client sites, which can communicate only with the servers, not with other clients
- Network management VPNs, which are used to manage customer edge (CE) devices in scenarios where the service provider owns and manages the devices

What Is the Central Services Extranet?

This topic describes the characteristics of central services extranet VPNs.

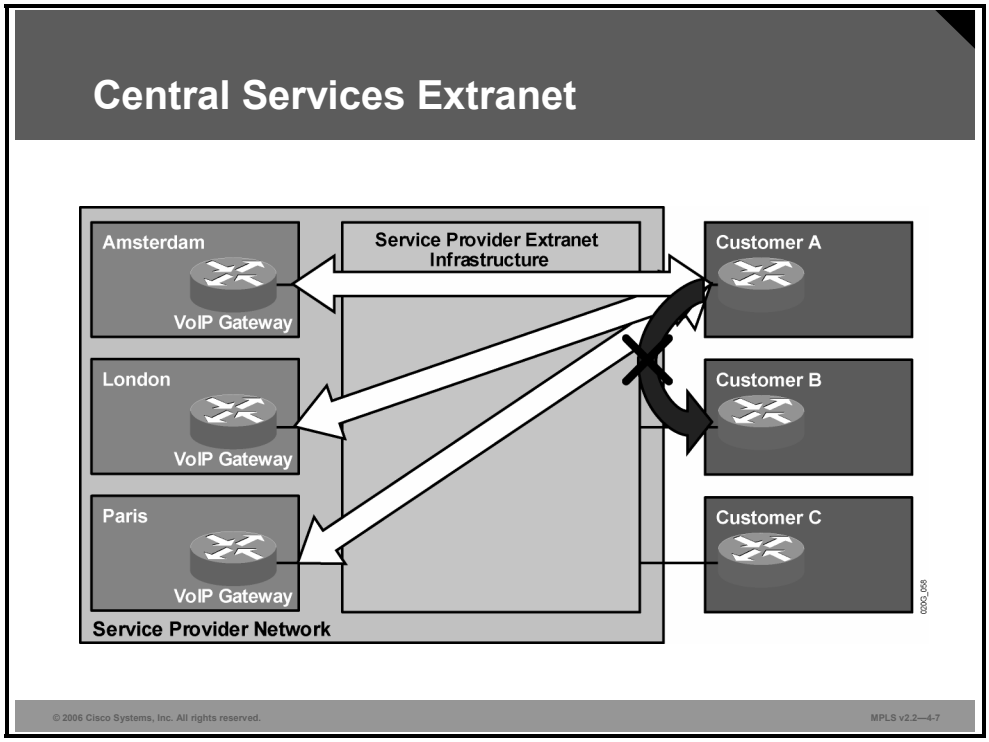

A service provider can integrate the business and connectivity attributes of VPNs to offer a central services extranet to its customers. For example, a service provider can provide international Voice over IP (VoIP) service.

Example: Central Services Extranet

The figure illustrates this example. Every customer of this service can access voice gateways in various countries but cannot access other customers using the same service.

What Is a Managed Network Implementation?

This topic describes the characteristics of managed network VPNs.

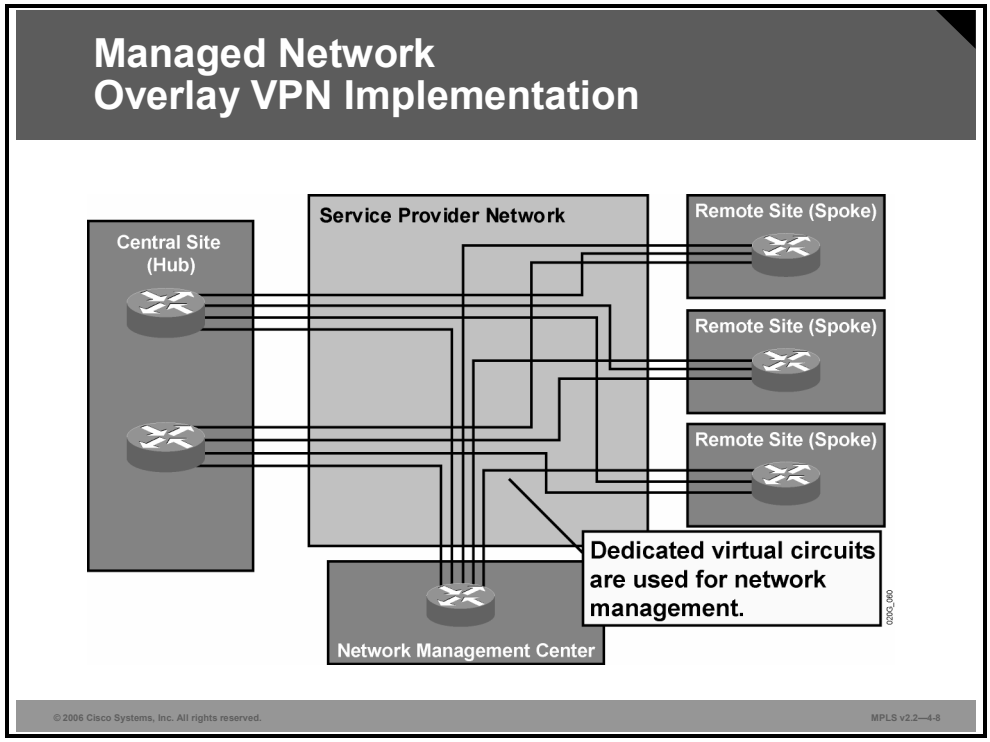

A managed network VPN is traditionally implemented in combination with overlay VPN services. Dedicated virtual circuits are deployed between any managed CE router and the central network management system (NMS) router to which the NMS is connected.

This managed network VPN implementation is sometimes called a “rainbow” implementation because the physical link between the NMS router and the core of the service provider network carries a number of virtual circuits—one circuit per managed router.

Example: Hybrid Implementation

The network diagram shows an interesting scenario in which a peer-to-peer VPN and an overlay VPN implementation can be used together to provide end-to-end service to the customer.

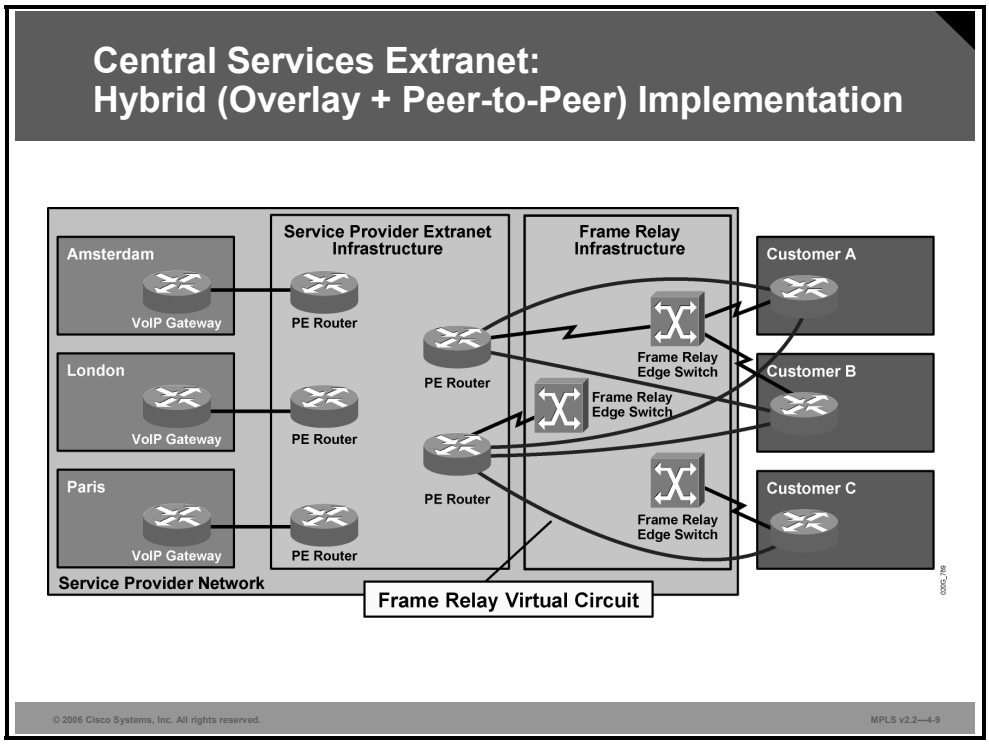

The VoIP service is implemented with a central services extranet topology, which is in turn implemented with a peer-to-peer VPN. Connectivity between provider edge (PE) routers in the peer-to-peer VPN and customer routers is implemented with an overlay VPN based on Frame Relay. The PE routers of the peer-to-peer VPN and the CE routers act as CE devices of the Frame Relay network.

Summary

This topic summarizes the key points that were discussed in this lesson.

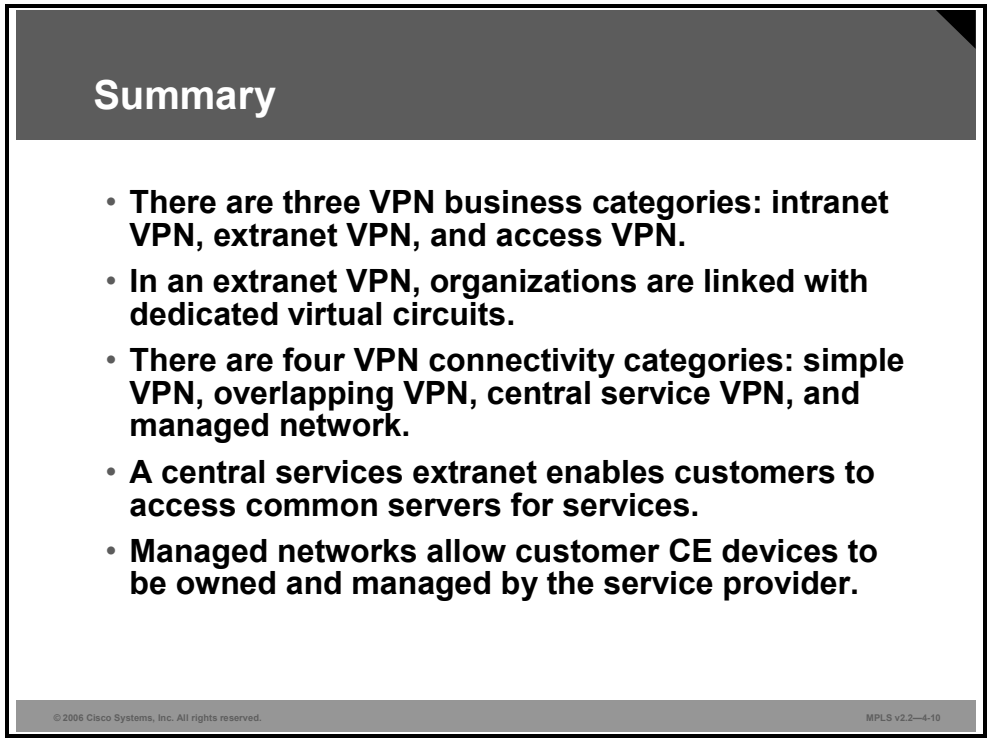

Lesson 3

Introducing MPLS VPN Architecture

Overview

This lesson explains the Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) architecture, route information propagation, route distinguishers (RDs), route targets (RTs), and virtual routing tables.

It is important to understand how the MPLS VPN architecture is structured, what the components of that architecture are, and how the components are used. This knowledge will help later when you begin to look at design issues and configuration parameters.

Objectives

Upon completing this lesson, you will be able to describe the major architectural components of an MPLS VPN. This ability includes being able to meet these objectives:

- Describe the drawbacks of traditional peer-to-peer VPNs
- Describe the features of the MPLS VPN architecture
- Describe the architecture of a PE router in an MPLS VPN
- Describe the different methods of propagating routing information across the provider network
- Describe the features of RDs
- Describe the features of RTs
- Describe how complex VPNs have redefined the meaning of VPNs
- Describe the impact of complex VPN topologies on virtual routing tables

What Are the Drawbacks of Traditional Peer-to-Peer VPNs?

This topic describes the drawbacks of the traditional peer-to-peer VPN implementation model.

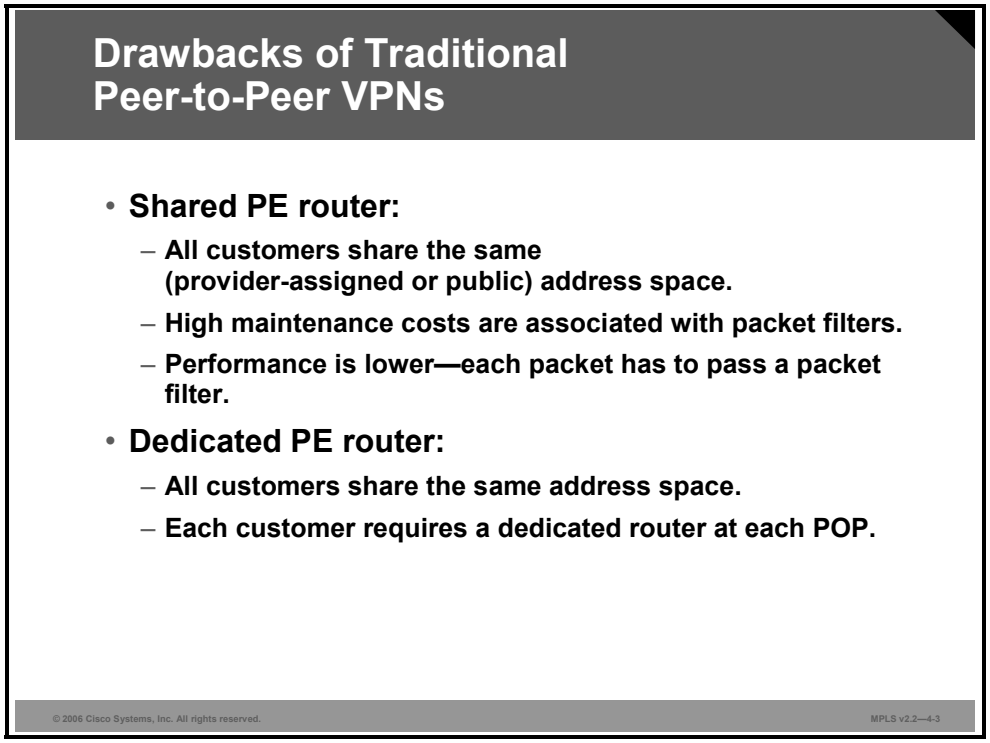

Pre-MPLS implementations of peer-to-peer VPNs all share a common drawback. Customers have to share the same global address space, either using their own public IP addresses or relying on provider-assigned IP addresses. In both cases, connecting a new customer to a peer-to-peer VPN service usually requires IP renumbering inside the customer network (C-network)—an operation most customers are reluctant to perform.

Peer-to-peer VPNs based on packet filters also incur high operational costs associated with packet filter maintenance and performance degradation because of heavy use of packet filters.

Peer-to-peer VPNs implemented with per-customer provider edge (PE) routers are easier to maintain and can provide optimum routing performance, but they are usually more expensive because every customer requires a dedicated router in every point of presence (POP). Thus, this approach is usually used if the service provider has only a small number of large customers.

What Is the MPLS VPN Architecture?

This topic describes the features of the MPLS VPN architecture.

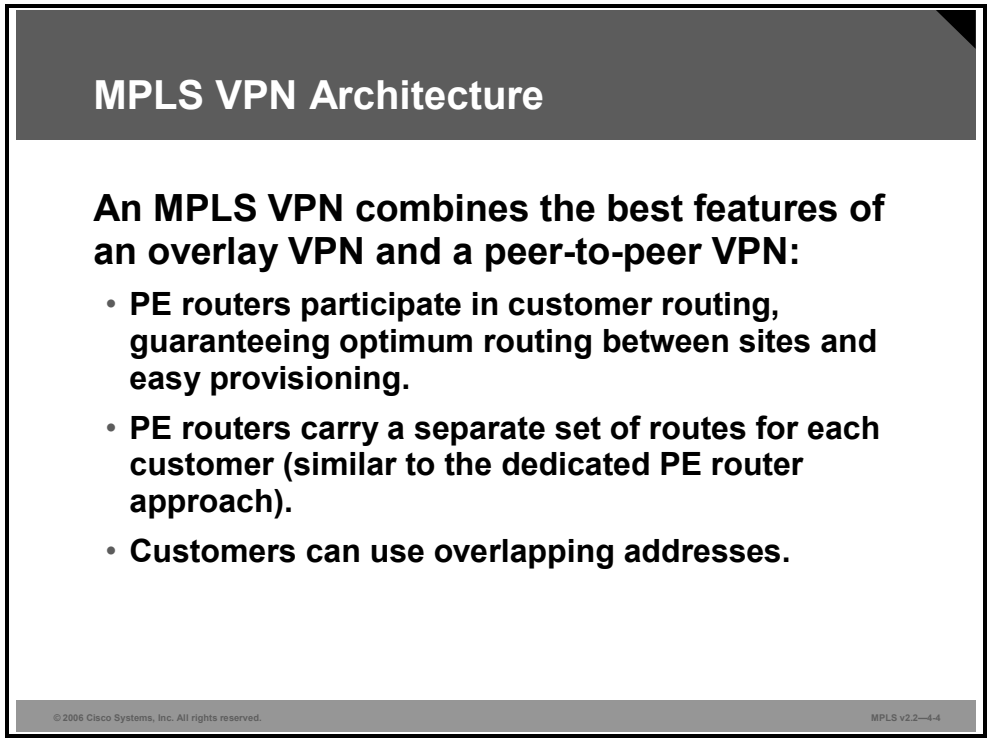

The MPLS VPN architecture offers service providers a peer-to-peer VPN architecture that combines the best features of overlay VPNs (support for overlapping customer address spaces) with the best features of peer-to-peer VPNs. This list describes these characteristics:

- PE routers participate in customer routing, guaranteeing optimum routing between customer sites.

Note	In an MPLS VPN implementation, the PE router is the edge label switch router (edge LSR).

- PE routers carry a separate set of routes for each customer, resulting in perfect isolation between customers.
- Customers can use overlapping addresses.

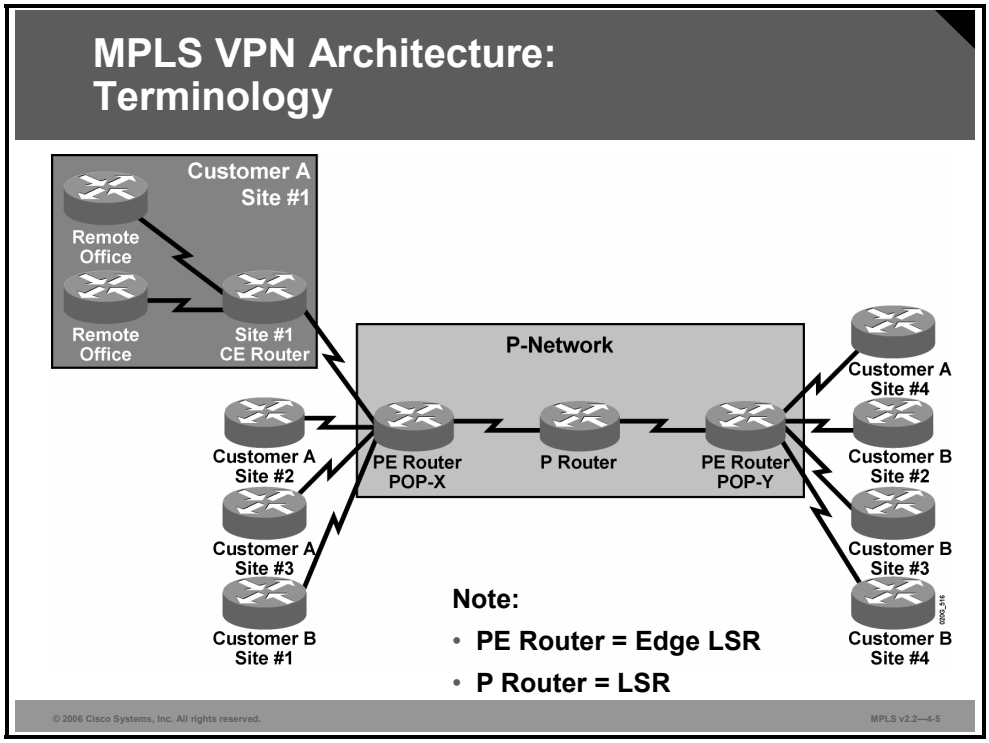

MPLS VPN terminology divides the overall network into a customer-controlled part (the C-network) and a provider-controlled part (the provider network [P-network]). Contiguous portions of the C-network are called sites and are linked with the P-network via customer edge (CE) routers. The CE routers are connected to the PE routers, which serve as the edge devices of the P-network. The core devices in the P-network, the provider routers (P routers), provide transit transport across the provider backbone and do not carry customer routes.

Note	In an MPLS VPN implementation, the P router is the LSR.

What Is the Architecture of a PE Router in an MPLS VPN?

This topic describes the architecture of a PE router in an MPLS VPN.

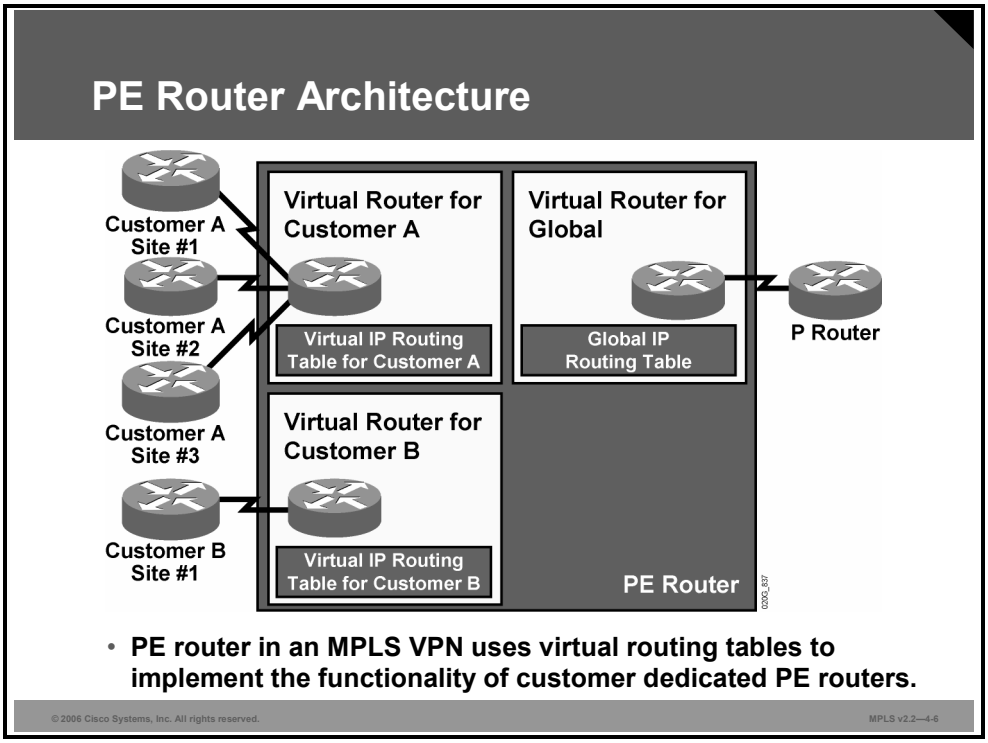

The architecture of a PE router in an MPLS VPN is very similar to the architecture of a POP with customer-dedicated PE routers used in the dedicated-router peer-to-peer VPN model. The only difference is that the whole architecture is condensed into one physical device with the PE router in an MPLS VPN. Each customer is assigned an independent routing table (virtual routing table or VRF) that corresponds to the customer dedicated PE router in the traditional peer-to-peer model. Routing across the provider backbone is performed by another routing process that uses a global IP routing table corresponding to the intra-POP P router in the traditional peer-to-peer model.

Note	Cisco IOS software implements isolation between customers via virtual routing and forwarding tables (VRF). The whole PE router is still configured and managed as a single device, not as a set of virtual routers.

What Are the Methods of Propagating Routing Information Across the P-Network?

This topic describes the different methods of propagating routing information across the P-network.

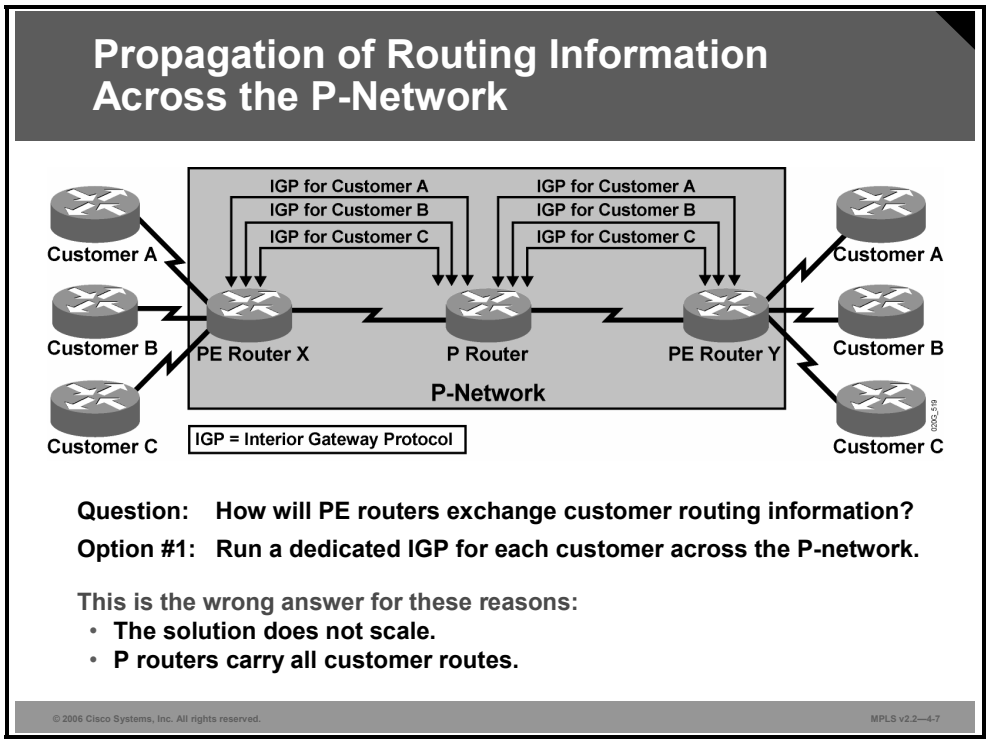

Although virtual routing tables provide isolation between customers, the data from these routing tables still needs to be exchanged between PE routers to enable data transfer between sites attached to different PE routers. Therefore, a routing protocol is needed that will transport all customer routes across the P-network, while maintaining the independence of individual customer address spaces.

An obvious solution, implemented by various VPN vendors, is to run a separate routing protocol for each customer. There are two common implementations that require that a per-customer routing protocol be run between PE routers:

1. The P routers participate in customer routing and pass the customer routing information between PE routers.
2. The PE routers are connected via point-to-point tunnels, for example IP Security (IPsec), thereby hiding the customer routing from the P routers.

The separate routing protocol for each customer is very simple to implement (and often used by some customers), but is not appropriate in service provider environments because it simply does not scale. The specific problems are as follows:

- The PE routers have to run a large number of routing protocols.
- The P routers have to carry all customer routes.

Propagation of Routing Information Across the P-Network (Cont.)

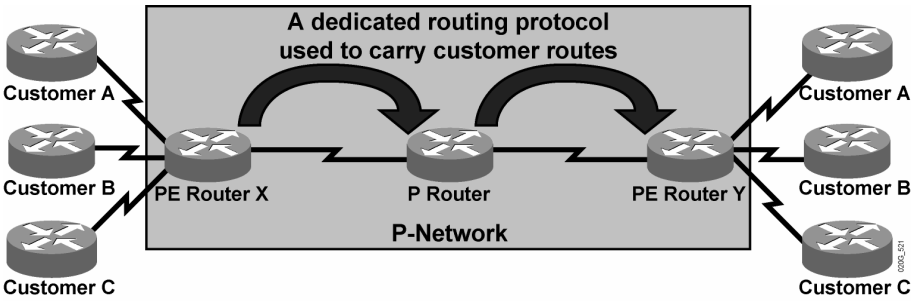

Question: **How will PE routers exchange customer routing information?**

Option #2: **Run a single routing protocol that will carry all customer routes inside the provider backbone.**

Better answer, but still not good enough:

- **P routers carry all customer routes.**

A better approach to the route propagation problem is to deploy a single routing protocol that can exchange all customer routes across the P-network. Although this approach is better than the previous one, the P routers are still involved in customer routing; therefore, the proposal retains some of the same scalability issues of the previous one.

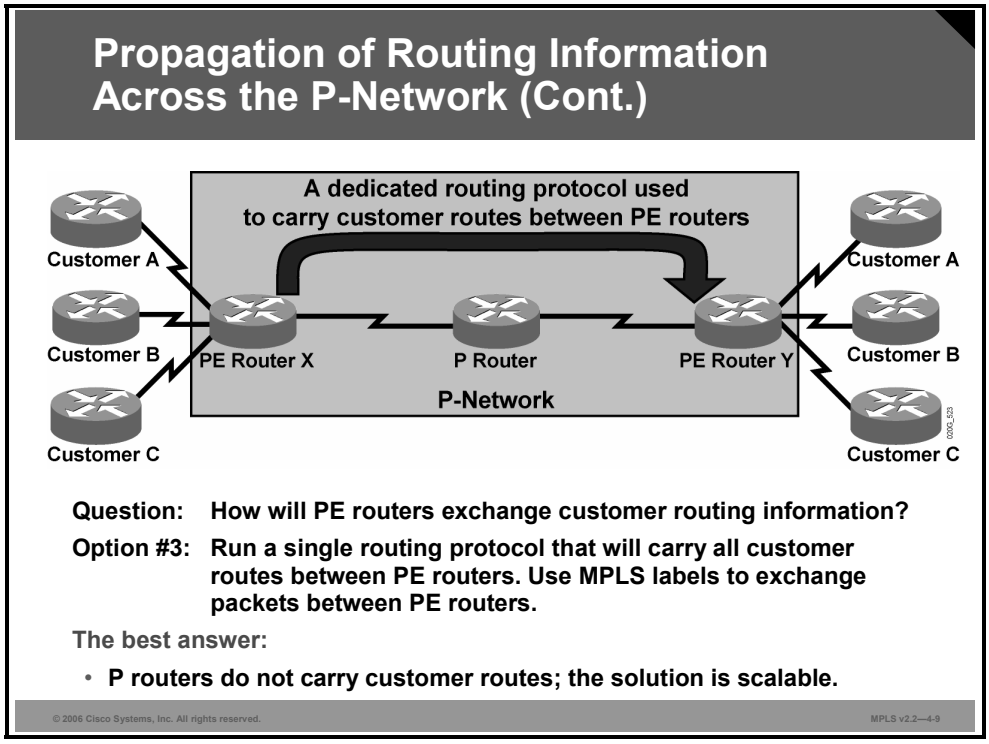

The best solution to the customer route propagation issue is to run a single routing protocol between PE routers that will exchange all customer routes without the involvement of the P routers. This solution is scalable. Some of the benefits of this approach are as follows:

- The number of routing protocols running between PE routers does not increase with an increasing number of customers.
- The P routers do not carry customer routes.

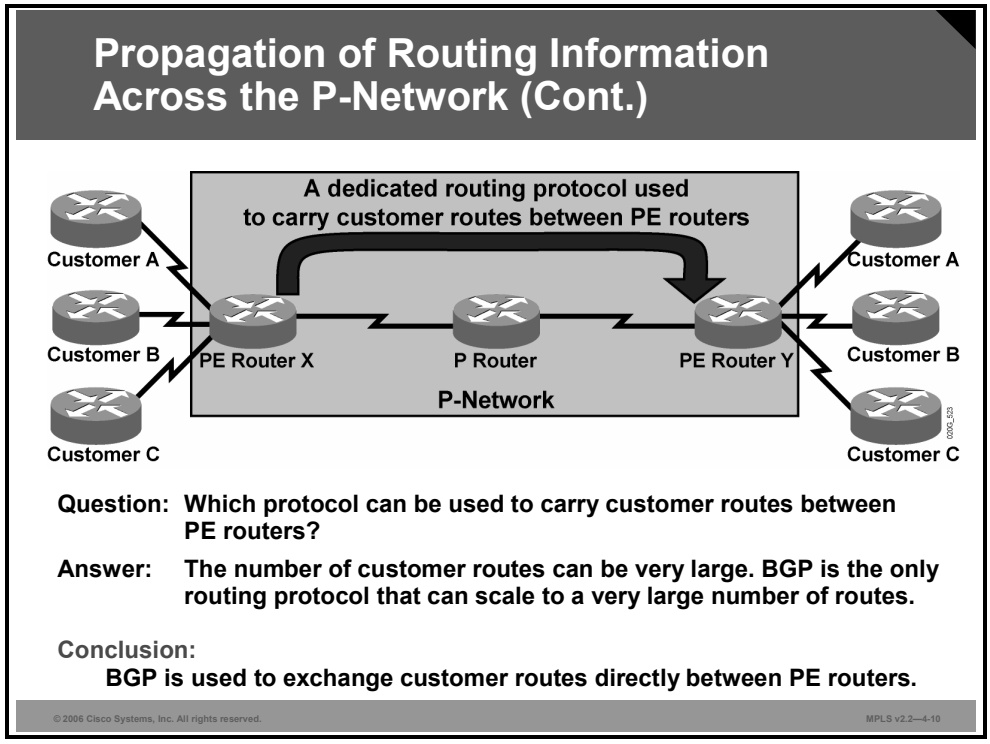

The next design decision to be made is the choice of the routing protocol running between PE routers. Given that the total number of customer routes is expected to be very large, the only well-known protocol with the required scalability is Border Gateway Protocol (BGP). In fact, BGP is used in the MPLS VPN architecture to transport customer routes directly between PE routers.

Propagation of Routing Information Across the P-Network (Cont.)

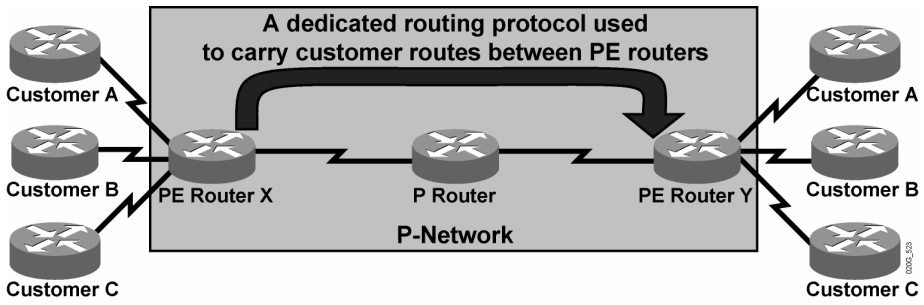

Question: How will information about the overlapping subnetworks of two customers be propagated via a single routing protocol?

Answer: Extend the customer addresses to make them unique.

MPLS VPN architecture differs in an important way from traditional peer-to-peer VPN solutions: MPLS VPNS support overlapping customer address spaces.

With the deployment of a single routing protocol (BGP) exchanging all customer routes between PE routers, an important issue arises: how can BGP propagate several identical prefixes, belonging to different customers, between PE routers?

The only solution to this dilemma is the expansion of customer IP prefixes with a unique prefix that makes them unique even if they had previously overlapped. A 64-bit prefix called the RD is used in MPLS VPNs to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE routers.

What Are RDs?

This topic describes the features of an RD.

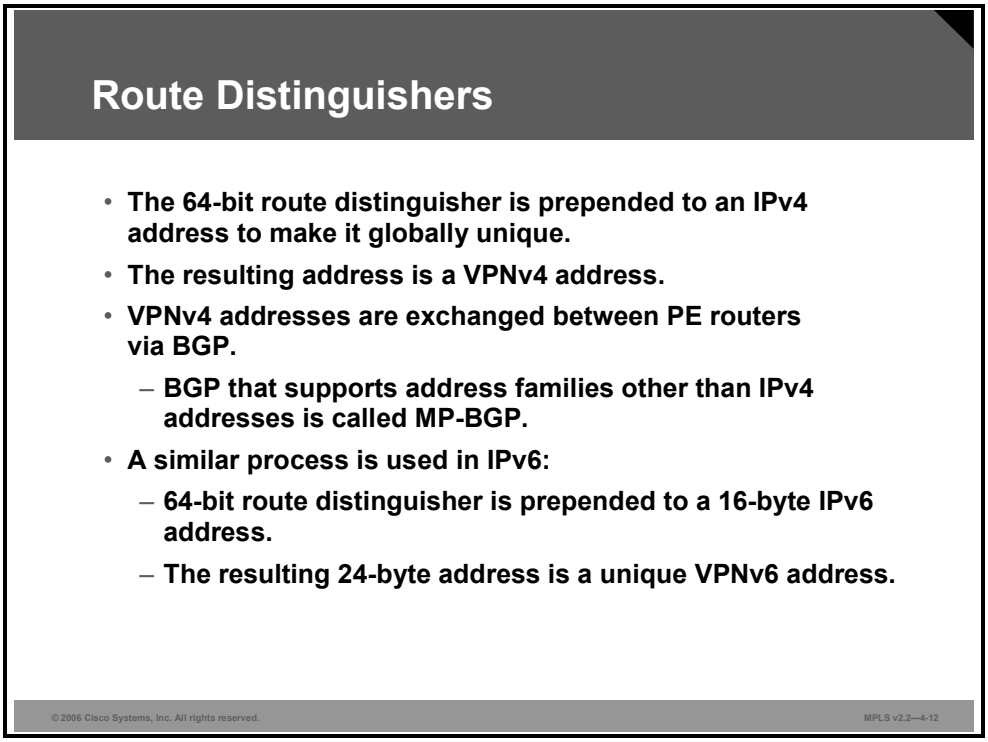

The RD is used only to transform non-unique 32-bit customer IP version 4 (IPv4) addresses into unique 96-bit VPN version 4 (VPNv4) addresses (also called VPN IPv4 addresses).

Note	Although the course will focus on VPNv4, in an IP version 6 (IPv6) implementation, the theory is the same. Multiprotocol Border Gateway Protocol (MP-BGP) is enhanced to carry IPv6 in a VPN known as VPN version 6 (VPNv6), which uses a new VPNv6 address family. The VPNv6 address family consists of an 8-byte RD followed by a 16-byte IPv6 prefix. This combination forms a unique VPNv6 identifier of 24 bytes.

VPNv4 addresses are exchanged only between PE routers; they are never used between CE routers. Between PE routers, BGP must therefore support the exchange of traditional IPv4 prefixes and the exchange of VPNv4 prefixes. A BGP session between PE routers is consequently called an MP-BGP session.

Note	The MPLS VPN implementation in Cisco IOS Release 12.4 and earlier supports only MPLS VPN services within a single autonomous system (AS). In such a scenario, the BGP session between PE routers is always an Internal Border Gateway Protocol (IBGP) session.

Route Distinguishers (Cont.)

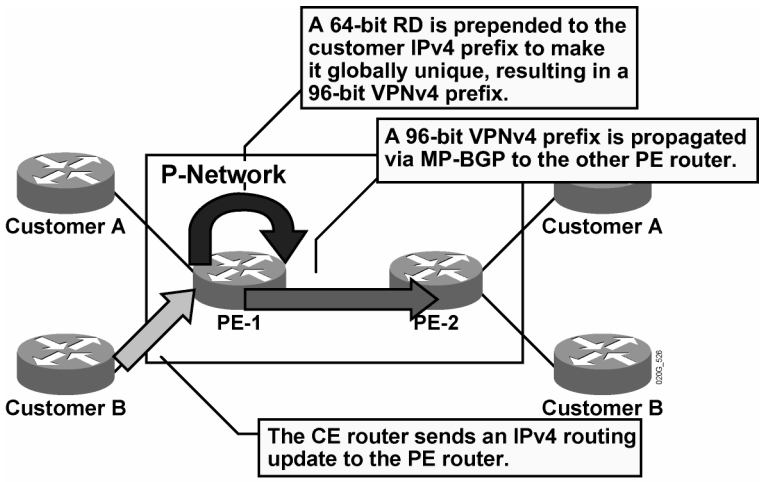

Customer route propagation across an MPLS VPN network is done using this process:

Step 1 The CE router sends an IPv4 routing update to the PE router.

Step 2 The PE router prepends a 64-bit RD to the IPv4 routing update, resulting in a globally unique 96-bit VPNv4 prefix.

Step 3 The VPNv4 prefix is propagated via a Multiprotocol Internal Border Gateway Protocol (MP-IBGP) session to other PE routers.

Route Distinguishers (Cont.)

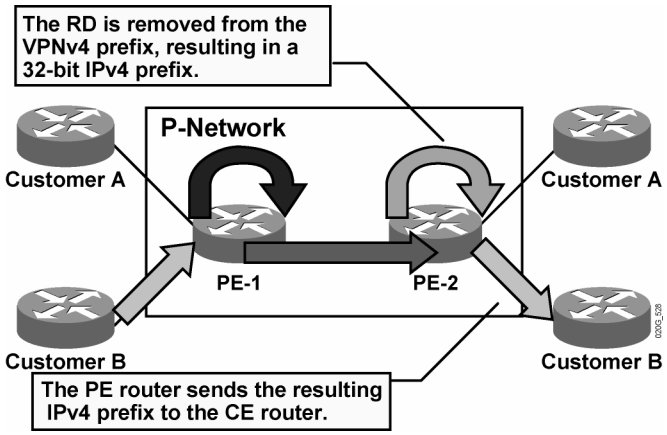

Step 4 The receiving PE routers strip the RD from the VPNv4 prefix, resulting in an IPv4 prefix.

Step 5 The IPv4 prefix is forwarded to other CE routers within an IPv4 routing update.

RDs: Usage in an MPLS VPN

- **The RD has no special meaning.**
- **The RD is used only to make potentially overlapping IPv4 addresses globally unique.**
- **The RD is used as a VPN identifier, but this design could not support all topologies required by the customers.**

The RD has no special meaning or role in MPLS VPN architecture; its only function is to make overlapping IPv4 addresses globally unique. The RD value has a local significance on the router where it is configured.

Note	Because there has to be a unique one-to-one mapping between RD and virtual routing and forwarding instances (VRFs), the RD could be viewed as the virtual routing and forwarding (VRF) identifier in the Cisco implementation of an MPLS VPN.

The RD is configured at the PE router as part of the setup of the VPN site. The RD is not configured on the CE router, and is not visible to the customer.

Simple VPN topologies require only one RD per customer, raising the possibility that the RD could serve as a VPN identifier. This design, however, would not allow implementation of more complex VPN topologies, such as when a customer site belongs to multiple VPNs.

Is the RD Enough?

This topic describes why RDs are not sufficient to identify VPNs.

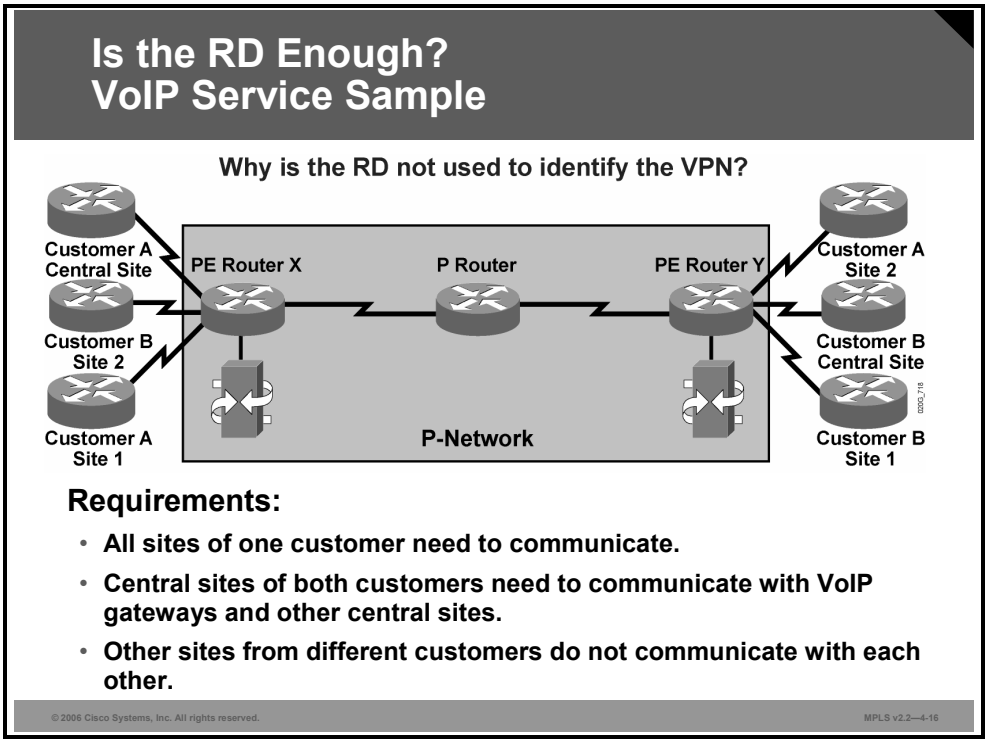

To illustrate the need for a more versatile VPN indicator than the RD, consider the Voice over IP (VoIP) service.

Example: VoIP Service Sample

The figure illustrates the need for a more versatile VPN indicator than the RD. The connectivity requirements of the VoIP service are as follows:

- All sites of a single customer need to communicate.
- The central sites of different customers subscribed to the VoIP service need to communicate with the VoIP gateways (to originate and receive calls in the public voice network) and also with other central sites to exchange intercompany voice calls.

Note	Additional security measures would have to be put in place at central sites to ensure that the central sites exchange only VoIP calls with other central sites. Otherwise, the corporate network of a customer could be compromised by another customer who is using the VoIP service.

Example: Connectivity Requirements

The connectivity requirements of the VoIP service are illustrated in the figure.

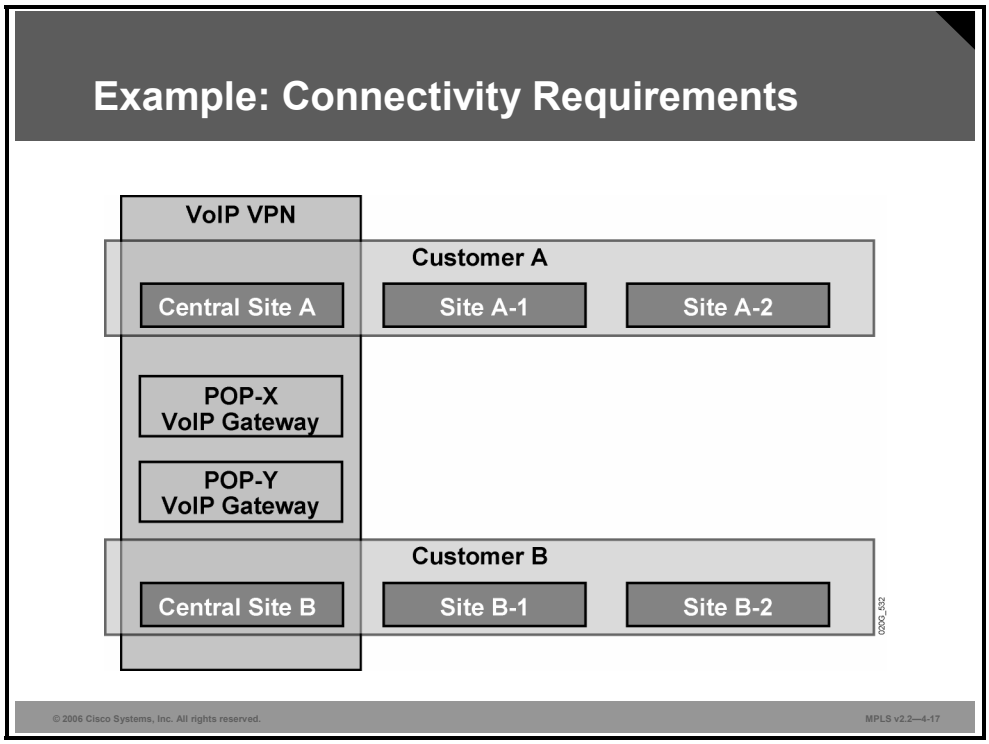

Three VPNs are needed to implement the desired connectivity: two customer VPNs and a shared VoIP VPN. Central customer sites participate in the customer VPN and in the VoIP VPN.

Note	The POP-X and POP-Y VoIP gateways were attached to PE router X and PE router Y in the previous graphic.

The RD (again, a single entity prepended to an IPv4 route) cannot indicate that a site participates in more than one VPN. A method is needed in which a set of VPN identifiers can be attached to a route to indicate its membership in several VPNs.

What Are RTs?

This topic describes the features of an RT.

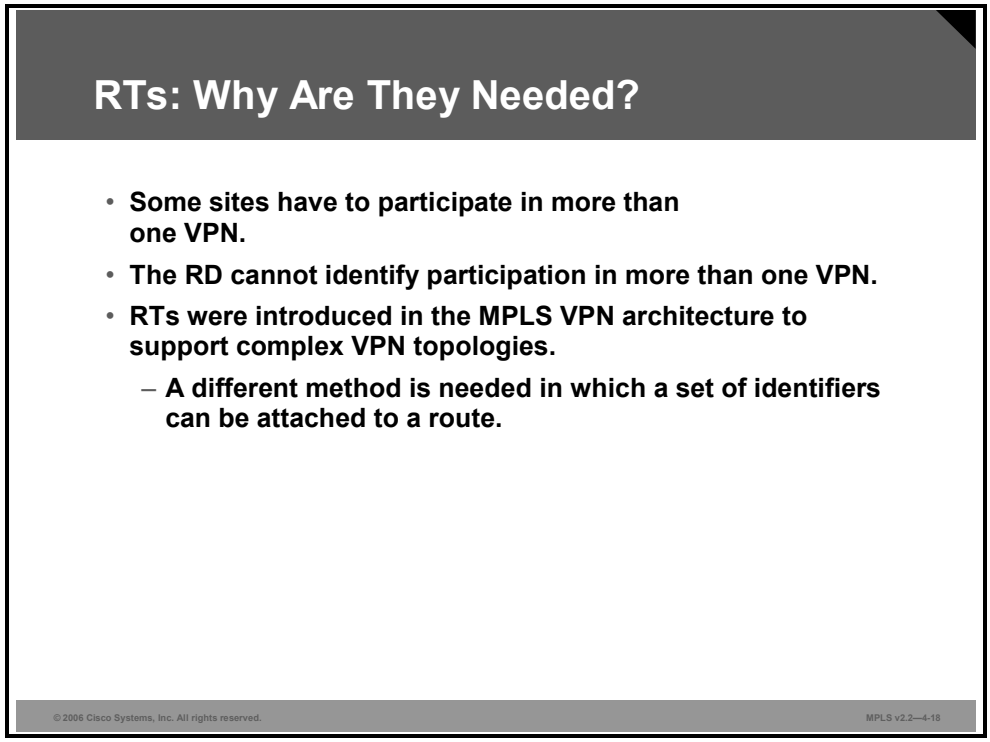

RTs were introduced into the MPLS VPN architecture to support identifying a site that participates in more than one VPN.

RTs: What Are They?

- **RTs are additional attributes attached to VPNv4 BGP routes to indicate VPN membership.**
- **Extended BGP communities are used to encode these attributes.**
  - **Extended communities carry the meaning of the attribute together with its value.**
- **Any number of RTs can be attached to a single route.**

RTs are attributes that are attached to a VPNv4 BGP route to indicate its VPN membership. The extended BGP communities of routing updates are used to carry the RT of that update, thus identifying to which VPN the update belongs.

As with standard BGP communities, a set of extended communities can be attached to a single BGP route, satisfying the requirements of complex VPN topologies.

Extended BGP communities are 64-bit values. The semantics of the extended BGP community are encoded in the high-order 16 bits of the value, making those bits useful for a number of different applications, such as MPLS VPN RTs.

RTs: How Do They Work?

- **Export RTs:**
  - **Identifying VPN membership**
  - **Appended to the customer route when it is converted into a VPNv4 route**
- **Import RTs:**
  - **Associated with each virtual routing table**
  - **Select routes to be inserted into the virtual routing table**

MPLS VPN RTs are attached to a customer route at the moment that it is converted from an IPv4 route to a VPNv4 route by the PE router. The RTs attached to the route are called export RTs and are configured separately for each virtual routing table in a PE router. Export RTs identify a set of VPNs in which sites associated with the virtual routing table belong.

When the VPNv4 routes are propagated to other PE routers, those routers need to select the routes to import into their virtual routing tables. This selection is based on import RTs. Each virtual routing table in a PE router can have a number of configured import RTs that identify the set of VPNs from which the virtual routing table is accepting routes.

In overlapping VPN topologies, RTs are used to identify VPN membership. Advanced VPN topologies (for example, central services VPNs) use RTs in more complex scenarios.

How Have Complex VPNs Redefined the Meaning of VPNs?

This topic describes how complex VPNs have redefined the meaning of VPNs.

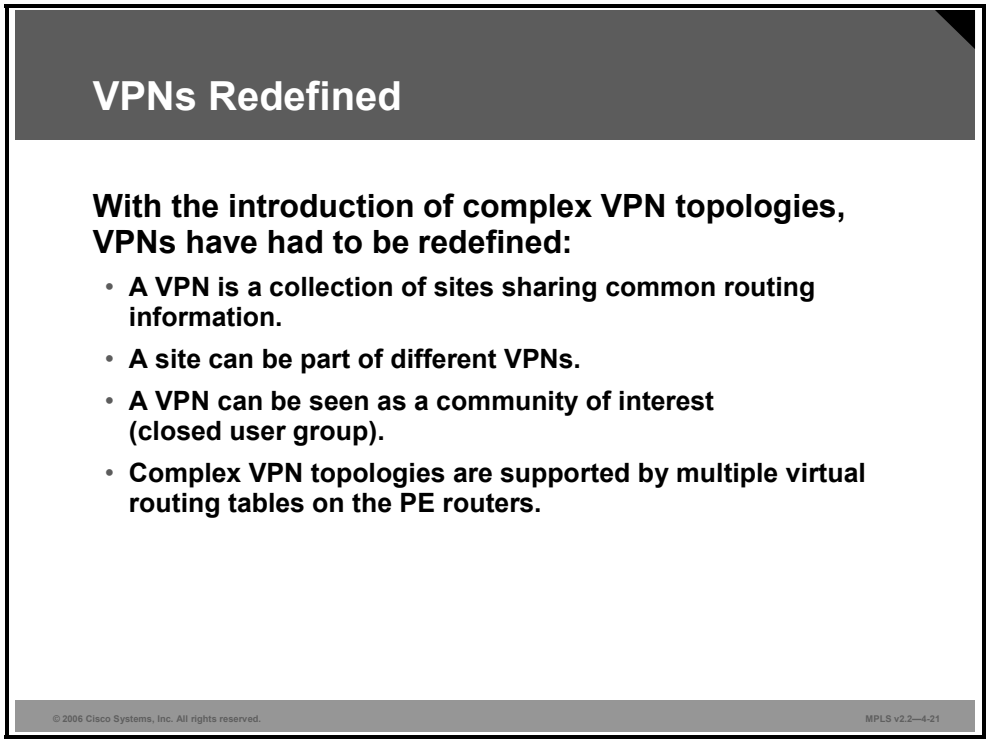

With the introduction of complex VPN topologies, the definition of a VPN has needed to be changed. A VPN is simply a collection of sites sharing common routing information. In traditional switched WAN terms (for example, in X.25 terminology), such a concept would be called a closed user group (CUG).

In the classic VPN, all sites connected to a VPN shared a common routing view. In complex VPNs, however, a site can be part of more than one VPN. This results in differing routing requirements for sites that belong to a single VPN and those that belong to more than one VPN. These routing requirements have to be supported with multiple virtual routing tables on the PE routers.

What Is the Impact of Complex VPN Topologies on Virtual Routing Tables?

This topic describes the impact of complex VPN topologies on virtual routing tables.

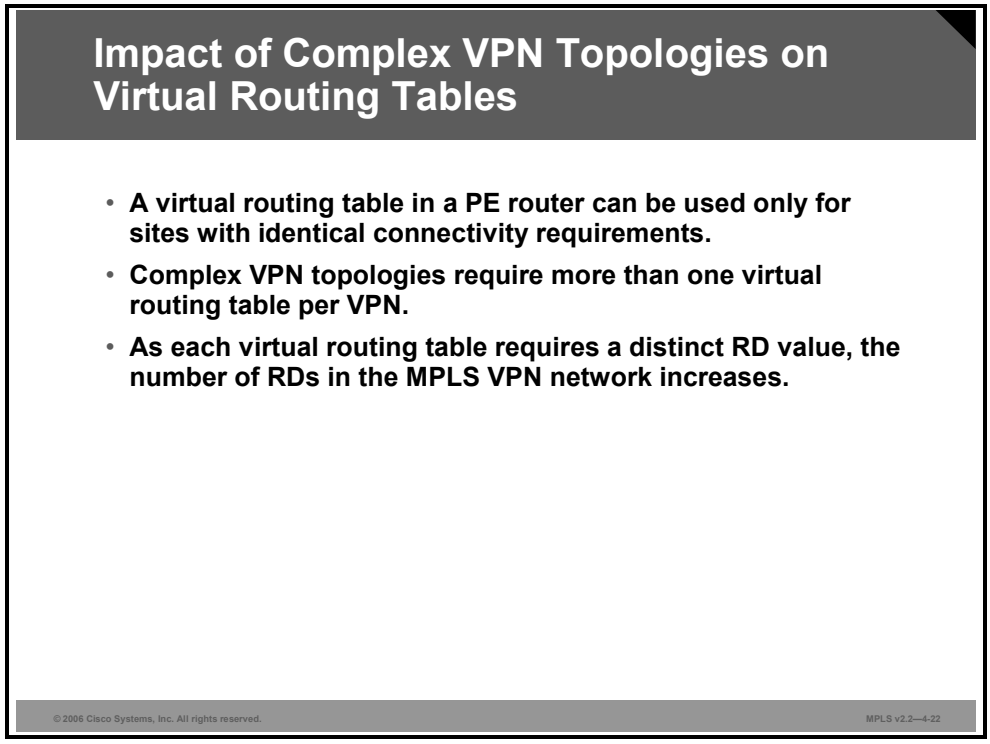

A single virtual routing table can be used only for sites with identical connectivity requirements. Complex VPN topologies, therefore, require more than one virtual routing table per VPN.

Note	If sites with different requirements are associated with the same virtual routing table, some of the sites might be able to access destinations that should not be accessible to them.

Because each virtual routing table requires a distinctive RD, the number of RDs in an MPLS VPN network increases with the introduction of overlapping VPNs. Moreover, the simple association between RD and VPN that was true for simple VPNs is also gone.

Example: Impact of Complex VPN Topologies on Virtual Routing Tables

To illustrate the requirements for multiple virtual routing tables, consider a VoIP service with three VPNs (customer A, customer B, and a VoIP VPN).

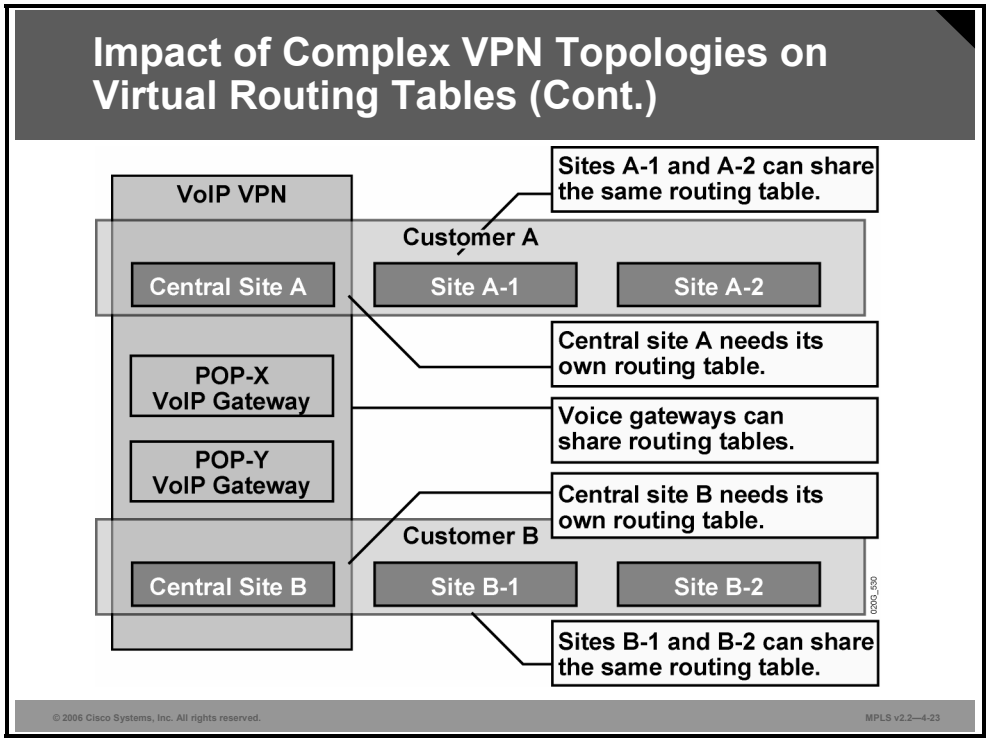

The virtual routing table needs of this service are as follows:

1. All sites of customer A (apart from the central site) can share the same virtual routing table because they belong to a single VPN.
2. The same is true for all sites of customer B (apart from the central site).
3. The VoIP gateways participate only in the VoIP VPN and can belong to a single virtual routing table.
4. Central site A has unique connectivity requirements—it has to see sites of customer A and sites in the VoIP VPN and, consequently, requires a dedicated virtual routing table.
5. Likewise, central site B requires a dedicated virtual routing table.

Therefore, in this example, five different VRF tables are needed to support three VPNs. There is no one-to-one relationship between the number of VRFs and the number of VPNs.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- There are several drawback to traditional peer-to-peer VPNs.
- MPLS VPN architecture combines the best features of the overlay and peer-to-peer VPN models.
- The architecture of a PE router in an MPLS VPN uses separate virtual routers containing the routes of each customers inside one physical router.
- The most scalable method of exchanging customer routes across a provider network is the use of a single BGP routing protocol from PE router to PE router.

Summary (Cont.)

- Route distinguishers transform non-unique 32-bit addresses into 96-bit unique addresses.
- Route targets are used to identify VPN membership in overlapping topologies.
- VPNs are now considered a collection of sites sharing common routing information.
- Placing sites with different routing requirements in the same virtual routing table will result in inconsistent routing.

Lesson 4

Introducing the MPLS VPN Routing Model

Overview

This lesson explains the routing requirements for Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). The lesson offers address and routing perspectives from the customer and service provider side, and it discusses how routing tables appear on provider edge (PE) routers. This lesson also discusses MPLS VPN end-to-end information flow, Multiprotocol Border Gateway Protocol (MP-BGP), updates, and display formats.

It is important to understand how information is routed in an MPLS VPN, and how the routing tables are viewed and interpreted. This lesson will help you to get a clear understanding of the similarities and differences between the global routing table and the virtual routing tables that are created in an MPLS VPN.

Objectives

Upon completing this lesson, you will be able to identify the routing requirements for MPLS VPNs. This ability includes being able to meet these objectives:

- Describe the routing requirements for MPLS VPNs
- Describe the MPLS VPN routing model for CE routers, PE routers, and P routers
- Describe how IPv4 is used to provide support for existing Internet routing
- Identify the routing tables implemented in the PE router to support MPLS VPNs
- Describe the end-to-end flow of routing updates in an MPLS VPN
- Describe how an MPLS VPN determines which routes are distributed to a CE router

MPLS VPN Routing Requirements and Model

This topic describes the routing requirements and model for MPLS VPNs.

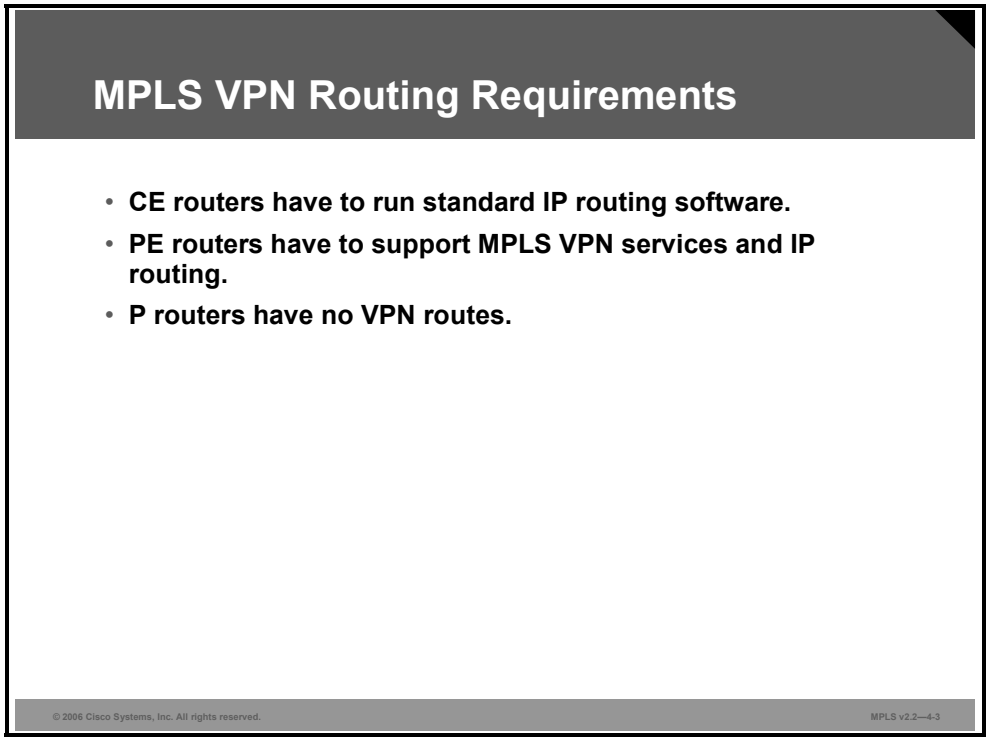

The designers of MPLS VPN technology were faced with these routing requirements:

- Customer edge (CE) routers should not be MPLS VPN-aware; CE routers should run standard IP routing software.
- PE routers must support MPLS VPN services and traditional Internet services.
- To make the MPLS VPN solution scalable, provider routers (P routers) must not carry VPN routes.

What Is the MPLS VPN Routing Model?

This section describes the MPLS VPN routing model for CE routers, PE routers, and P routers.

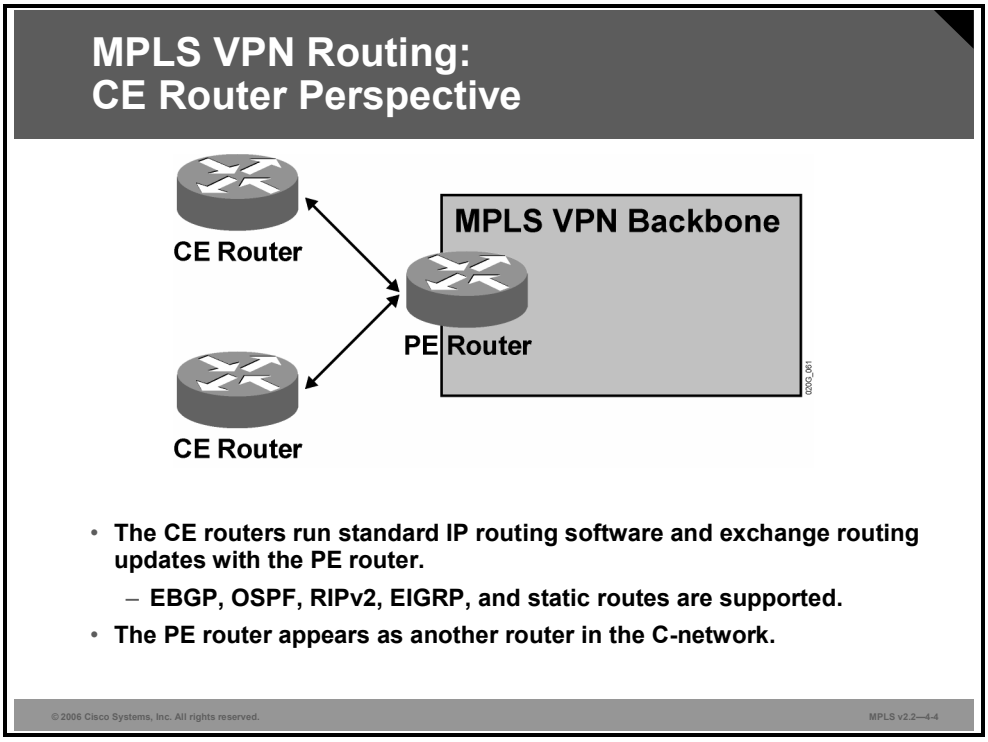

The MPLS VPN backbone should look like a standard corporate backbone to the CE routers. The CE routers run standard IP routing software and exchange routing updates with the PE routers, which appear to them as normal routers in the customer network (C-network).

Note	Since Cisco IOS Release 12.2, the choice of routing protocols that can be run between a CE router and a PE router includes Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol version 2 (RIPv2), Open Shortest Path First (OSPF), External Border Gateway Protocol (EBGP), and static routes.

MPLS VPN Routing: Overall Customer Perspective

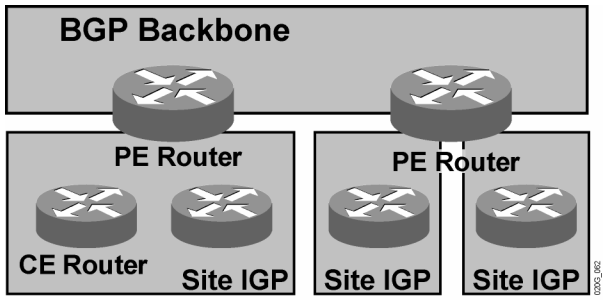

- **To the customer, the PE routers appear as core routers connected via a BGP backbone.**
- **The usual BGP and IGP design rules apply.**
- **The P routers are hidden from the customer.**

From the customer perspective, the MPLS VPN backbone looks like an intracompany Border Gateway Protocol (BGP) backbone with PE routers performing route redistribution between individual sites and the core backbone. The standard design rules used for enterprise BGP backbones can be applied to the design of the C-network.

The P routers are hidden from customer view; the internal topology of the BGP backbone is therefore transparent to the customer.

MPLS VPN Routing: P Router Perspective

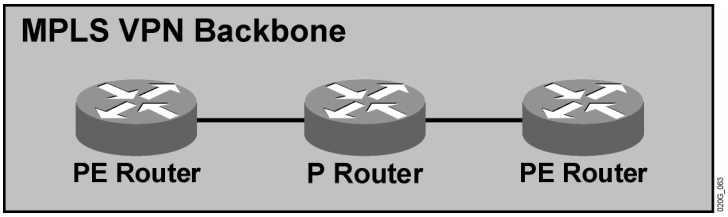

- **P routers do not participate in MPLS VPN routing and do not carry VPN routes.**
- **P routers run backbone IGP with the PE routers and exchange information about global subnetworks (core links and loopbacks).**

From the P router perspective, the MPLS VPN backbone looks even simpler—the P routers do not participate in MPLS VPN routing and do not carry VPN routes. The P routers run only a backbone Interior Gateway Protocol (IGP) with other P routers and with PE routers, and exchange information about core subnetworks. BGP deployment on P routers is not needed for proper MPLS VPN operation; it might be needed, however, to support traditional Internet connectivity that has not yet been migrated to MPLS.

MPLS VPN Routing: PE Router Perspective

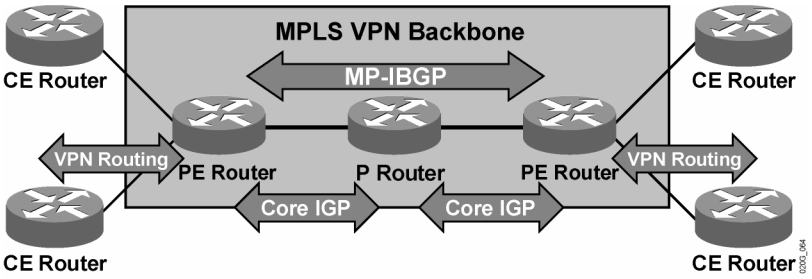

PE routers:

- **Exchange VPN routes with CE routers via per-VPN routing protocols**
- **Exchange core routes with P routers and PE routers via core IGP**
- **Exchange VPNv4 routes with other PE routers via MP-IBGP sessions**

The PE routers are the only routers in MPLS VPN architecture that see all routing aspects of the MPLS VPN. PE routers are able to perform these exchanges:

- PE routers exchange IP version 4 (IPv4) VPN routes with CE routers via various routing protocols running in the virtual routing tables.
- PE routers exchange VPN version 4 (VPNv4) routes via Multiprotocol Internal Border Gateway Protocol (MP-IBGP) sessions with other PE routers.
- PE routers exchange core routes with P routers and other PE routers via core IGP.

Existing Internet Routing Support

This topic describes how IPv4 is used to provide support for existing Internet routing.

Support for Existing Internet Routing

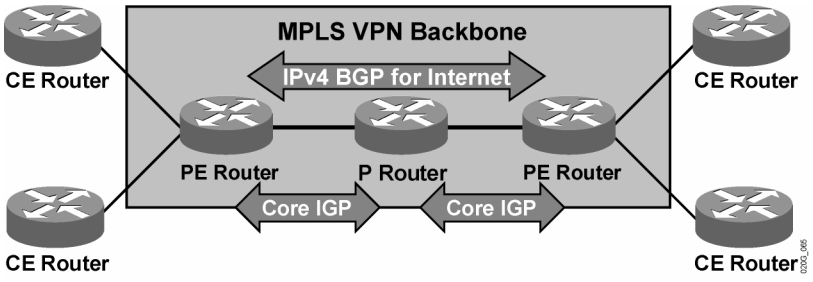

PE routers can run standard IPv4 BGP in the global routing table:

- **PE routers exchange Internet routes with other PE routers.**
- **CE routers do not participate in Internet routing.**
- **P routers do not need to participate in Internet routing.**

The routing requirements for PE routers also extend to supporting Internet connectivity—PE routers have to exchange Internet routes with other PE routers. The CE routers cannot participate in Internet routing if the Internet routing is performed in global address space. The P routers could participate in Internet routing; however, Internet routing should be disabled on the P routers to make the network core more stable.

Routing Tables on PE Routers

This topic identifies the routing tables implemented in the PE router to support MPLS VPNs.

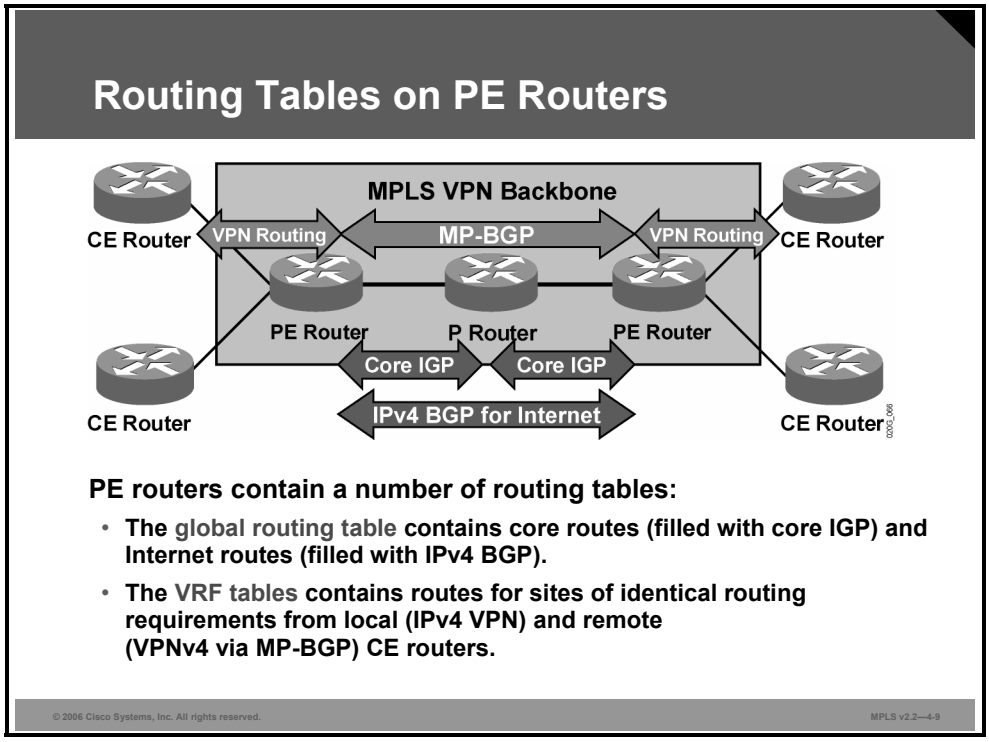

The PE routers fulfill various routing requirements imposed on them by using a number of IP routing tables. Here are some examples:

- The global IP routing table (the IP routing table that is always present in a Cisco IOS software-based router even if it is not supporting an MPLS VPN) contains all core routes (inserted by the core IGP) and the Internet routes (inserted from the global IPv4 BGP table).
- The virtual routing and forwarding (VRF) tables contain sets of routes for sites with identical routing requirements. The VRFs are filled with intra-VPN IGP information exchanged with the CE routers and with VPNv4 routes received through MP-BGP sessions from the other PE routers.

Identifying End-to-End Routing Update Flow

This topic describes the end-to-end flow of routing updates in an MPLS VPN.

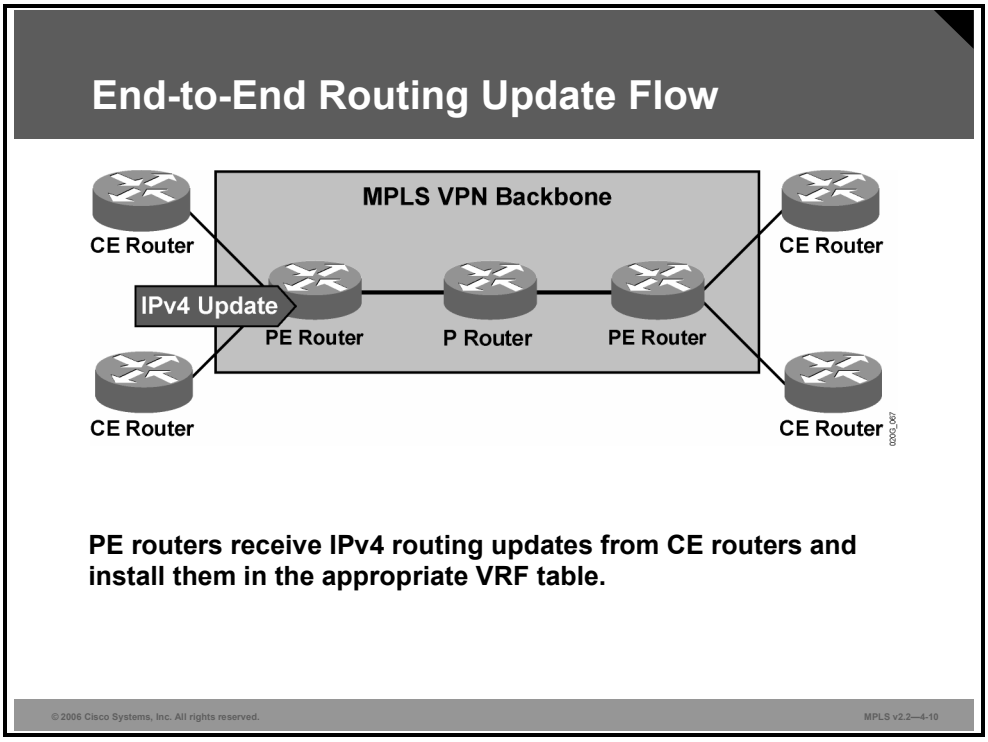

These figures provide an overview of end-to-end routing information flow in an MPLS VPN network.

Example: End-to-End Routing Update Flow

The figure here illustrates how PE routers receive IPv4 routing updates from the CE routers and install them in the appropriate VRF table.

End-to-End Routing Update Flow (Cont.)

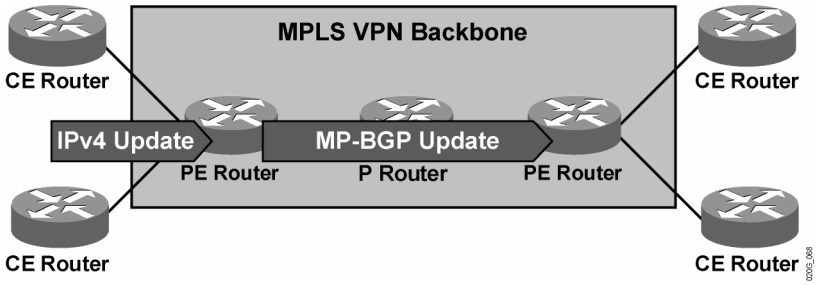

PE routers export VPN routes from VRF tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers.

Current MPLS VPN implementation in Cisco IOS software (up to Cisco IOS Release12.4) supports MPLS VPN services only within the scope of a single autonomous system (AS). The MP-BGP sessions between the PE routers are therefore Internal Border Gateway Protocol (IBGP) sessions and were subject to the IBGP split-horizon rules. Either a full mesh of MP-IBGP sessions is required between PE routers or route reflectors need to be used.

End-to-End Routing Update Flow: MP-BGP Update

An MP-BGP update contains these elements:

- **VPNv4 address**
- **Extended communities (route targets, optionally SOO)**
- **Label used for VPN packet forwarding**
- **Any other BGP attribute (for example, AS path, local preference, MED, standard community)**

An MP-BGP update exchange between PE routers contains these elements:

- VPNv4 address
- Extended BGP communities (route targets [RTs] are required; Site of Origin [SOO] is optional)
- Label used for VPN packet forwarding

Note	The “Forwarding MPLS VPN Packets” lesson explains how this label is used in the MPLS label stack.

- Mandatory BGP attributes (for example, AS path)

Optionally, the MP-BGP update can contain any other BGP attribute; for example, local preference, multi-exit discriminator (MED), or standard BGP community.

End-to-End Routing Update Flow (Cont.)

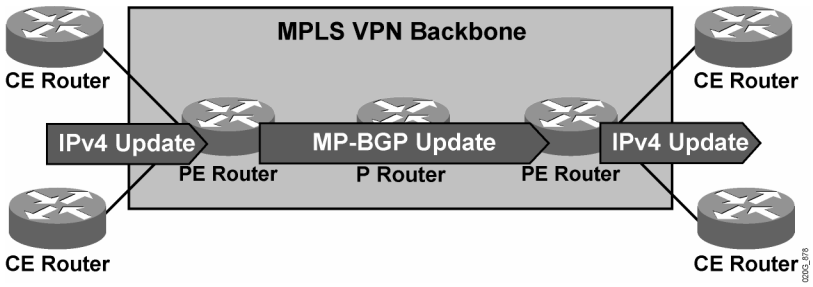

- **The receiving PE router imports the incoming VPNv4 routes into the appropriate VRF based on route targets attached to the routes.**
- **The routes installed in the VRFs are propagated to the CE routers.**

The PE routers receiving MP-BGP updates import the incoming VPNv4 routes into their VRFs based on RTs attached to the incoming routes and on import RTs configured in the VRFs. The VPNv4 routes installed in the VRFs are converted to IPv4 routes and then propagated to the CE routers.

Route Distribution to CE Routers

This topic describes how an MPLS VPN determines which routes are distributed to a CE router.

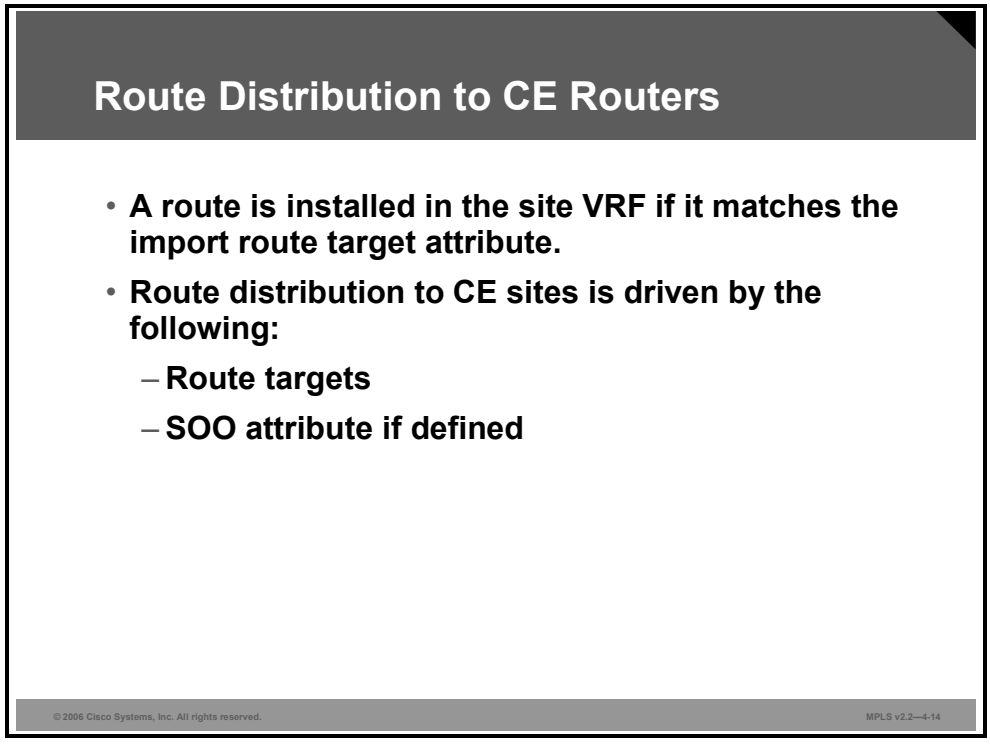

The RTs attached to a route and the import RTs configured in the VRF drive the propagation of the routes to the CE router.

Incoming VPNv4 routes are imported into VRFs on the receiving PE router only if at least one RT attached to the route matches at least one import RT configured in the VRF.

When BGP is used to connect the CE and PE, the SOO attribute attached to the VPNv4 route can also help control the IPv4 route propagation to the CE routers. A route inserted into a VRF is not propagated to a CE router if the SOO attached to the route is equal to the SOO attribute associated with the CE router. The SOO can thus be used to prevent routing loops in MPLS VPN networks with multihomed sites.

To be distributed to the CE, routes need to be installed in the VRF, and not have a conflicting SOO.

Example: Extending MPLS VPNs with VRF-Lite

This topic discusses extending MPLS VPNs with Multi-VRF CE (VRF-lite).

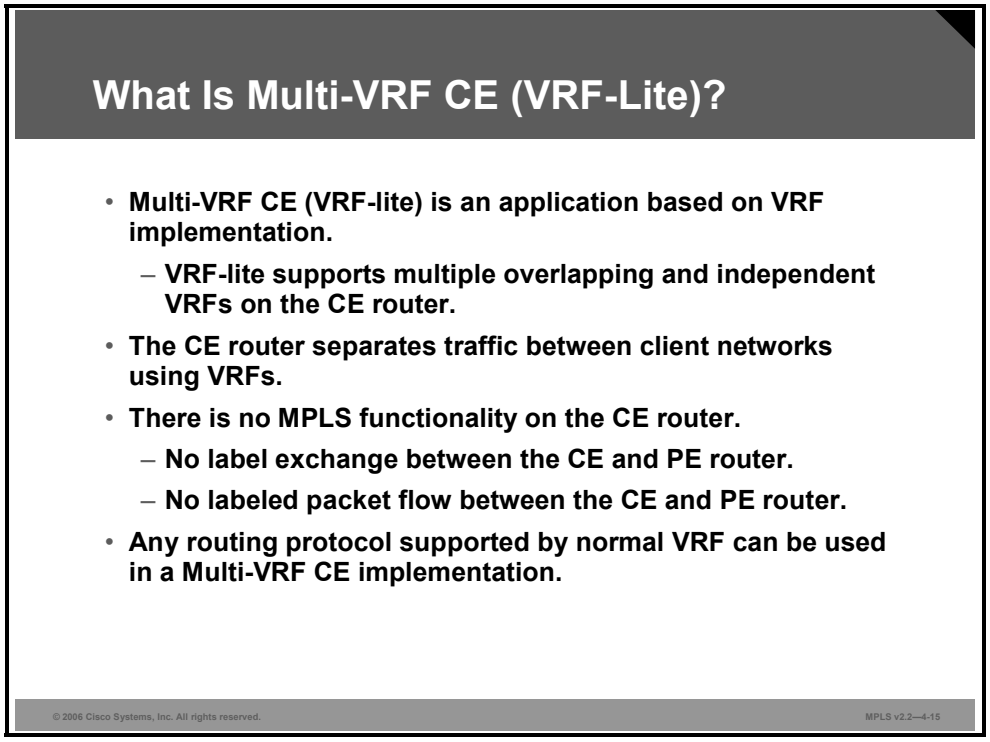

Although MPLS VPNs provide security and privacy as traffic travels through the provider network, the CE router has no mechanism to guarantee private networks across its LAN networks. To provide privacy, each client or organization is traditionally placed in a separate VLAN or on a separate CE router. VRF-lite extends limited PE functionality to a CE router. VRF-lite allows the CE router the ability to maintain separate VRF tables to extend the privacy and security of an MPLS VPN down to a branch office or interface.

The CE router using VRF-lite can isolate traffic by placing each client or organization in a separate VRF with its own IP address space. Each interface or subinterface contains its own IP address space to separate each different client.

Similar to MPLS VRFs, routes are installed in the appropriate VRF with VRF-lite. However, the CE router does not run MPLS.

Note	With VRF-lite, there is no label exchange, there is no Label Distribution Protocol (LDP) adjacency, and there is *no labeled* packet flow between PE and CE router.

The CE router needs a routing protocol or static routes to propagate routes from each specific VRF on the CE router to the same VRF on the PE router.

Note	Additional information on VRF-lite is outside the scope of this course.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- In MPLS VPNs:
  - CE routers run standard protocols (static, RIPv2, OSPF, EIGRP, EBGP) to the PE routers.
  - PE routers provide the VPN routing and services via MP-BGP.
  - P routers do not participate in VPN routing, and only provide core IGP backbone routing to the PE routers.
- The PE router functions are extended to carry regular Internet routing via IPv4 BGP in addition to the MP-BGP.

Summary (Cont.)

- PE routers separate the global IPv4 BGP routing table from each unique customer VPNv4 MP-BGP routing table.
- The ingress PE router receives CE customer IPv4 updates and exports these IPv4 routes to other PE routers via MP-BGP.
- The egress PE router imports the VPNv4 routes and forwards them to the CE router as an IPv4 update.

Lesson 5

Forwarding MPLS VPN Packets

Overview

This lesson explains how forwarding across a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) backbone occurs, identifies how labels get propagated, and explains the effects of summarization in the core.

It is important to understand how packets are forwarded across an MPLS VPN backbone, because this understanding will help you when you try to isolate problems in the network. This lesson explains how the far-end label is sent to the ingress provider edge (PE) router and how that information is shared.

Objectives

Upon completing this lesson, you will be able to describe how packets are forwarded in an MPLS VPN environment. This ability includes being able to meet these objectives:

- Describe the end-to-end MPLS VPN forwarding mechanisms
- Describe the operation of PHP in an MPLS VPN environment
- Describe how labels are propagated between PE routers
- Describe the effects of MPLS VPNs on label propagation
- Describe the effects of MPLS VPNs on packet forwarding

What Are the End-to-End VPN Forwarding Mechanisms?

This topic describes the end-to-end MPLS VPN forwarding mechanisms.

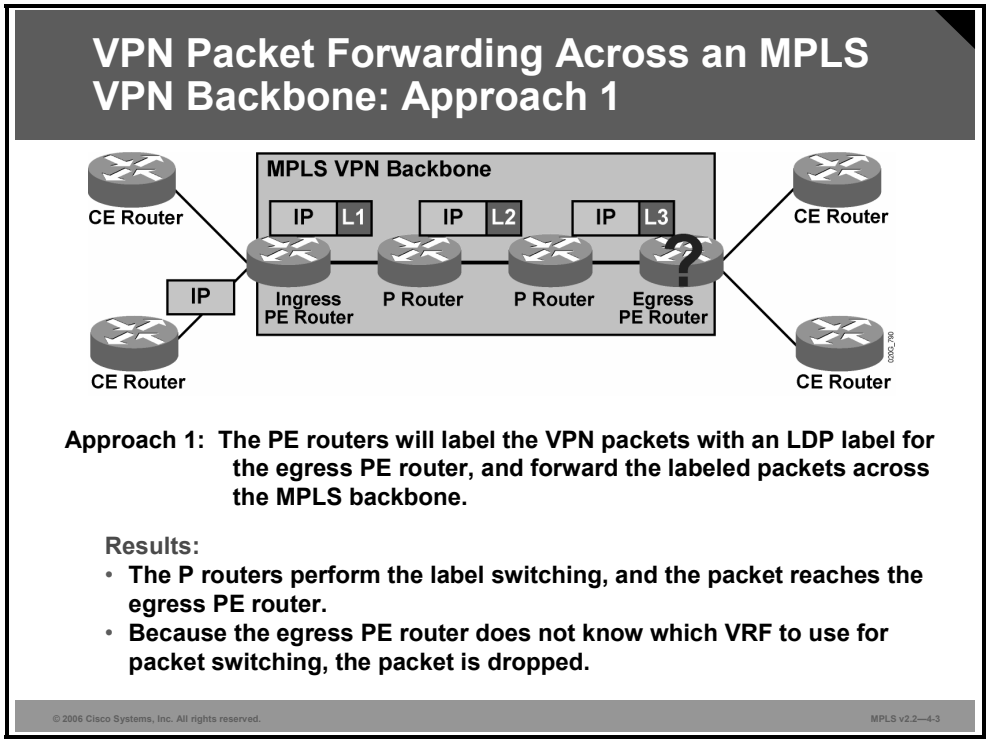

A simple MPLS-oriented approach to MPLS VPN packet forwarding across the MPLS VPN backbone would be to label the customer packet with the label assigned by Label Distribution Protocol (LDP) for the egress PE router. The core routers consequently would never see the customer IP packet; instead, the core routers would see just a labeled packet targeted toward the egress PE router. The core routers would perform simple label-switching operations, eventually delivering the customer packet to the egress PE router. Unfortunately, the customer IP packet would contain no VPN or virtual routing and forwarding (VRF) information that could be used to perform VRF lookup on the egress PE router. The egress PE router would not know which VRF to use for packet lookup and would have to drop the packet.

VPN Packet Forwarding Across an MPLS VPN Backbone: Approach 2

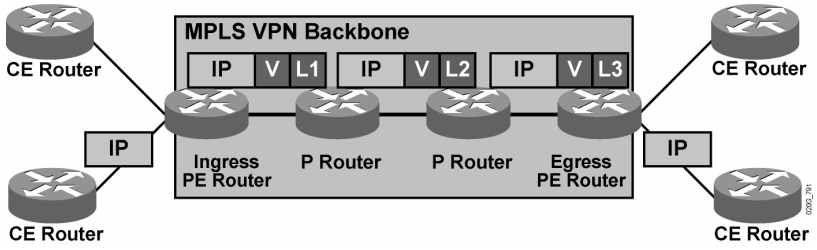

Approach 2: **The PE routers will label the VPN packets with a label stack, using the LDP label for the egress PE router as the top label, and the VPN label assigned by the egress PE router as the second label in the stack.**

Result:

- **The P routers perform label switching using the top label, and the packet reaches the egress PE router. The top label is removed.**
- **The egress PE router performs a lookup on the VPN label and forwards the packet toward the CE router.**

An MPLS label stack can be used to tell the egress PE router what to do with the VPN packet. When using the label stack, the ingress PE router labels the incoming IP packet with two labels. The top label in the stack is the LDP label for the egress PE router; this label guarantees that the packet will traverse the MPLS VPN backbone and arrive at the egress PE router. The second label in the stack is assigned by the egress PE router, and tells how to forward the incoming VPN packet. The second label could point directly toward an outgoing interface, in which case the egress PE router would perform label lookup only on the VPN packet. The second label could also point to a VRF, in which case the egress PE router would first perform a label lookup to find the target VRF, and then perform an IP lookup within the VRF.

Both methods of implementing second labels are used in Cisco IOS software. The second label in the stack points toward an outgoing interface whenever the customer edge (CE) router is the next hop of the VPN route. The second label in the stack points to the VRF table for aggregate VPN routes, VPN routes pointing to a null interface, and routes for directly connected VPN interfaces.

The two-level MPLS label stack satisfies these MPLS VPN forwarding requirements:

- The P routers perform label switching on the LDP-assigned label toward the egress PE router.
- The egress PE router performs label switching on the second label (which it has previously assigned) and either forwards the IP packet toward the CE router or performs another IP lookup in the VRF pointed to by the second label in the stack.

What Is VPN PHP?

This topic describes operation of penultimate hop popping (PHP) in an MPLS VPN environment.

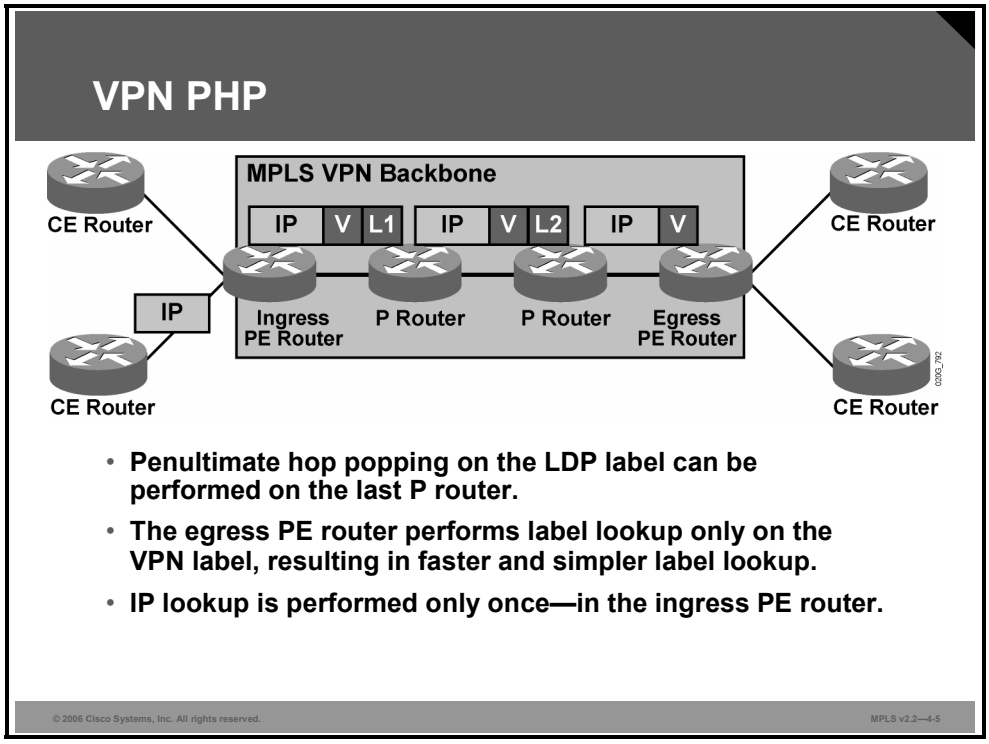

PHP is the removal of the top label in the stack on the hop prior to the egress router. PHP can be performed in frame-based MPLS networks. In these networks, the last provider router (P router) in the label-switched path (LSP) tunnel pops the LDP label (as previously requested by the egress PE router through LDP), and the PE router receives a labeled packet that contains only the VPN label. In most cases, a single label lookup performed on that packet in the egress PE router is enough to forward the packet toward the CE router. The full IP lookup through the Forwarding Information Base (FIB) is performed only once, in the ingress PE router, even without PHP.

Propagating VPN Labels Between PE Routers

This topic describes how labels are propagated between PE routers.

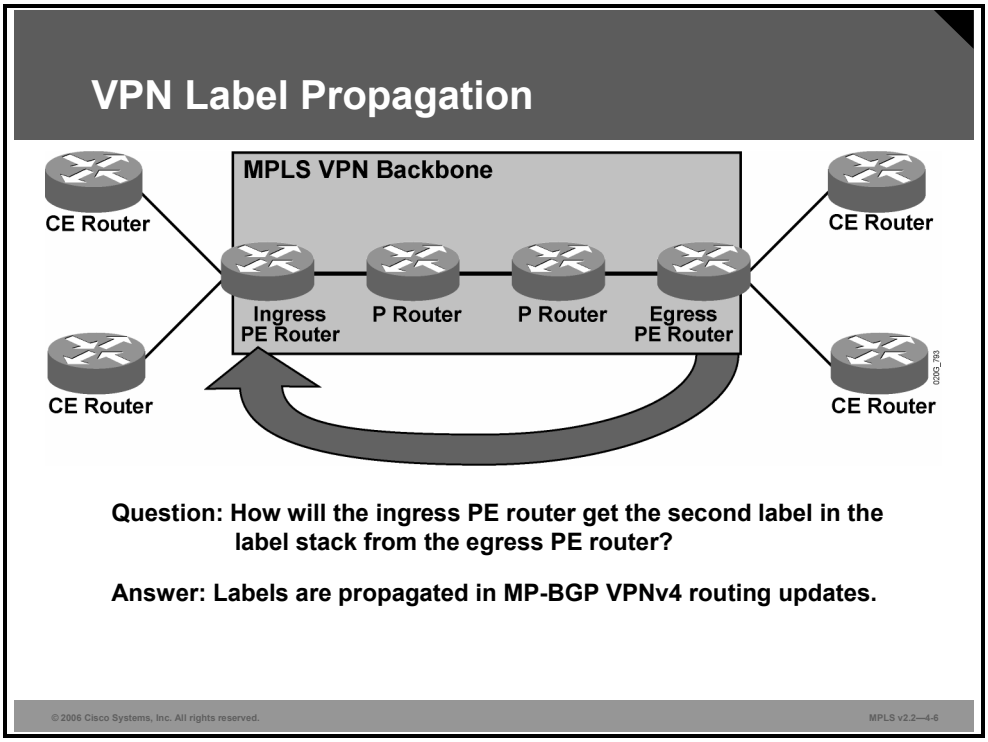

The previous figures showed that an MPLS label stack with the second label is required for proper MPLS VPN operation. This label was allocated by the egress PE router. This label has to be propagated from the egress PE router to the ingress PE routers to enable proper packet forwarding. Multiprotocol Border Gateway Protocol (MP-BGP) was chosen as the propagation mechanism. Every MP-BGP update thus carries a label assigned by the egress PE router together with the 96-bit VPN version 4 (VPNv4) prefix.

Example: VPN Label Propagation Between PE Routers

The figure illustrates VPN label propagation between PE routers.

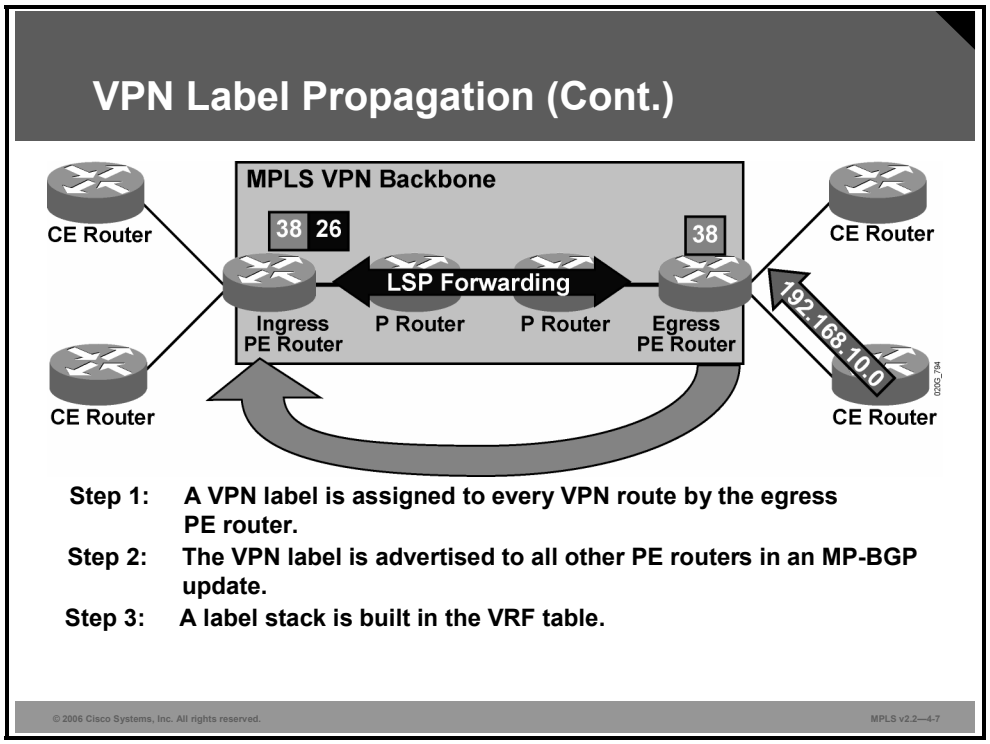

These steps describe the label propagation between PE routers:

Step 1 The egress PE router assigns a label to every VPN route received from the attached CE routers and to every summary route summarized inside the PE router. This label is then used as the second label in the MPLS label stack by the ingress PE routers when labeling VPN packets.

Note	In the graphic, the VPN label 38 for destination 192.168.10.0 is assigned by the egress PE router.

The VPN labels assigned locally by the PE router can be inspected with the **show mpls forwarding vrf** *vrf-name* command.

Step 2 The VPN labels assigned by the egress PE routers are advertised to all other PE routers together with the VPNv4 prefix in MP-BGP updates.

The labels can be inspected with the **show ip bgp vpnv4 all labels** command on the ingress PE router.

The routes that have an input label but no output label are the routes received from the CE routers (and the input label was assigned by the local PE router). The routes with an output label but no input label are the routes received from the other PE routers (and the output label was assigned by the remote PE router).

Step 3 The ingress PE router has two labels associated with a remote VPN route: a label for the Border Gateway Protocol (BGP) next hop assigned by the next-hop P router via LDP—and taken from the local Label Information Base (LIB)—and also the label assigned by the remote PE router and propagated via MP-BGP update. Both labels are combined in a label stack and installed in the VRF table.

The label stack in the VRF table can be inspected using the **show ip cef vrf** *vrf-name* **detail** command. The *tags imposed* values in the output displays the MPLS label stack. The first label in the MPLS label stack is the LDP label forwarded toward the egress PE router, and the second label is the VPN label advertised by the egress PE router.

What Are the Effects of MPLS VPNs on Label Propagation?

This topic describes the effects of MPLS VPNs on label propagation.

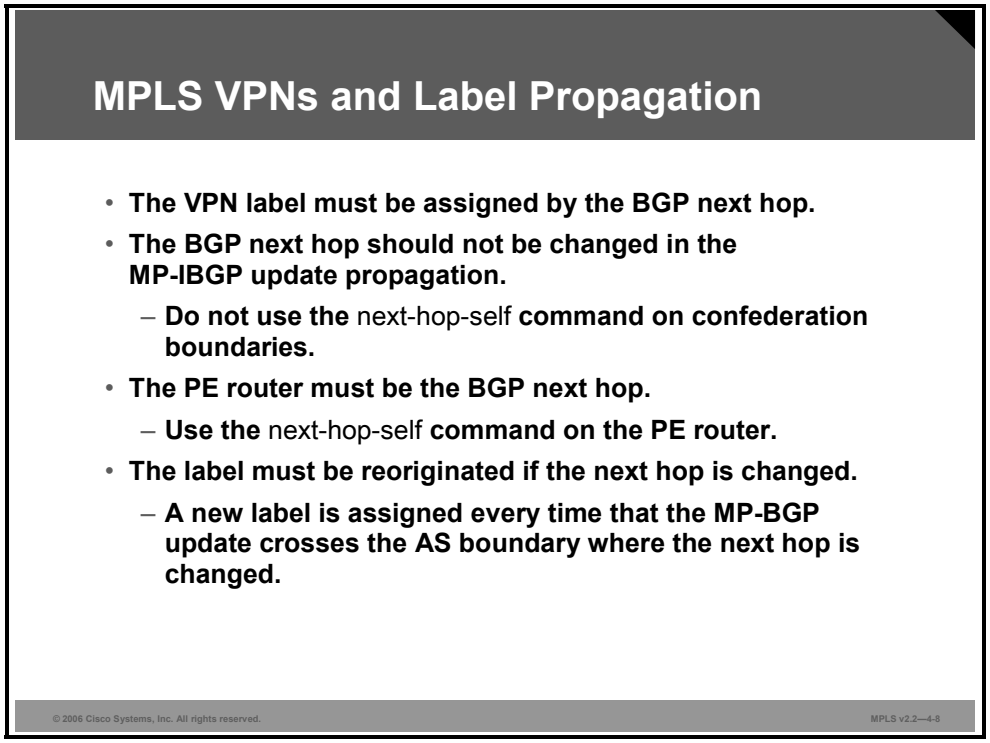

MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same router as the one that assigned the second label in the label stack. Here are three scenarios that can cause the BGP next hop to be different from the IP address of the PE router assigning the VPN label:

- If the customer route is received from the CE router via an External Border Gateway Protocol (EBGP) session, the next hop of the VPNv4 route is still the IP address of the CE router (the BGP next hop of an outgoing Internal Border Gateway Protocol [IBGP] update is always identical to the BGP next hop of the incoming EBGP update). You have to configure the **next-hop-self** command on the MP-BGP sessions between PE routers to make sure that the BGP next hop of the VPNv4 route is always the IP address of the PE router, regardless of the routing protocol used between the PE router and the CE router.
- The BGP next hop should not change inside an autonomous system (AS). It can change, however, if you use the **next-hop-self** command on an inter-AS boundary inside a BGP confederation or if you use inbound the **route-map** command on a PE route to change the next hop (a strongly discouraged practice). To prevent this situation, never change the BGP next hop with the **route-map** or **next-hop-self** commands inside an AS.
- The BGP next hop is always changed on an EBGP session. If the MPLS VPN network spans multiple public autonomous systems (not just autonomous systems within a BGP confederation), special provisions must be made in the AS boundary routers to reoriginate the VPN label at the same time that the BGP next hop is changed. This functionality is supported by Cisco IOS Releases 12.1(4)T, 12.2, and later.

What Are the Effects of MPLS VPNs on Packet Forwarding?

This topic describes the effects of MPLS VPNs on packet forwarding.

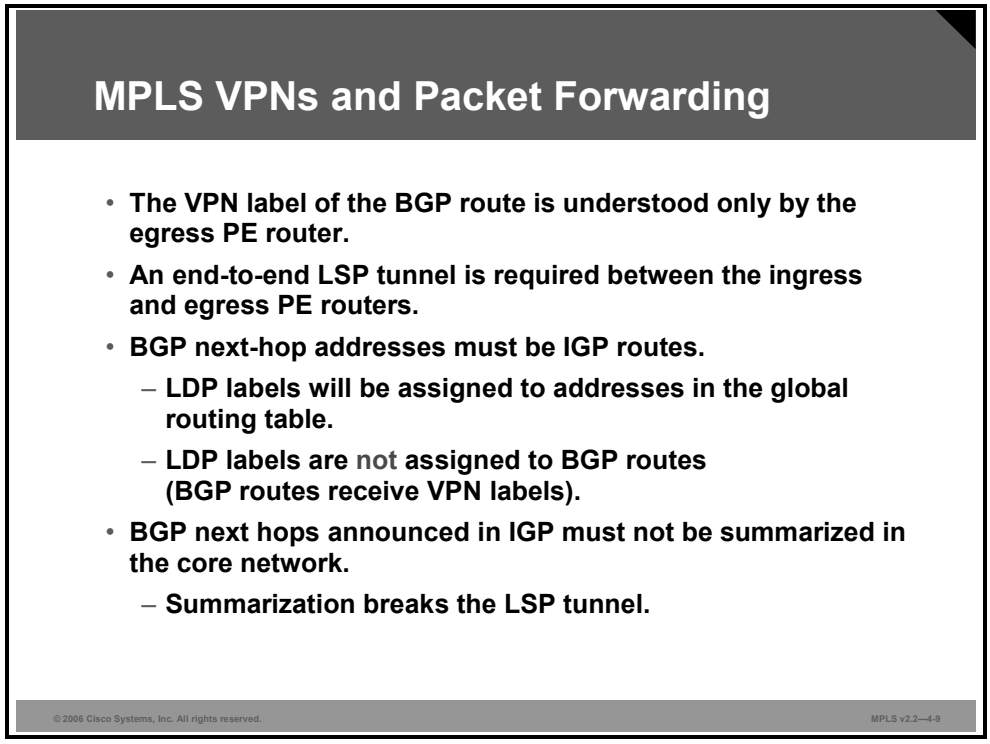

For successful propagation of MPLS VPN packets across an MPLS backbone, there must be an unbroken LSP tunnel between PE routers. This is because the second label in the stack is recognized only by the egress PE router that has originated it, and will not be understood by any other router should it ever become exposed.

Here are two scenarios that could cause the LSP tunnel between PE routers to break:

- If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label and the label stack will not be built correctly. The IP address of the PE router must be announced in the global routing table.
- If the P routers perform summarization of the address range within which the IP address of the egress PE router lies, the LSP tunnel will be disrupted at the summarization point, as illustrated in the figure.

Example: Summarization in the Core

In the figure, the P router summarizes the loopback address of the egress PE router.

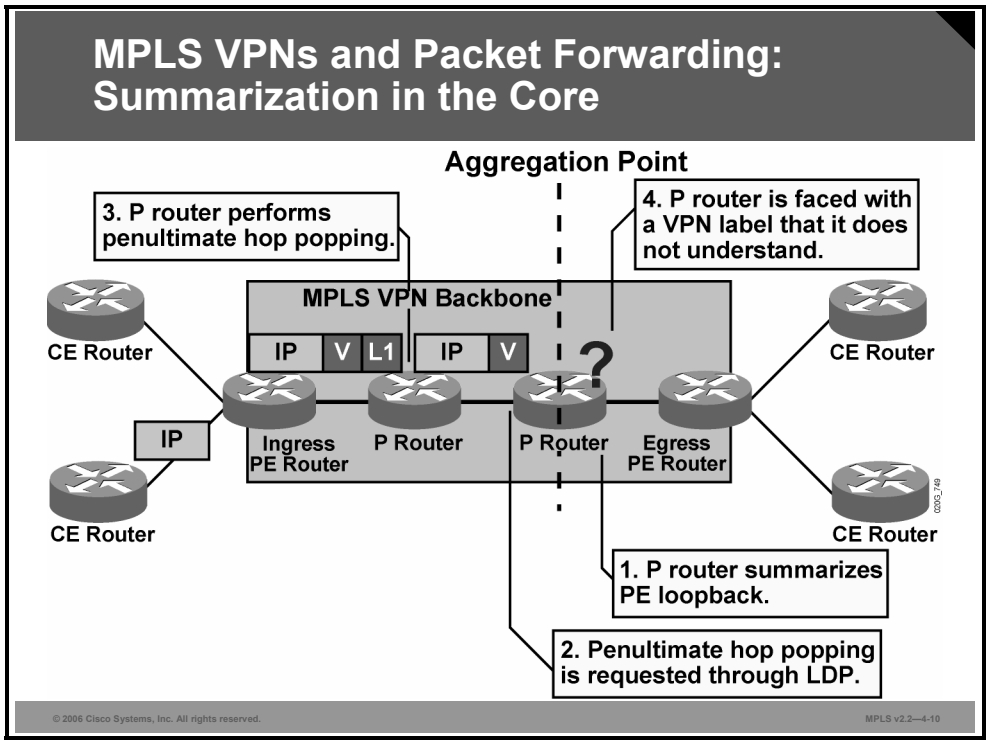

The LSP tunnel is broken at a summarization point, so the summarizing router needs to perform full IP lookup. In a frame-based MPLS network, the P router would request PHP for the summary route, and the upstream P router (or a PE router) would remove the LDP label, exposing the VPN label to the P router. Because the VPN label is assigned not by the P router but by the egress PE router, the label will not be understood by the P router and the VPN packet will be dropped or misrouted.

Summary

This topic summarizes the key points that were discussed in this lesson.

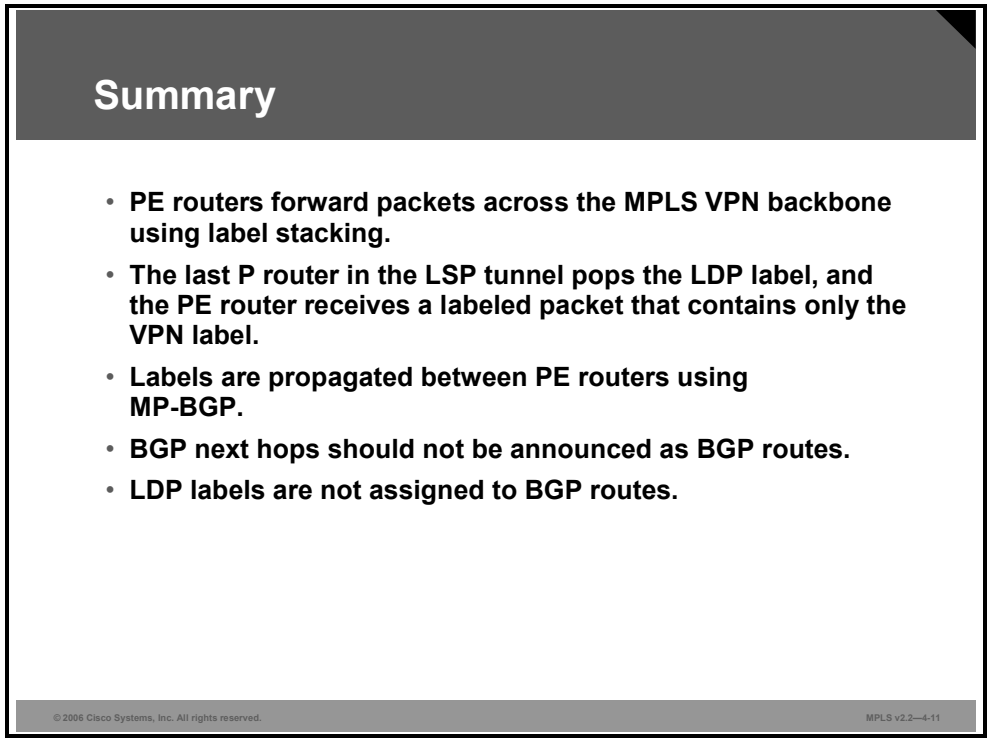

Module Summary

This topic summarizes the key points that were discussed in this module.

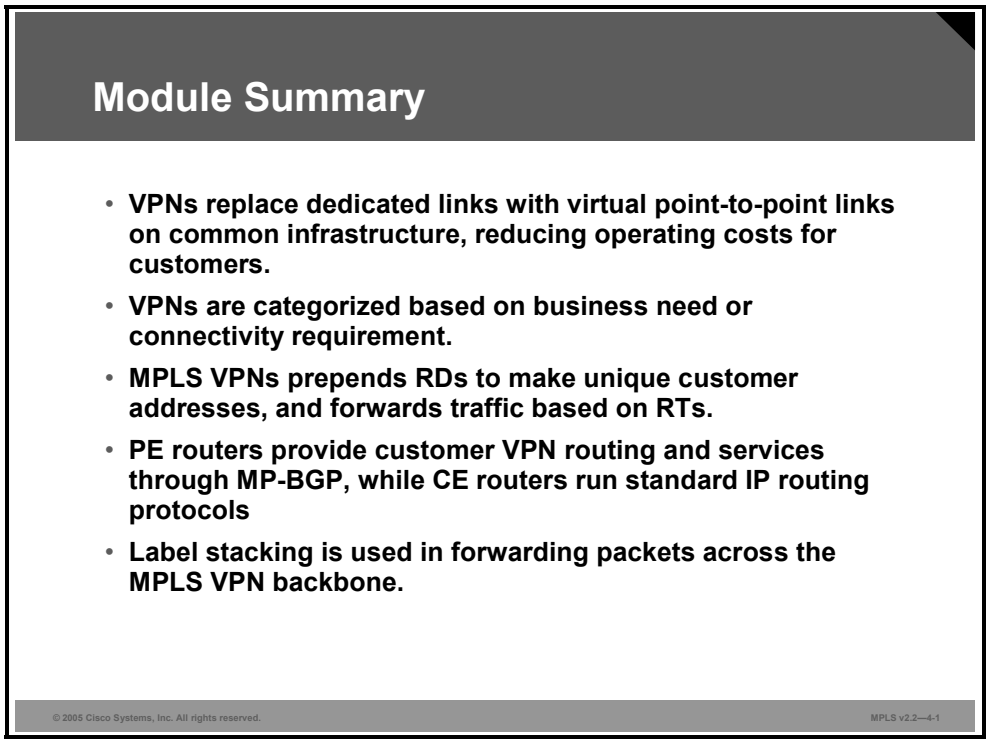

The two major Virtual Private Network (VPN) design options—overlay VPN and peer-to-peer VPN—have many benefits and drawbacks. The VPN topology categories and architectural components help determine the method for forwarding packets in a Multiprotocol Label Switching (MPLS) VPN environment.

References

For additional information, refer to these resources:

- Access Cisco.com for additional information about VPNs.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) Traditional router-based networks that connected customer sties were implemented using which type of links? (Source: Introducing VPNs)

A) PVC
B) dedicated point-to-point
C) SVC
D) emulated point-to-point

Q2) VPNs are implemented using which type of links? (Source: Introducing VPNs)

A) emulated point-to-point
B) dedicated point-to-point
C) PVC
D) PSTN

Q3) Which two network elements are contained in the P-network? (Choose two.) (Source: Introducing VPNs)

A) P device
B) CE device
C) PE device
D) CPE device

Q4) What is a characteristic of an overlay VPN? (Source: Introducing VPNs)

A) PE routers carry all routes from all customers.
B) An overlay VPN guarantees optimum routing between customer sites.
C) The service provider participates in the customer routing.
D) The service provider provides virtual point-to-point links between customer sites.

Q5) In the traditional switched WAN model for Layer 2 VPN implementation, what is the service provider responsible for? (Source: Introducing VPNs)

A) packet filtering
B) transport of Layer 2 frames
C) routing updates
D) encapsulation of protocols

Q6) The peer-to-peer VPN concept was introduced to help overcome what type of drawback? (Source: Introducing VPNs)

_____

_____

_____

Q7) How is a peer-to-peer VPN implemented using packet filters? (Source: Introducing VPNs)

_____

_____

_____

Q8) How do you implement a peer-to-peer VPN based on controlled route distribution? (Source: Introducing VPNs)

_____

_____

_____

Q9) Which VPN type does NOT require the service provider to participate in customer routing? (Source: Introducing VPNs)

A) overlay
B) peer-to-peer
C) central services
D) access VPNs

Q10) For which VPN type is it easier to provision an additional VPN? (Source: Introducing VPNs)

A) overlay
B) peer-to-peer
C) central services
D) access VPNs

Q11) Which VPN type requires the PE router to carry all routes from all customers? (Source: Introducing VPNs)

A) overlay
B) peer-to-peer
C) central services
D) access VPNs

Q12) Which VPN type requires the service provider to participate in customer routing? (Source: Introducing VPNs)

A) overlay
B) peer-to-peer
C) central services
D) access VPNs

Q13) Describe the use of address space and packet routing in each of these peer-to-peer implementations. (Source: Introducing VPNs)

Shared PE router

_____

_____

Dedicated PE router

_____

_____

Q14) Which connectivity category should you use if all sites must have connectivity with each other? (Source: Introducing VPNs)

A) simple
B) overlapping
C) peer-to-peer
D) hub-and-spoke
E) central services

Q15) Which connectivity category should you use if all sites must have connectivity to a server provided by the service provider? (Source: Introducing VPNs)

A) simple
B) overlapping
C) peer-to-peer
D) hub-and-spoke
E) central services

Q16) What are the connectivity requirements of a managed network VPN? (Source: Introducing VPNs)

A) The service provider is restricted to access of the P-network.
B) The service provider is granted access to the entire C-network.
C) The service provider is restricted to access of the managed CE routers.
D) The service provider grants the customer access to the PE routers but not the P routers.

Q17) Which VPN topology has many sites connecting to a central site? (Source: Categorizing VPNs)

A) simple
B) overlapping
C) peer-to-peer
D) hub-and-spoke
E) central services

Q18) When you are using a dynamic routing protocol such as RIP in a redundant hub-and-spoke topology, which statement is true? (Source: Categorizing VPNs)

A) Static routing must be used to provide connectivity from remote site to remote site.
B) Split-horizon updates must be disabled at the hub router if static routing is used.
C) Split-horizon updates must be disabled at the hub router if point-to-point subinterfaces are not used.
D) Split-horizon updates must be enabled at the remote site router when point-to-point subinterfaces are not used.

Q19) Identify the criteria that a customer should consider when determining where virtual circuits are established in a partial mesh topology. (Source: Categorizing VPNs)

_____

_____

_____

Q20) Which component of the VPN business category is used to connect different organizations? (Source: Categorizing VPNs)

A) intranet VPNs
B) Internet VPNs
C) access VPNs
D) extranet VPNs

Q21) Which component of the VPN business category relies on security mechanisms to ensure protection of participating individual organizations? (Source: Categorizing VPNs)

A) intranet VPNs
B) Internet VPNs
C) access VPNs
D) extranet VPNs

Q22) Which implementation of the VPN business category provides the most cost-effective model? (Source: Categorizing VPNs)

A) overlay
B) peer-to-peer
C) central services
D) access VPNs

Q23) Which component of the VPN connectivity category provides full connectivity between sites? (Source: Categorizing VPNs)

A) simple
B) overlapping
C) central services
D) managed services

Q24) Describe the connectivity in a central services extranet. (Source: Categorizing VPNs)

Q25) Describe the connectivity in a managed network VPN. (Source: Categorizing VPNs)

Q26) Which routers are MPLS VPNs aware of? (Source: Introducing MPLS VPN Architecture)

Q27) Which traditional VPN module can the architecture of a PE router in an MPLS VPN be compared to? (Source: Introducing MPLS VPN Architecture)

Q28) Which protocol is used to transport customer routes directly between PE routers? (Source: Introducing MPLS VPN Architecture)

A) RIP
B) VPN
C) BGP
D) OSPF

Q29) What is the function of the RD in an MPLS VPN? (Source: Introducing MPLS VPN Architecture)

Q30) What is the function of the RT in MPLS VPNs? (Source: Introducing MPLS VPN Architecture)

Q31) How has the introduction of complex VPN topologies redefined the meaning of a VPN? (Source: Introducing MPLS VPN Architecture)

Q32) What could happen if two different sites with different requirements are associated with the same virtual routing table? (Source: Introducing MPLS VPN Architecture)

Q33) In which two ways do MPLS VPNs support overlapping customer address spaces? (Choose two.) (Source: Introducing MPLS VPN Architecture)

A) by implementing unique RDs for each customer
B) by implementing unique RTs for each customer
C) by implementing different LSPs for each customer
D) by implementing virtual routing spaces for each customer

Q34) Which statement is true if you use the P-network IGP to propagate customer routing information across the P-network? (Source: Introducing MPLS VPN Architecture)

A) The PE router must be VPN-aware.
B) The P router must be VPN-aware.
C) Customers can use overlapping address spaces.
D) The P router must carry all of the customer routes.

Q35) Why do MPLS VPNs implement route targets? (Source: Introducing MPLS VPN Architecture)

A) to identify different customer VPNs
B) to allow a site to participate on more than one VPN
C) to convert a customer address to an MP-BGP address
D) to convert a non-unique IP address into a unique VPNv4 address

Q36) Which routing protocol does the CE router run? (Source: Introducing the MPLS VPN Routing Model)

A) any IP routing protocol
B) any VPN-aware BGP protocol
C) any VPN-aware IP routing protocol
D) any VPN-aware link-state protocol

Q37) Which type of routers exchange VPNv4 routes? (Source: Introducing the MPLS VPN Routing Model)

A) P
B) CE
C) PE

Q38) Which protocol would a PE router use to support an existing Internet routing scheme? (Source: Introducing the MPLS VPN Routing Model)

A) IS-IS
B) EIGRP
C) BGP IPv4
D) BGP VPNv4

Q39) Identify the routing tables implemented in the PE router to support an MPLS VPN and describe their contents. (Source: Introducing the MPLS VPN Routing Model)

Q40) What BGP function do MPLS VPNs use to transport RTs? (Source: Introducing the MPLS VPN Routing Model)

Q41) How does the PE router know in which VRF table to install received routes for a customer? (Source: Introducing the MPLS VPN Routing Model)

Q42) What is the impact of an MPLS VPN on CE routers? (Source: Introducing the MPLS VPN Routing Model)

A) The CE routers must support BGP.
B) The CE routers must run a link-state protocol.
C) The CE routers can run any standard IP routing protocol.
D) The IGP of the CE routers must be upgraded to a VPN-aware IGP.

Q43) Why would IPv4 routing be enabled on the PE router? (Source: Introducing the MPLS VPN Routing Model)

A) to support the MPLS VPN route update
B) to support the MPLS VPN route target exports
C) to support an existing Internet routing scheme
D) to support the transport of MP-BGP extended communities

Q44) Which two types of routes would an MPLS VPN install into the VRF? (Choose two.) (Source: Introducing the MPLS VPN Routing Model)

A) those routes received via an IPv4 update
B) those routes received via a VPNv4 update
C) those routes received via the core IGP update
D) those routes received via the customer IGP update

Q45) What will happen if the SOO attached to the route is equal to the SOO attribute associated with the CE router? (Source: Introducing the MPLS VPN Routing Model)

A) The route will not be inserted into the VRF.
B) The route will not be inserted into the global table.
C) The route will be inserted into a VRF but not propagated to a CE router.
D) The route will be inserted into a VRF but not propagated to neighboring PE routers.

Q46) Why does the label stack contain two labels when supporting MPLS VPNs? (Source: Forwarding MPLS VPN Packets)

Q47) Why is the VPN label not popped during the PHP process? (Source: Forwarding MPLS VPN Packets)

Q48) Which protocol is used to transport VPN labels between PE routers? (Source: Forwarding MPLS VPN Packets)

A) LDP
B) RSVP
C) MP-BGP
D) the core IGP

Q49) In MPLS VPNs, why must the BGP next hop be set to the egress router in all MP-IBGP updates? (Source: Forwarding MPLS VPN Packets)

Q50) What scenarios would cause the LSP tunnel between PE routers to break? (Source: Forwarding MPLS VPN Packets)

Q51) How can P routers forward VPN packets if they do not have VPN routes? (Source: Forwarding MPLS VPN Packets)

A) They forward based upon the LSP label.
B) They forward based upon the VPN label.
C) They forward based upon the MP-BGP next hop.
D) They forward based upon a routing table lookup of the IP address.

Q52) Which router assigns the VPN label? (Source: Forwarding MPLS VPN Packets)

A) P
B) egress CE
C) egress PE
D) ingress CE
E) ingress PE

Q53) What is used to identify the label that will be used to transport the VPN packet to the egress router? (Source: Forwarding MPLS VPN Packets)

A) the IGP least-cost path
B) the EBGP next-hop address
C) the MP-IBGP next-hop address
D) the VPN label entry in the LFIB

Q54) What is the impact of changing a BGP next hop on an MP-BGP update at confederation boundaries? (Source: Forwarding MPLS VPN Packets)

A) The packet will be forwarded but over a suboptimal route.
B) Packet forwarding for the affected destination will be interrupted.
C) The first P router of the confederation that receives the packet will have to perform a routing table lookup to identify the MP-IBGP next hop.
D) The ingress PE router will forward an MPLS packet to the router identified as the next hop, where it will be converted to an IP packet and forwarded via MP-IBGP.

Module Self-Check Answer Key

Q1) B

Q2) A

Q3) A, C

Q4) D

Q5) B

Q6) the need for customers to establish point-to-point links or virtual circuits between sites

Q7) The service provider allocates portions of its address space to the customers and manages the packet filters on the PE routers to ensure full reachability between sites of a single customer and isolation between customers.

Q8) The core service provider routers (P routers) contain all customer routes, and the PE routers contain only routes of a single customer.

Q9) A

Q10) B

Q11) B

Q12) B

Q13) Shared PE router: All customers share the same (provider-assigned or public) address space. The PE router contains all customer routes. Packet filters are used to provide isolation between customers.

Dedicated PE router: All customers share the same address space. The P routers contain all customer routes. A route filter is used to forward the routes of each customer to the dedicated PE router of that customer.

Q14) A

Q15) E

Q16) C

Q17) D

Q18) C

Q19) The virtual circuits in a partial mesh can be established based on a wide range of criteria, such as traffic pattern between sites, availability of physical infrastructure, and cost considerations.

Q20) D

Q21) D

Q22) B

Q23) A

Q24) All customer sites can connect to the server sites.
All server sites cannot connect to the customer sites.
Customer sites can connect to each other.

Q25) Dedicated virtual circuits are deployed between any managed CE router and the central NMS router.

Q26) P routers

Q27) the dedicated PE router peer-to-peer model

Q28) C

Q29) The RD is used to transform the non-unique IP addresses of the customer into unique VPNv4 addresses.

Q30) The RT attaches a set of VPN identifiers to a route that indicate its membership in several VPNs. This capability allows one site to be a member of more than one VPN.

Q31) A site can be part of more than one VPN, resulting in differing routing requirements for sites that belong to a single VPN and those belonging to multiple VPNs.

Q32) Some of the sites might be able to access destinations that they should not be able to access.

Q33) A, D

Q34) D

Q35) B

Q36) A

Q37) C

Q38) C

Q39) global IP routing table—contains all core IGP routes and the IPv4 routes; VRFs—contain CE routes and VPNv4 routes

Q40) extended communities

Q41) Customer routes are identified by the RT contained in the extended BGP community.

Q42) C

Q43) C

Q44) B, D

Q45) C

Q46) The first label indicates the LSP that will be used to reach the egress router. The second label indicates the VPN that the packet belongs to.

Q47) The egress router needs the label to identify which VPN the packet belongs to.

Q48) C

Q49) The BGP next hop is used to identify which LSP will be used to get to the egress router. If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label and the label stack will not be built correctly.

Q50) If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label and the label stack will not be built correctly.
If the P routers perform summarization of the address range within which the IP address of the egress PE router lies, the LSP tunnel will be disrupted at the summarization point.

Q51) A

Q52) C

Q53) C

Q54) B

