MPLS

Implementing Cisco MPLS

Volume 1

Version 2.1

Student Guide

Text Part Number: ILSG Production Services: 11.18.04

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site 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 ***1***

Overview 1
Learner Skills and Knowledge 2
Course Goal and Objectives 3
Course Flow 4
Additional References 5
Cisco Glossary of Terms 5
Your Training Curriculum 6

MPLS Concepts ***1-1***

Overview 1-1
Module Objectives 1-1

Introducing Basic MPLS Concepts **1-3**

Overview 1-3
Objectives 1-3
What Are the Drawbacks of Traditional IP Routing? 1-4
Example: Traditional IP Forwarding 1-5
Example: IP over ATM 1-6
Example: Traffic Engineering 1-7
What Are the Basic MPLS Features? 1-8
Example: MPLS Concepts 1-9
What Are the Differences Between MPLS and IP over ATM? 1-10
Example: MPLS vs. IP over ATM 1-10
What Is Traffic Engineering? 1-11
What Are the MPLS Architecture Components? 1-12
Example: Control Plane Components 1-13
What Are MPLS Labels? 1-14
Example: MPLS Labels—Frame-Mode MPLS 1-16
What Are the Label Switch Router Functions? 1-18
Summary 1-23

Introducing MPLS Labels and Label Stack **1-25**

Overview 1-25
Objectives 1-25
Where Are MPLS Labels Inserted? 1-26
What Is the MPLS Label Format? 1-27
What Is an MPLS Label Stack? 1-28
What Is MPLS Forwarding? 1-30
Example: MPLS Network—Frame-Mode MPLS 1-31
Summary 1-33

Identifying MPLS Applications **1-35**

Overview 1-35
Objectives 1-35
Which Applications Are Used with MPLS? 1-36
What Is Unicast IP Routing? 1-37
What Is Multicast IP Routing? 1-38
Using MPLS Traffic Engineering 1-39
What Is Quality of Service? 1-40
What Are Virtual Private Networks? 1-41
What Are the Interactions Between MPLS Applications? 1-42
Example: Interactions Between MPLS Applications 1-42
Summary 1-43
Module Summary 1-44
References 1-44
Module Self-Check 1-45
Module Self-Check Answer Key 1-49

Label Assignment and Distribution 2-1

Overview 2-1
Module Objectives 2-1

Introducing Typical Label Distribution in Frame-Mode MPLS 2-3

Overview 2-3
Objectives 2-3
Propagating Labels Across a Network 2-4
Example: Building Blocks for IP Forwarding 2-5
Example: Using the FIB Table to Forward Packets 2-6
Example: Using LDP 2-7
What Are Label-Switched Paths? 2-8
Example: IGP Propagates Routing Information 2-9
Example: LFIB and LIB Tables 2-10
Propagating Labels Using PHP 2-11
Example: PHP—Before 2-11
Example: PHP—After 2-12
What Is the Impact of IP Aggregation on Label-Switched Paths? 2-14
Example: MPLS IP Aggregation Problem 2-14
Allocating Labels in a Frame-Mode MPLS Network 2-16
Example: Label Allocation 2-17
Distributing and Advertising Labels 2-20
Example: Label Distribution and Advertisement 2-20
Example: Interim Packet Propagation Through an MPLS Network 2-22
Example: LDP Update Sent to All Adjacent Routers 2-23
Populating LFIB 2-25
Example: Populating LFIB 2-25
Propagating Packets Across an MPLS Network 2-26
Example: Packet Propagation Through an MPLS Network 2-26
Detecting Frame-Mode Loops 2-27
Example: Normal TTL Operation 2-28
Example: TTL and Loop Detection 2-29
Example: Traceroute with Disabled TTL Propagation 2-31
Allocating Per-Platform Labels 2-34
Example: Per-Platform Label Allocation 2-34
Summary 2-36

Introducing Convergence in Frame-Mode MPLS 2-37

Overview 2-37
Objectives 2-37
What Is the MPLS Steady-State Operation? 2-38
What Happens in a Link Failure? 2-39
Example: Link Failure Actions 2-39
What Is the Routing Protocol Convergence After a Link Failure? 2-40
Example: Routing Protocol Convergence 2-40
What Is the MPLS Convergence After a Link Failure? 2-41
What Happens in Link Recovery? 2-43
Example: Link Recovery Actions 2-43
Summary 2-46

Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge 2-47

Overview 2-47
Objectives 2-47
What Are Cell-Mode MPLS Network Issues? 2-48
Building the IP Routing Table 2-49
Example: Building the IP Routing Table 2-49
Building the IP Forwarding Table 2-50
Requesting a Label 2-51
Example: Requesting a Label 2-51

Allocating a Label 2-52
Example: Allocating a Table 2-52
Example: Additional LSRs 2-54
What Are Cell Interleave Issues? 2-55
Example: Additional Label Allocation 2-56
What Is VC Merge? 2-57
Example: VC Merge 2-57
Detecting Loops in Cell-Mode MPLS Networks 2-59
Example: LDP Hop Count 2-61
Example: Traceroute Through ATM LSRs 2-62
What Is Per-Interface Label Allocation? 2-65
Summary 2-67
Introducing MPLS Label Allocation, Distribution, and Retention Modes 2-69
Overview 2-69
Objectives 2-69
What Are Label Distribution Parameters? 2-70
What Is Label Space? 2-71
Distributing Labels 2-73
Example: Unsolicited Downstream 2-73
Allocating Labels 2-75
Example: Ordered Control 2-76
Retaining Labels 2-77
Example: Liberal Retention Mode 2-77
Example: Conservative Retention Mode 2-78
What Are Standard Parameter Sets in MPLS Implementation? 2-79
Summary 2-81
Discovering LDP Neighbors 2-83
Overview 2-83
Objectives 2-83
Establishing an LDP Session 2-84
What Are LDP Hello Messages? 2-85
Negotiating Label Space 2-86
Example: Label Space Negotiation 2-87
Discovering LDP Neighbors 2-88
Example: LDP Neighbor Discovery 2-88
Negotiating LDP Sessions 2-89
Establishing LDP Sessions Between ATM LSRs 2-90
Example: LDP Sessions Between ATM LSRs 2-90
Discovering Nonadjacent Neighbors 2-91
Summary 2-92
Module Summary 2-93
References 2-93
Module Self-Check 2-94
Module Self-Check Answer Key 2-101
Frame-Mode and Cell-Mode MPLS Implementation on Cisco IOS Platforms 3-1
Overview 3-1
Module Objectives 3-1
Introducing CEF Switching 3-3
Overview 3-3
Objectives 3-3
What Are Cisco IOS Platform Switching Mechanisms? 3-4
Using Standard IP Switching 3-5
Example: Standard IP Switching 3-5
What Is CEF Switching Architecture? 3-6
Configuring IP CEF 3-7
ip cef 3-7
Syntax Description 3-7

ip route-cache cef 3-8
Syntax Description 3-8
Defaults 3-8
Monitoring IP CEF 3-9
show ip cef 3-9
Summary 3-11
Configuring Frame-Mode MPLS on Cisco IOS Platforms 3-13
Overview 3-13
Objectives 3-13
What Are MPLS Configuration Tasks? 3-14
Configuring the MPLS ID on a Router 3-15
mpls ldp router-id 3-15
Defaults 3-15
Configuring MPLS on a Frame-Mode Interface 3-16
mpls ip 3-16
Syntax Description 3-16
Defaults 3-16
mpls label protocol [tdp | ldp | both] 3-17
Defaults 3-17
Example: Configuring MPLS on a Frame-Mode Interface 3-18
Configuring a Label-Switching MTU 3-20
mpls mtu 3-20
Defaults 3-20
Configuring IP TTL Propagation 3-22
mpls ip propagate-ttl 3-22
Syntax Description 3-22
Defaults 3-22
Example: Configuring IP TTL Propagation 3-23
Example: Disabling IP TTL Propagation 3-24
mpls ip propagate-ttl 3-25
Defaults 3-26
Command Modes 3-26
Usage Guidelines 3-26
Configuring Conditional Label Distribution 3-28
mpls ldp advertise-labels 3-28
Example: Conditional Label Distribution Configuration 3-29
Example: Enabling Conditional Label Advertisement 3-31
Configuring Frame-Mode MPLS on Switched WAN Media 3-32
Summary 3-36
Monitoring Frame-Mode MPLS on Cisco IOS Platforms 3-37
Overview 3-37
Objectives 3-37
Monitoring MPLS 3-38
show mpls ldp parameters 3-38
show mpls interfaces 3-38
show mpls ldp discovery 3-39
show mpls ldp parameters 3-39
Syntax Description 3-39
show mpls interfaces 3-41
show mpls ldp discovery 3-43
Monitoring LDP 3-45
show mpls ldp neighbor 3-45
show mpls ldp bindings 3-46
show mpls ldp neighbor 3-47
Usage Guidelines 3-47
show mpls ldp bindings 3-49
Usage Guidelines 3-50
Examples 3-50

Monitoring Label Switching 3-51
show mpls forwarding-table 3-51
show ip cef 3-51
show mpls forwarding-table 3-52
Examples: show mpls forwarding table Command Output 3-53
show ip cef detail 3-55
Usage Guidelines 3-55
Debugging MPLS and LDP 3-56
debug mpls packets 3-57
Summary 3-58
Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms 3-59
Overview 3-59
Objectives 3-59
What Are Common Frame-Mode MPLS Issues? 3-60
Solving LDP Session Startup Issues 3-61
Solving Label Allocation Issues 3-65
Solving Label Distribution Issues 3-66
Solving Packet Labeling Issues 3-67
show cef interface 3-68
Usage Guidelines 3-68
Solving Intermittent MPLS Failures 3-70
Solving Packet Propagation Issues 3-71
Summary 3-72
Configuring LC-ATM MPLS 3-73
Overview 3-73
Objectives 3-73
What Are the Configuration Tasks for MPLS on LC-ATM Interfaces? 3-74
Configuring an LC-ATM Interface on a Router 3-75
mpls label protocol [tdp | ldp | both] 3-76
Defaults 3-76
Configuring an LC-ATM Interface on a Catalyst ATM Switch 3-77
mpls ip 3-77
mpls label protocol [tdp | ldp | both] 3-78
Configuring MPLS Between a Router and a Switch 3-79
Configuring Additional LC-ATM Parameters 3-80
mpls atm control-vc 3-80
Defaults 3-81
mpls atm vpi 3-81
Defaults 3-81
Example: Configuring Additional LC-ATM Parameters 3-82
mpls ldp atm vc-merge 3-83
mpls ldp maxhops 3-84
Disabling VC Merge 3-85
Summary 3-86
Configuring LC-ATM MPLS over ATM Virtual Path 3-87
Overview 3-87
Objectives 3-87
What Is ATM Virtual Path? 3-88
ATM Virtual Path Usages 3-89
Configuring MPLS over ATM Virtual Path—Switches 3-92
Example: Configuring MPLS over ATM Virtual Path—Switches 3-92
Example: Configuration of Both MPLS-Enabled ATM Switches 3-93
Configuring MPLS over ATM Virtual Path—Routers 3-94
Example: Configuring MPLS Over ATM Virtual Path—Routers 3-94
Summary 3-96

Monitoring LC-ATM MPLS on Cisco IOS Platforms 3-97
Overview 3-97
Objectives 3-97
How to Monitor Specific LC-ATM Label-Switching Functions 3-98
show mpls atm-ldp summary 3-98
show mpls atm-ldp bindings 3-98
show mpls atm-ldp capability 3-99
How to Display Summary Information About ATM Entries 3-100
How to Display Current Label Bindings 3-102
How to Display MPLS ATM Capabilities by LDP 3-104
Debugging Specific ATM LDP Functions 3-106
debug mpls atm-ldp routes 3-106
debug mpls atm-ldp states 3-106
Summary 3-107
Module Summary 3-108
References 3-109
Module Self-Check 3-110
Module Self-Check Answer Key 3-116

MPLS Virtual Private Network Technology 4-1

Overview 4-1
Module Objectives 4-1
Introducing Virtual Private Networks 4-3
Overview 4-3
Objectives 4-3
Traditional Router-Based Network Connectivity 4-4
Advantages of Virtual Private Networks 4-5
Example: Virtual Private Networks 4-5
What Are VPN Network Elements? 4-6
How Are Virtual Circuits Used in Switched WANs? 4-8
Summary 4-9
Introducing Overlay and Peer-to-Peer VPNs 4-11
Overview 4-11
Objectives 4-11
What Are the VPN Implementation Technologies? 4-12
What Are the Overlay VPN Implementation Techniques? 4-13
Example: Overlay VPS—Frame Relay 4-15
What Are the Implementation Techniques for Peer-to-Peer VPNs? 4-19
Example: Controlled Route Distribution 4-21
What Are the Benefits of VPN Implementations? 4-22
What Are the Drawbacks of VPN Implementations? 4-23
What Are the Drawbacks of Traditional Peer-to-Peer VPNs? 4-24
Summary 4-25
Categorizing VPNs 4-27
Overview 4-27
Objectives 4-27
What Are the Overlay VPN Categories? 4-28
What Is the Hub-and-Spoke Overlay VPN Topology? 4-29
What Is the Partial Mesh Overlay VPN Topology? 4-31
What Are the VPN Business Categories? 4-32
What Are Extranet VPNs? 4-33
Example: Overlay VPN—Extranet VPNs 4-33
Example: Peer-to-Peer VPN—Extranet VPNs 4-34
What Is the VPN Connectivity Category? 4-35
What Is the Central Services Extranet? 4-36
Example: Central Services Extranet 4-36

Example: Hybrid Implementation 4-37
What Is a Managed Network Implementation? 4-38
Summary 4-39
Introducing MPLS VPN Architecture 4-41
Overview 4-41
Objectives 4-41
What Is the MPLS VPN Architecture? 4-42
What Is the Architecture of a PE Router in an MPLS VPN? 4-44
What Are the Methods of Propagation Across the P-Network? 4-45
What Are Route Distinguishers? 4-50
What Are Route Targets? 4-54
Example: VoIP Service Sample 4-54
Example: Connectivity Requirements 4-55
What Is the New Meaning of VPNs? 4-59
What Is the Impact of Complex VPN Topologies on Virtual Routing Tables? 4-60
Example: Impact of Complex VPN Topologies on Virtual Routing Tables 4-61
Summary 4-62
Introducing the MPLS VPN Routing Model 4-63
Overview 4-63
Objectives 4-63
MPLS VPN Routing Requirements 4-64
What Is the MPLS VPN Routing Model? 4-65
Existing Internet Routing Support 4-69
Routing Tables on PE Routers 4-70
Identifying End-to-End Routing Update Flow 4-71
Example: End-to-End Routing Update Flow 4-71
Route Distribution to CE Routers 4-75
Summary 4-76
Forwarding MPLS VPN Packets 4-77
Overview 4-77
Objectives 4-77
What Are the End-to-End VPN Forwarding Mechanisms? 4-78
What Is VPN PHP? 4-80
Propagating VPN Labels Between PE Routers 4-81
Example: VPN Label Propagation Between PE Routers 4-82
What Are the Effects of MPLS VPNs on Label Propagation? 4-84
What Are the Effects of MPLS VPNs on Packet Forwarding? 4-85
Example: Summarization in the Core 4-86
Summary 4-87
Module Summary 4-88
References 4-88
Module Self-Check 4-89
Module Self-Check Answer Key 4-98

Course Introduction

Overview

Service providers today 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 service providers 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) v2.1 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 in order to benefit fully from this course.

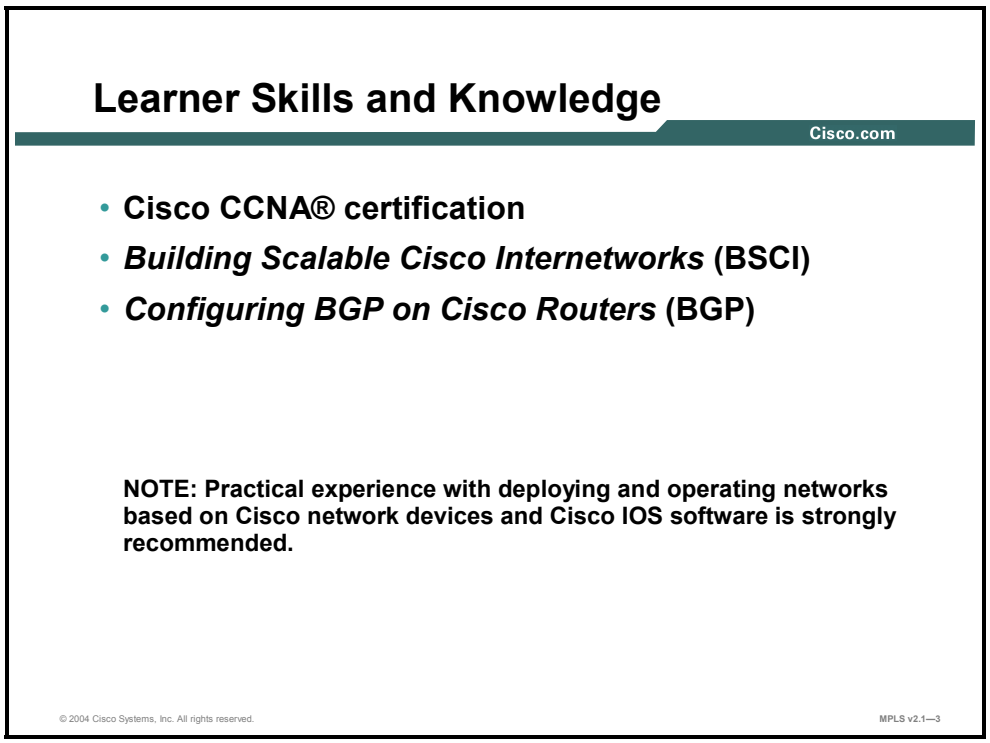

Course Goal and Objectives

This topic describes the course goal and objectives.

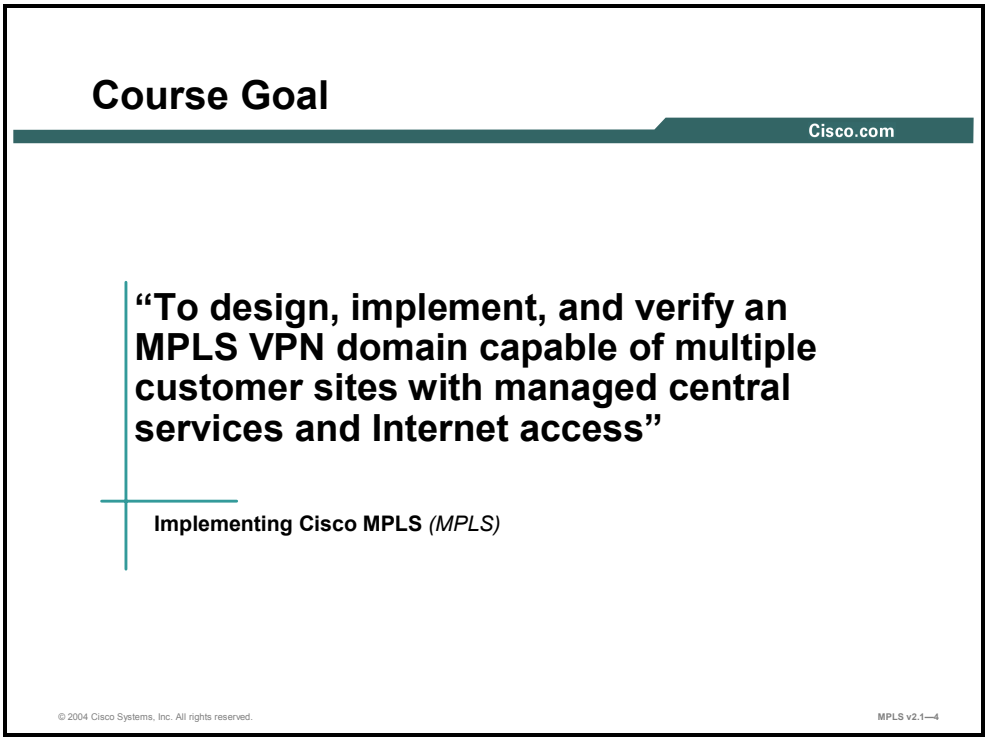

Upon completing this course, you will be able to meet these objectives:

- Describe the features of MPLS
- Describe how MPLS labels are assigned and distributed
- Describe the tasks and commands necessary to implement MPLS on frame-mode and LC-ATM 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

Course Flow

This topic presents the suggested flow of the course materials.

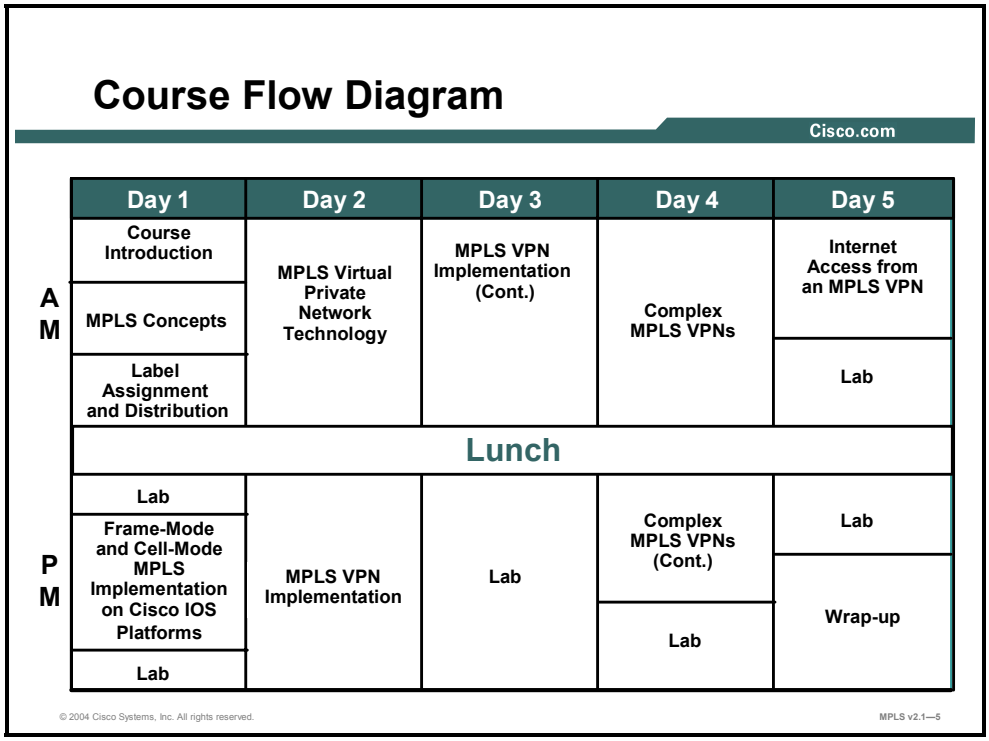

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

Additional References

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

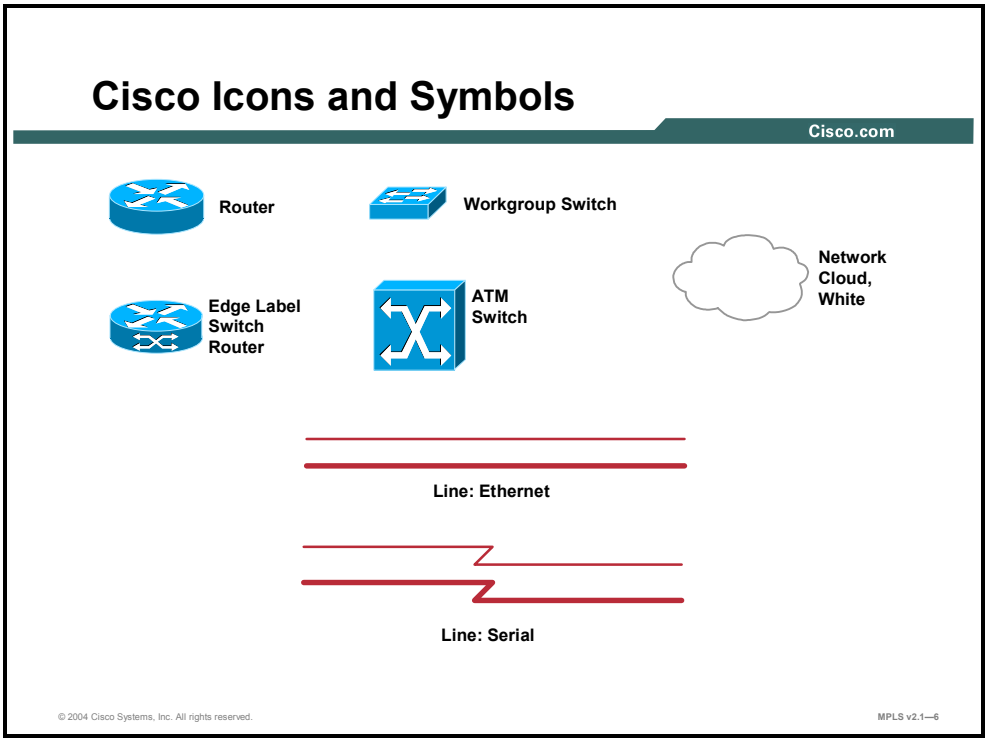

Cisco Glossary of Terms

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.

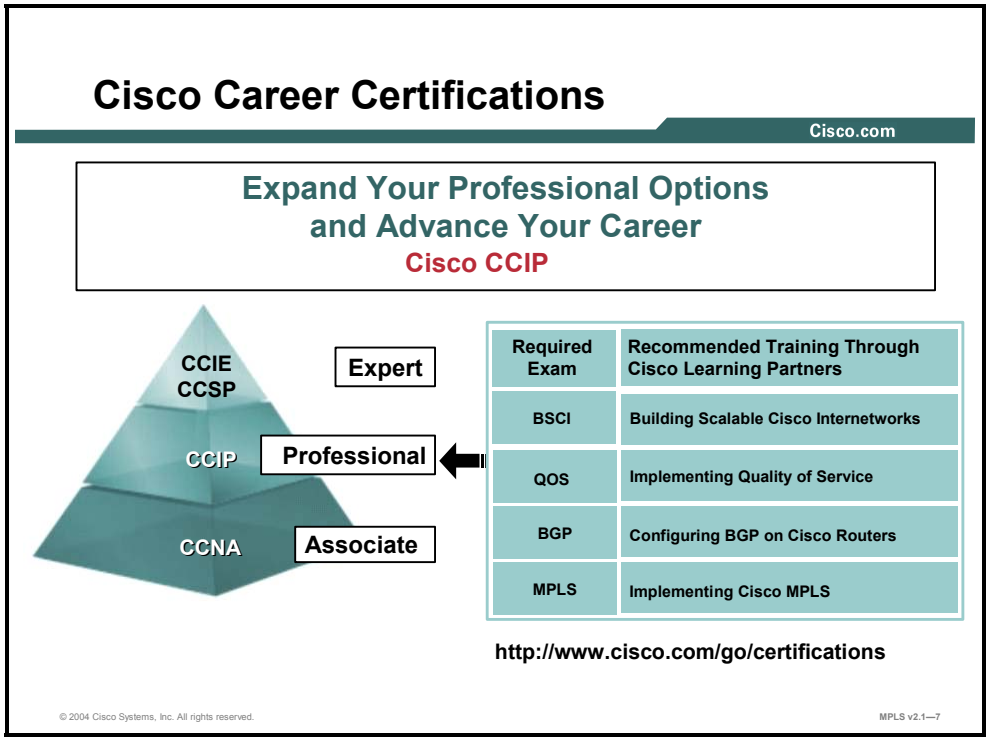

You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP™, or CCSP®). It provides a gathering place for Cisco-certified professionals to share questions, suggestions, 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 ATM and 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 MPLS. The lesson provides information about some of the MPLS components and labels. This lesson lays the foundation for subsequent lessons that cover key areas, such as Cisco MPLS Traffic Engineering (MPLS TE) 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 the learner have a clear picture of how to differentiate between the roles of certain devices and 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 the drawbacks in traditional IP routing. This ability includes being able to meet these objectives:

- Describe the drawbacks of traditional IP routing
- Describe the basic features of MPLS
- Describe the differences between MPLS and IP over ATM
- Describe the features of traffic engineering
- Describe the main components of the MPLS architecture
- Describe the features of MPLS labels
- Describe the function of the different types of LSRs

What Are the Drawbacks of Traditional IP Routing?

This topic describes the drawbacks of traditional IP routing.

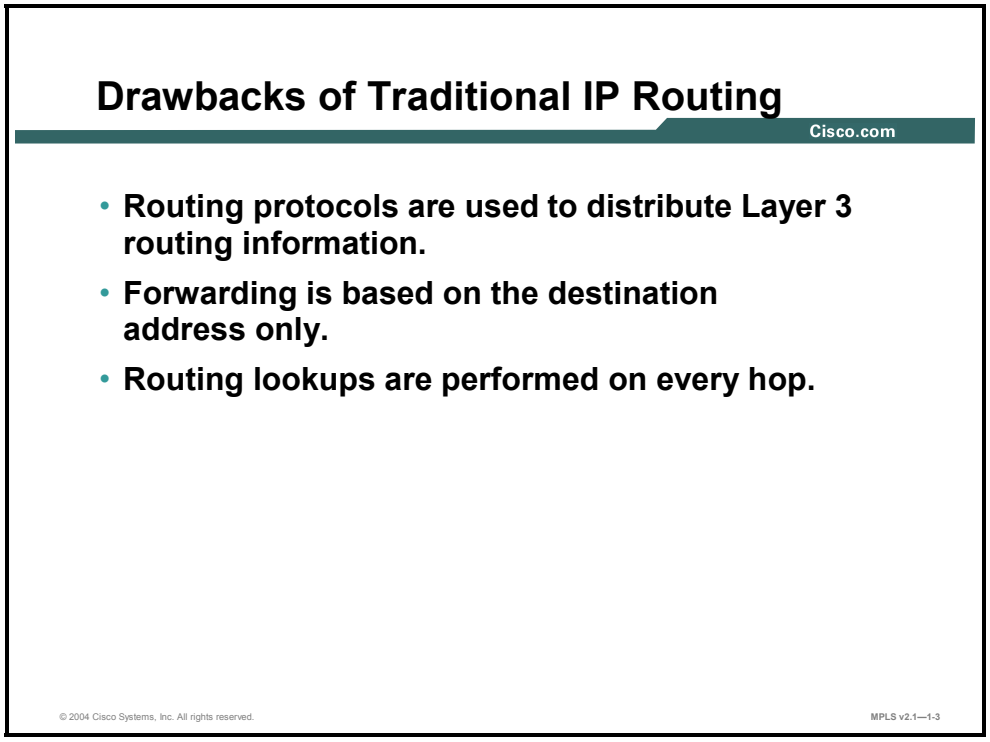

Before basic MPLS functionality is explained, the following three drawbacks of traditional IP routing need to be highlighted:

- Routing protocols are used on all devices to distribute routing information.
- Regardless of the routing protocol, routers always forward packets based on the destination address only. The only exception is policy-based routing (PBR), which bypasses the destination-based routing lookup.
- Routing lookups are performed on every router. Each router in the network makes an independent decision when forwarding packets.

MPLS helps reduce the number of routing lookups and can change the forwarding criteria. This capability eliminates the need to run a particular routing protocol on all the devices.

Example: Traditional IP Forwarding

The figure shows how routers in a service provider network forward packets based on their destination addresses. The figure also shows that all the routers need to run a routing protocol—Border Gateway Protocol (BGP)—to get the entire Internet routing information.

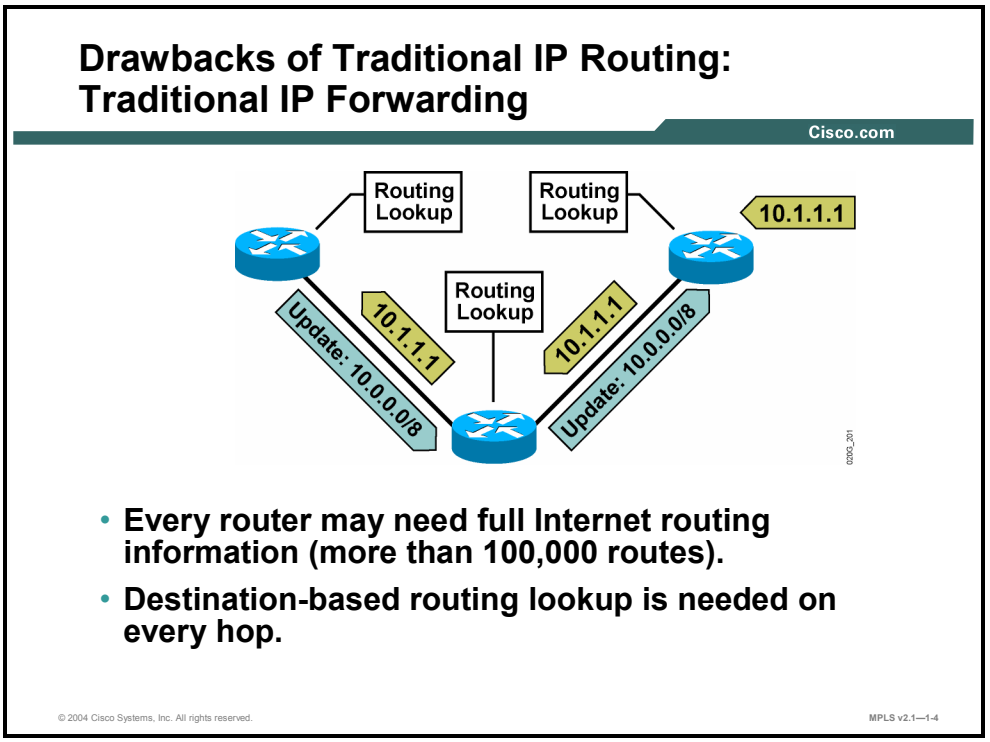

Every router in the path performs a destination-based routing lookup in a large forwarding table. Forwarding complexity is usually related to the size of the forwarding table and to the switching mechanism.

Example: IP over ATM

The figure shows a worst-case scenario where Layer 2 and Layer 3 topologies do not overlap.

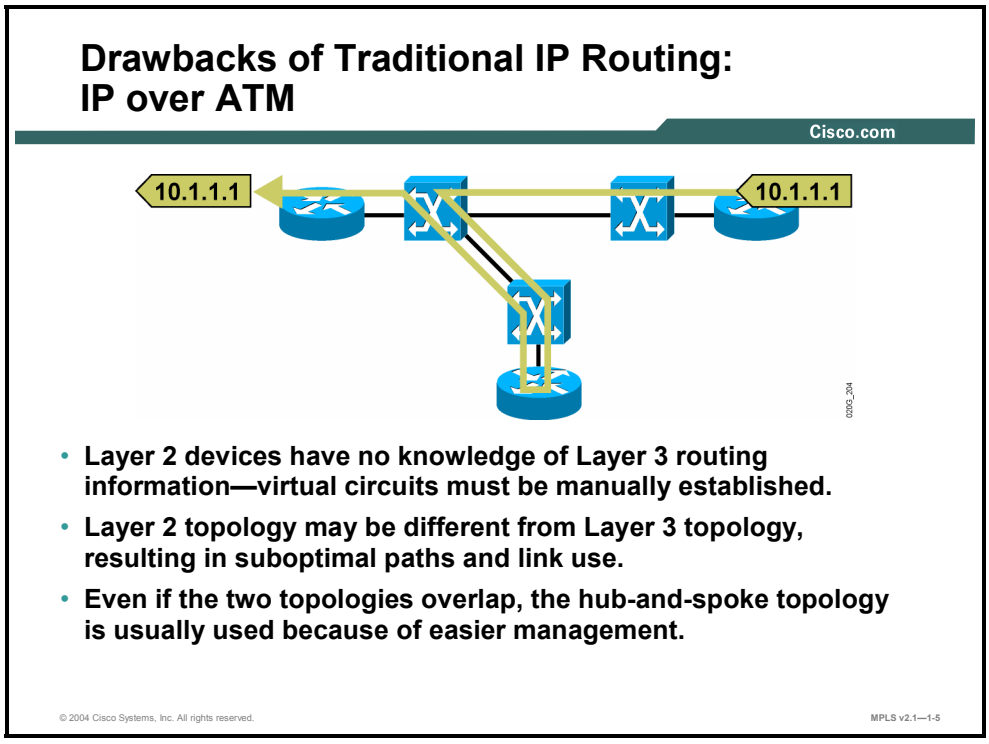

The result is that a single packet, which could be propagated with three Layer 2 hops, instead requires seven hops. The reason for this is that Layer 2 devices have static information about how to interconnect Layer 3 devices. Routers use a routing protocol to propagate Layer 3 routing information through the intermediary router.

Example: Traffic Engineering

The figure shows a topology with unequal links.

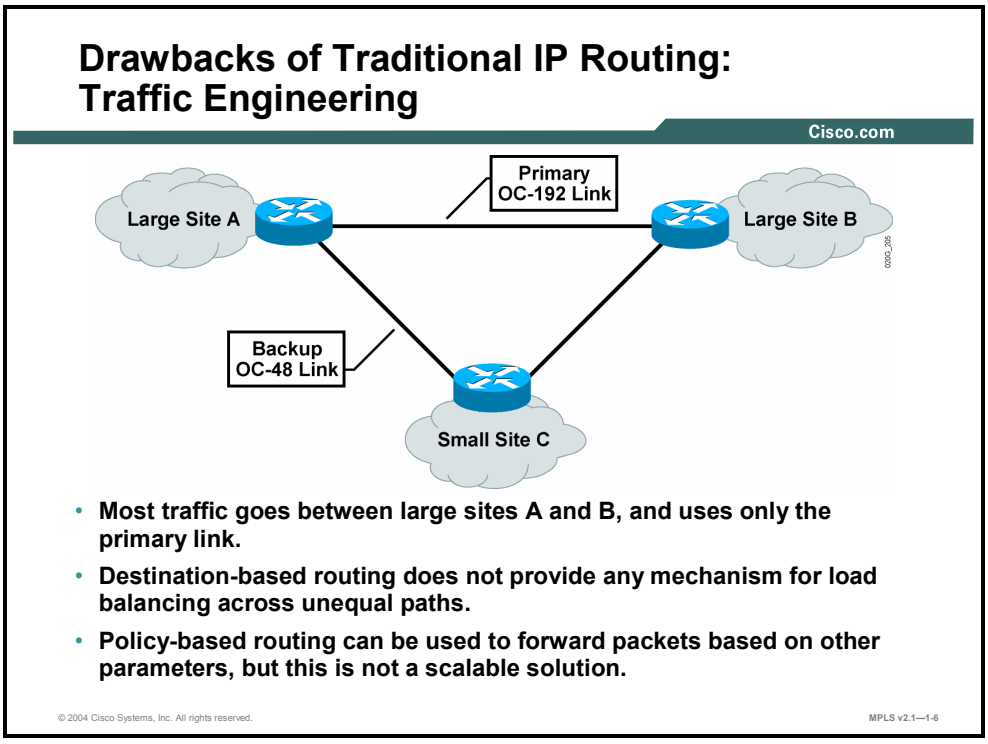

Traffic patterns illustrate that most of the traffic goes between sites A and B. Traditional IP forwarding does not have a scalable mechanism to allow use of the backup link. This situation results in unequal load balancing.

What Are the Basic MPLS Features?

This topic describes the basic features of MPLS.

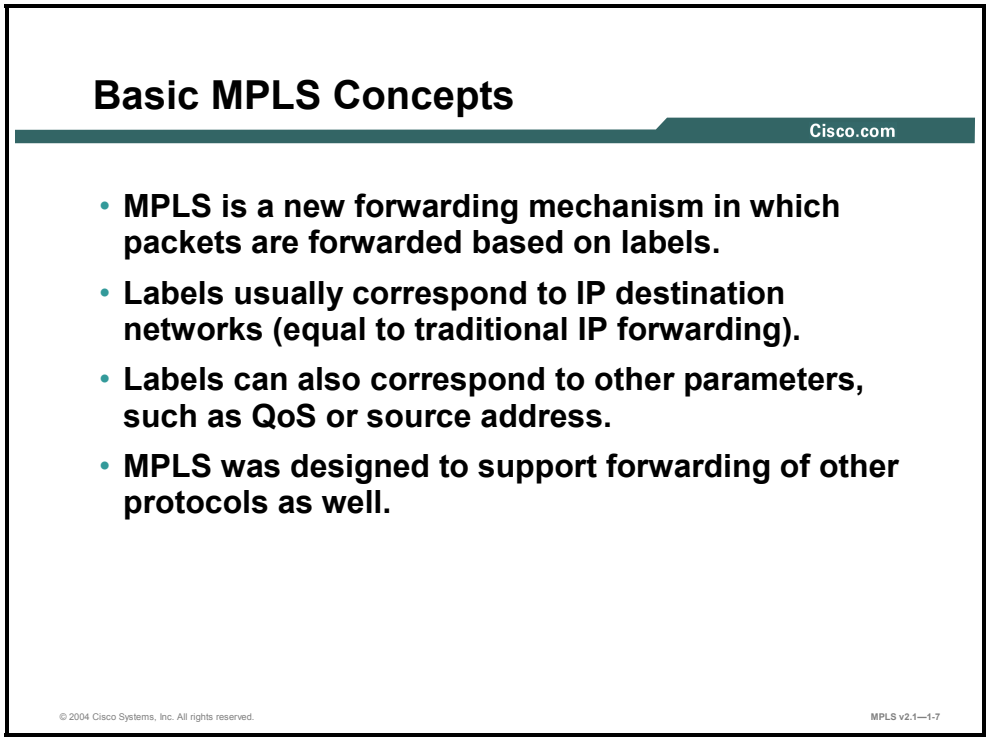

MPLS is a new switching mechanism that uses labels (numbers) to forward packets.

Labels usually 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.

MPLS was designed to support forwarding of other protocols as well. Label switching is performed regardless of the Layer 3 protocol.

Example: MPLS Concepts

The figure illustrates a situation in which the intermediary router 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 received label (23).

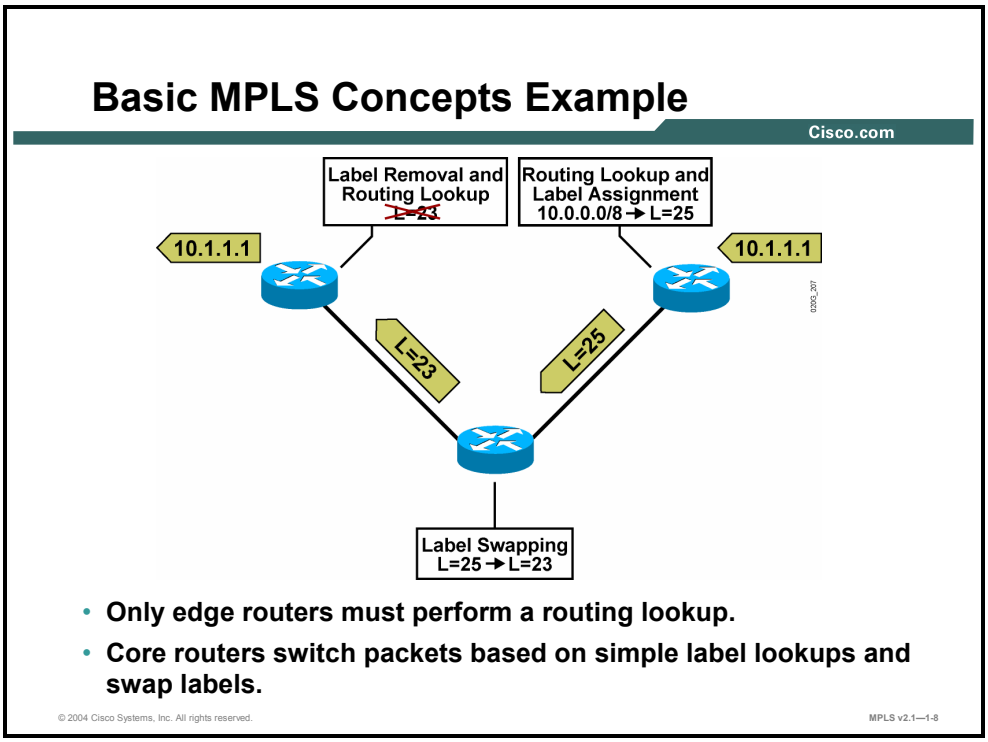

In larger networks, the result of MPLS labeling is that only the edge routers perform a routing lookup. All the core routers forward packets based on the labels.

What Are the Differences Between MPLS and IP over ATM?

This topic describes the differences between MPLS and IP over ATM.

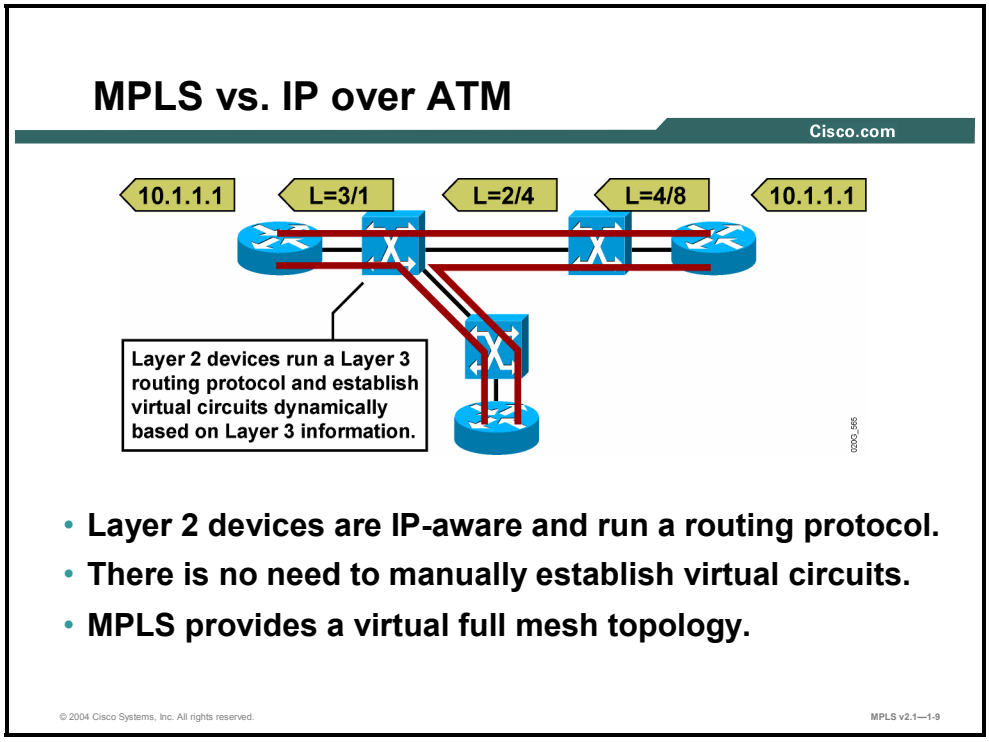

MPLS is used in ATM networks to provide optimal routing across Layer 2 ATM switches.

Example: MPLS vs. IP over ATM

For MPLS to work with ATM switches, the switches must be Layer 3-aware. In other words, ATM switches must run a Layer 3 routing protocol.

Another benefit of this setup is that there is no longer a need to manually establish virtual circuits. ATM switches automatically create a full mesh of virtual circuits based on Layer 3 routing information.

What Is Traffic Engineering?

This topic describes the features of traffic engineering (TE).

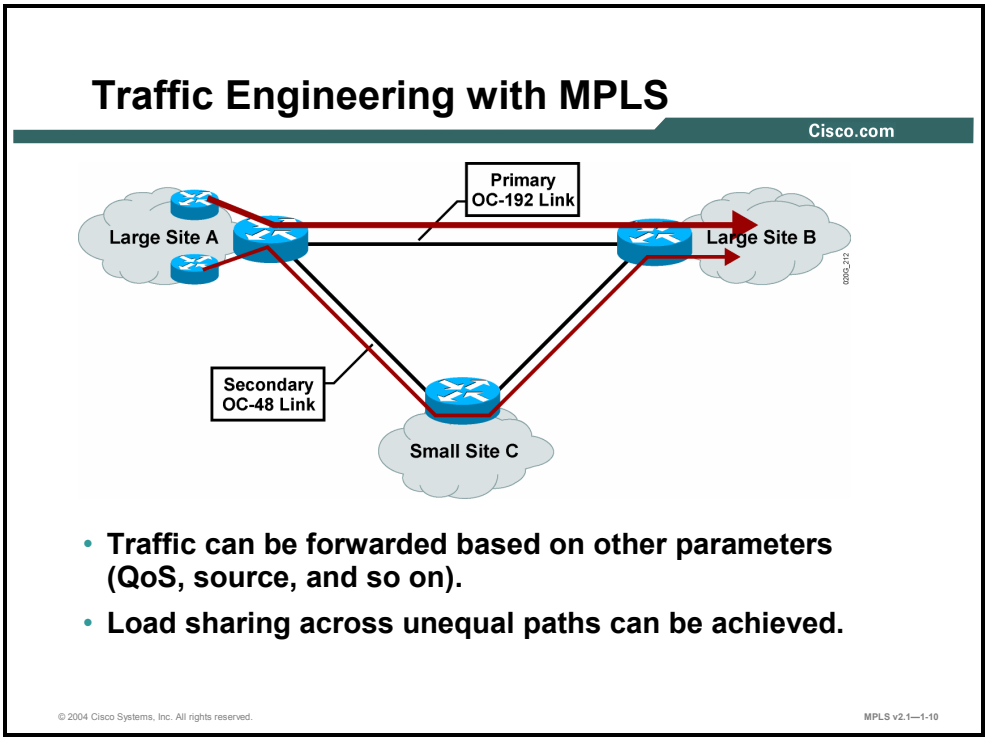

MPLS also supports TE. Traffic-engineered tunnels can be created based on traffic analysis to provide load balancing across unequal paths.

Multiple TE tunnels can lead to the same destination but can use different paths. Traditional IP forwarding would force all traffic to use the same path based on the destination-based forwarding decision. TE determines the path at the source based on additional parameters, such as available resources and constraints in the network.

What Are the MPLS Architecture Components?

This topic describes the main components of the MPLS architecture.

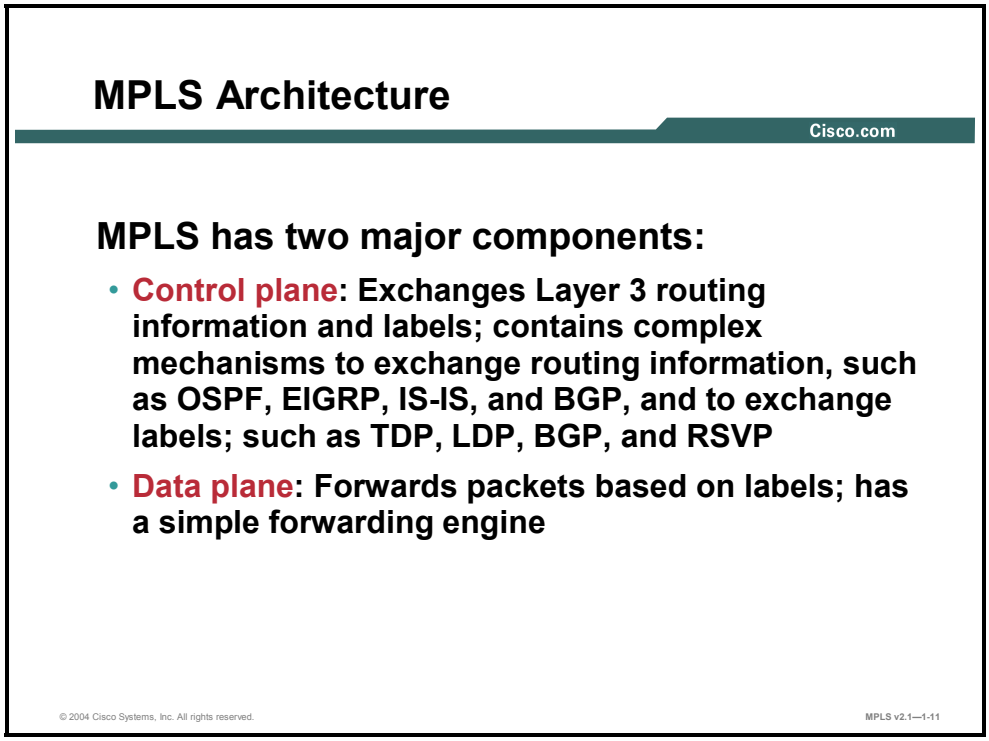

MPLS consists of the following two major components:

- **Control plane:** The control plane takes care of the routing information exchange and the label exchange between adjacent devices
- **Data plane:** The data plane takes care of forwarding based on either destination addresses or labels; this is also known as the forwarding plane.

A large number of different 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.

The control plane also requires protocols such as the label exchange protocols, Tag Distribution Protocol (TDP), MPLS Label Distribution Protocol (LDP), BGP (used by MPLS VPN), to exchange labels. Resource Reservation Protocol (RSVP) is used by MPLS TE to accomplish this exchange.

The data plane, however, is a simple label-based forwarding engine that is independent of the type of routing protocol or label exchange protocol. The label forwarding information base (LFIB) table is used to forward packets based on labels. The LFIB table is populated by the label exchange protocols (TDP or LDP, or both) used.

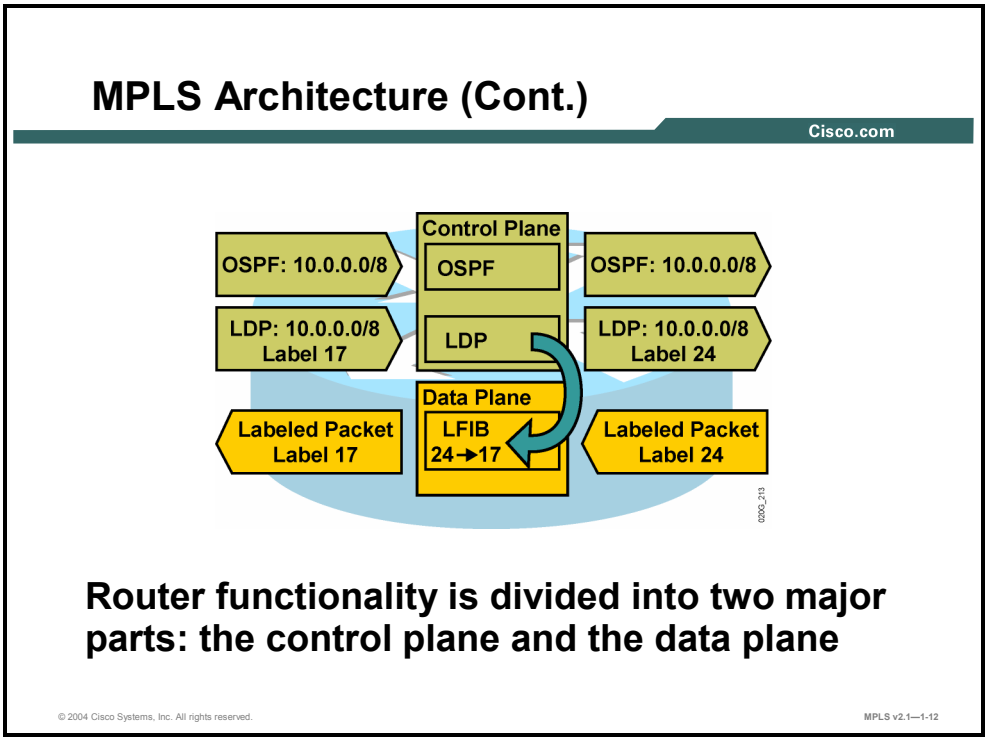

MPLS implements destination-based forwarding that uses labels to make forwarding decisions.

A Layer 3 routing protocol is still needed to propagate Layer 3 routing information. A label exchange mechanism is simply an add-on to propagate labels that are used for Layer 3 destinations.

Example: Control Plane Components

The figure illustrates the two components of the control plane.

- OSPF, which receives and forwards IP network 10.0.0.0/8.
- LDP, which receives label 17 to be used for packets with destination address 10.x.x.x. A local label 24 is generated and sent to upstream neighbors so that these neighbors can label packets with the appropriate label. LDP inserts an entry into the data plane LFIB table, where label 24 is mapped to label 17.

The data plane then forwards all packets with label 24 through the appropriate interfaces and replaces label 24 with label 17.

What Are MPLS Labels?

This topic describes the features of MPLS labels.

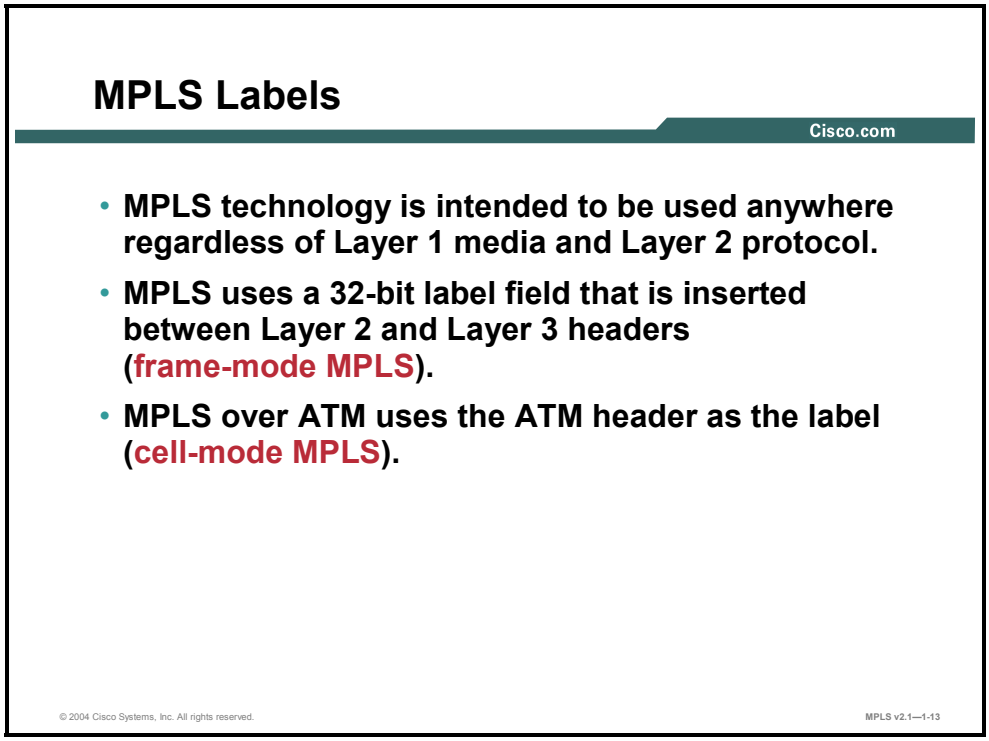

MPLS is designed for use on virtually any media and Layer 2 encapsulation. Most Layer 2 encapsulations are frame-based, and MPLS simply inserts a 32-bit label between the Layer 2 and Layer 3 headers (“frame-mode” MPLS).

ATM is a special case where fixed-length cells are used and a label cannot be inserted on every cell. MPLS uses the virtual path identifier/virtual channel identifier (VPI/VCI) fields in the ATM header as a label (“cell-mode” MPLS).

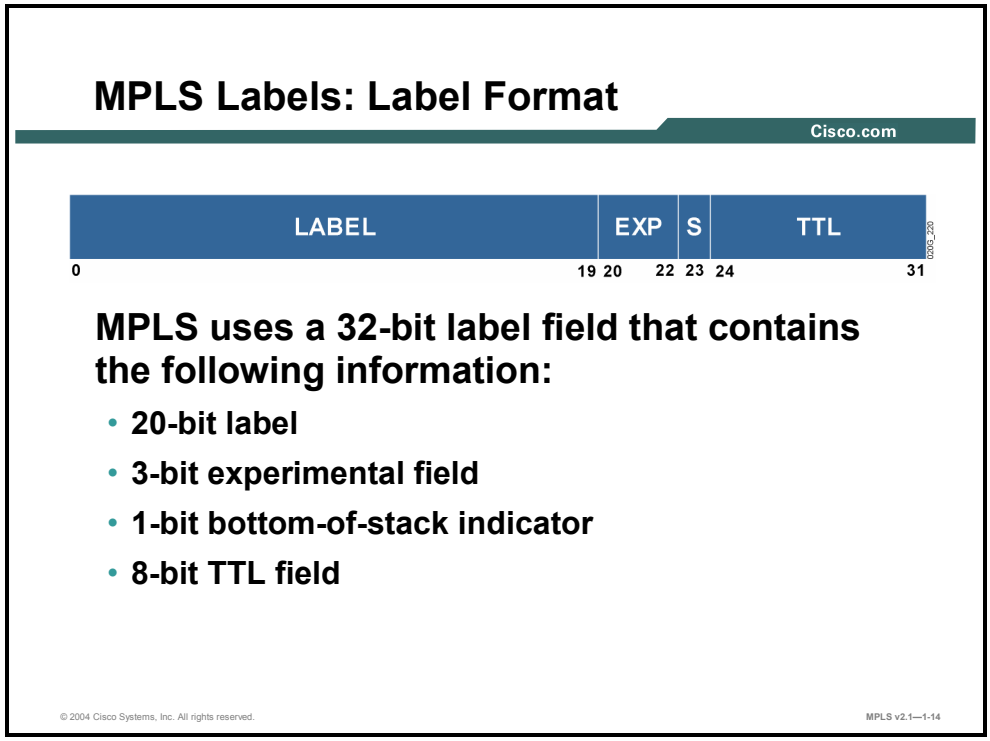

This table describes the fields contained in the 32-bit label.

32-Bit Label Fields

Field	Description
`20-bit label`	The actual label. Values 0 to 15 are reserved.
`3-bit experimental field`	Used to define a class of service (CoS) (IP precedence).
`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`	Has the same purpose as the TTL (time-to-live) field in the IP header.

Example: MPLS Labels—Frame-Mode MPLS

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

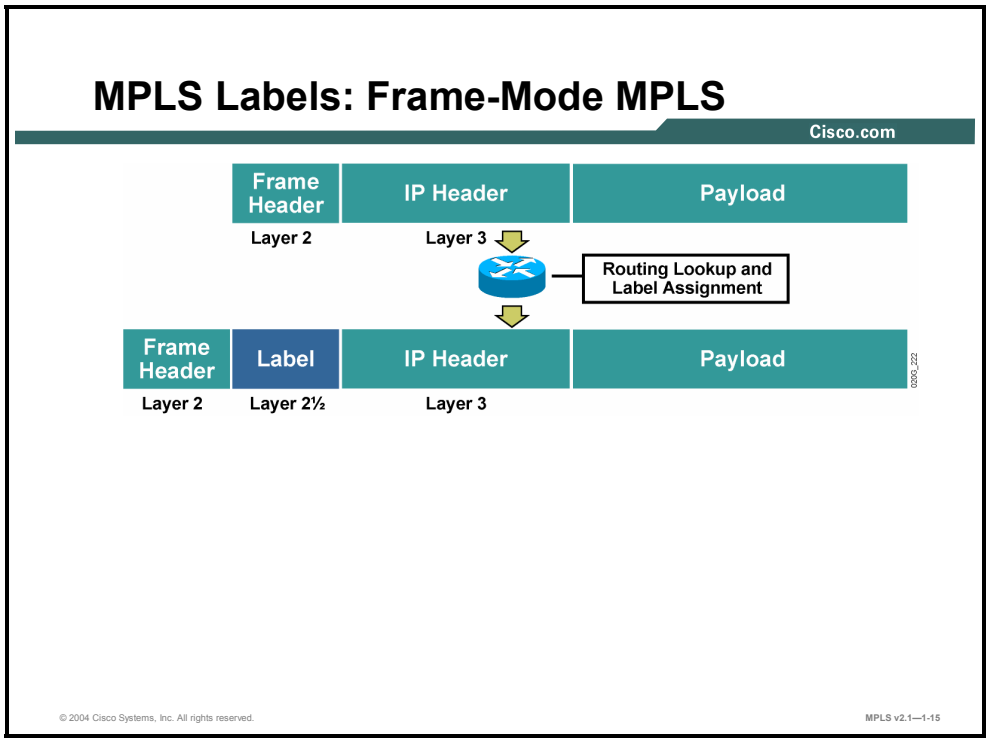

The router then does the following tasks:

- 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. The router then changes the Layer 2 Ethertype value to indicate that this is a labeled packet.
- The router sends the labeled packet.

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

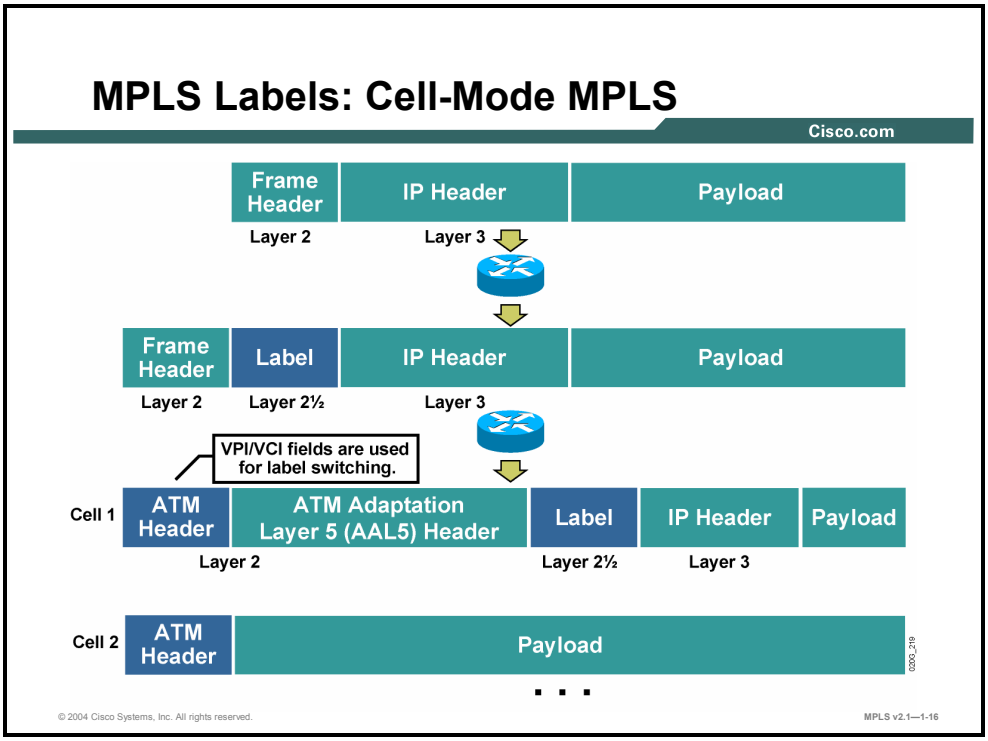

Cell-mode MPLS uses the ATM header VPI/VCI field for forwarding decisions. The 32-bit label is preserved in the frame but is not used in the ATM network. The original label is present only in the first cell of a packet.

What Are the Label Switch Router Functions?

This topic describes the function of the different types of LSRs.

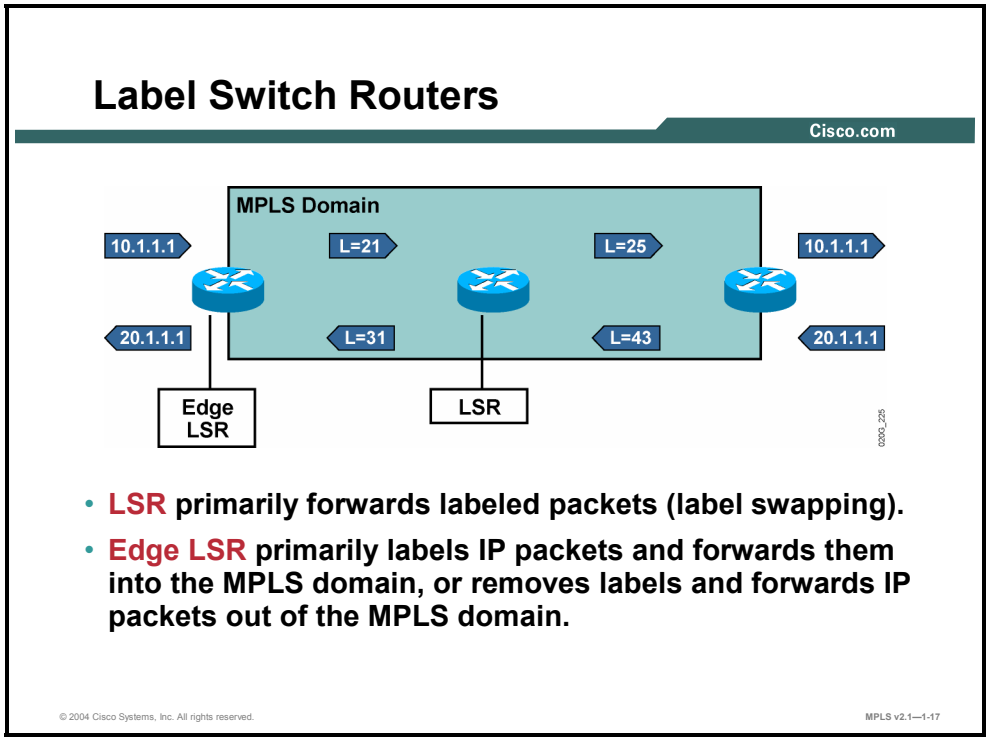

In preparation for a detailed description of MPLS, here is some of the terminology used in this course:

- **LSR:** A device that forwards packets primarily based on labels
- **Edge LSR:** A device that primarily labels packets or removes labels

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—autonomous system (AS). These routers also forward packets based on IP destination addresses and label them if the outgoing interface is enabled for MPLS.

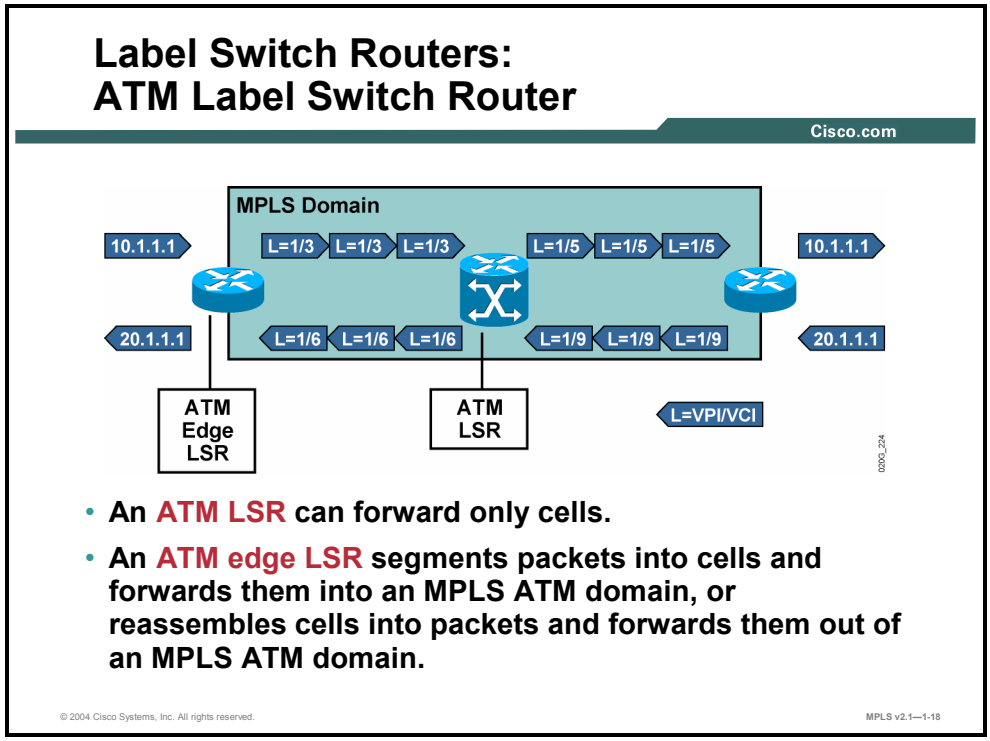

LSRs that perform cell-mode MPLS are divided into the following categories:

- ATM LSRs, if they are ATM switches. All interfaces are enabled for MPLS, and forwarding is done based *only* on labels.
- ATM edge LSRs, if they are routers connected to an MPLS-enabled ATM network.

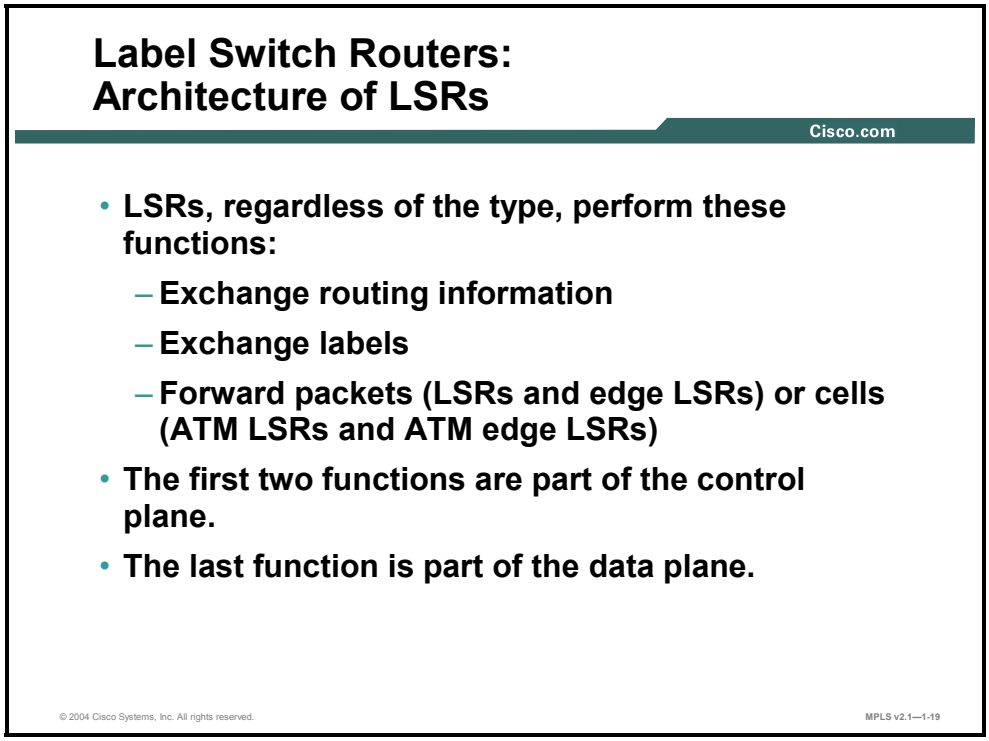

LSRs of all types must perform these functions:

- Exchange Layer 3 routing information; ATM LSRs must also exchange Layer 3 routing information (control plane).
- Exchange labels (control plane).
- Forward packets or cells (data plane).
- Frame-mode MPLS forwards packets based on the 32-bit label.
- Cell-mode MPLS forwards packets based on labels encoded into the VPI/VCI fields in the ATM header.

The data plane performs the following functions:

- Exchanges routing information regardless of the type of LSR
- Exchanges labels according to the type of MPLS (frame-mode or cell-mode)

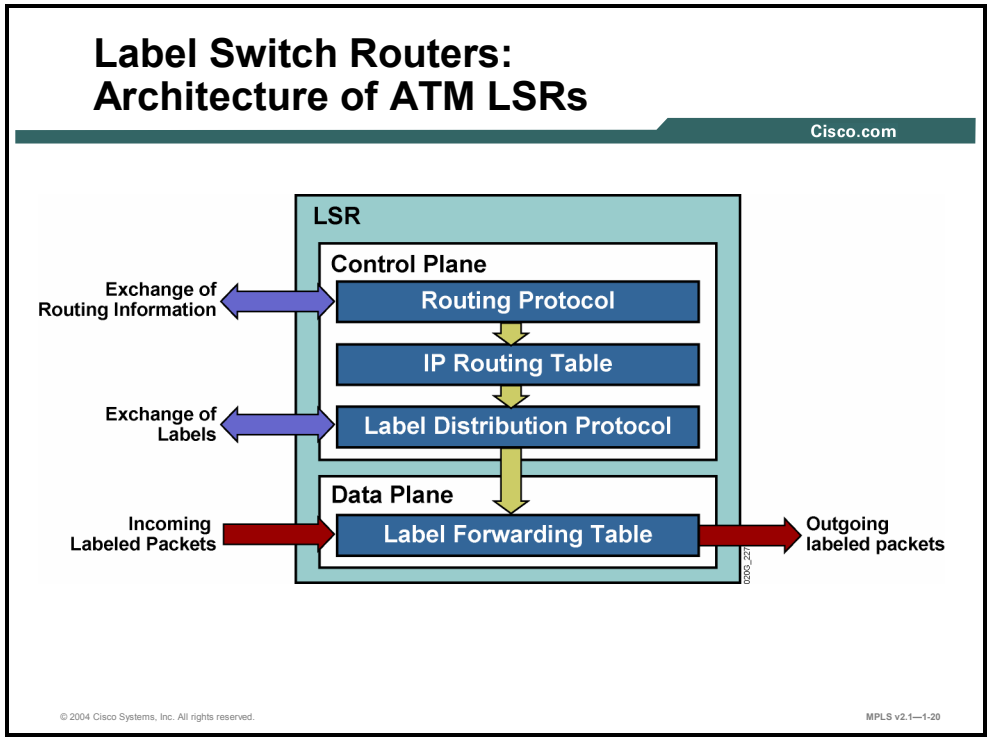

The primary function of an LSR is to forward labeled packets. Therefore, every LSR needs a Layer 3 routing protocol (for example, OSPF, EIGRP, IS-IS) and a label distribution protocol (for example, LDP, TDP).

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 either because they are ATM LSRs or because they do not have all of the routing information.

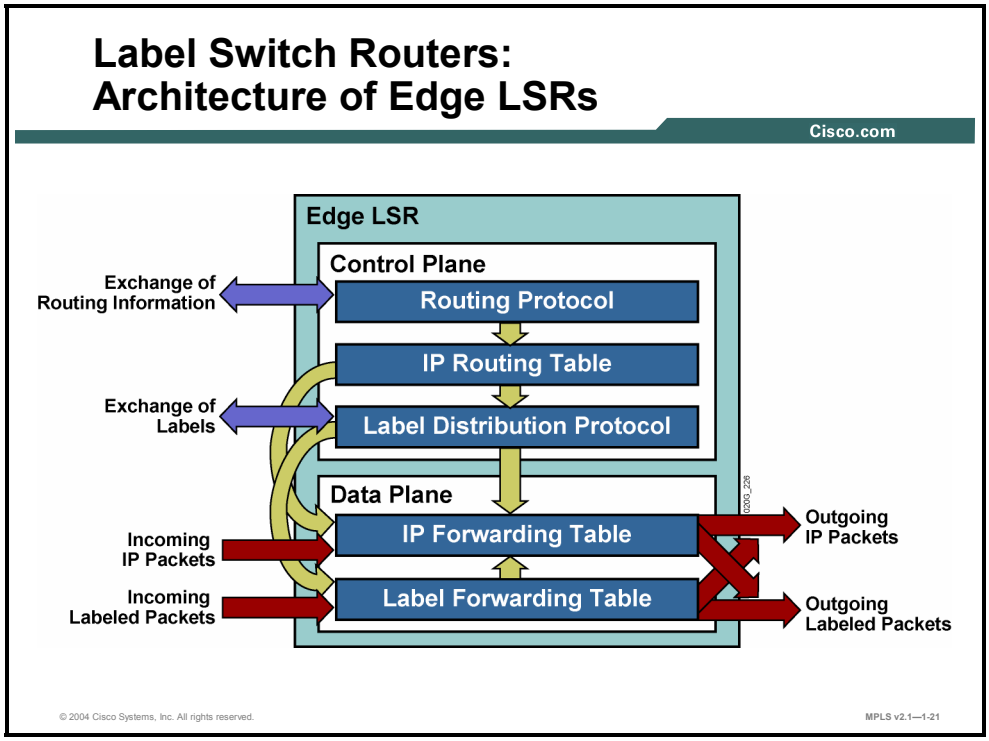

Edge LSRs also forward IP packets based on their IP destination addresses and optionally label them if a label exists.

The following combinations are possible:

- A received IP packet is forwarded based on the IP destination address and sent as an IP packet.
- A received IP packet is forwarded based on the IP destination address and sent as a labeled packet.
- A received labeled packet is forwarded based on the label; the label is changed and the packet is sent.

The following 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 Forwarding Information Base (FIB).
- A received IP packet is dropped if the destination is not found in the IP forwarding table (FIB table), even if there is an MPLS label-switched path toward the destination.

Summary

This topic summarizes the key points discussed in this lesson.

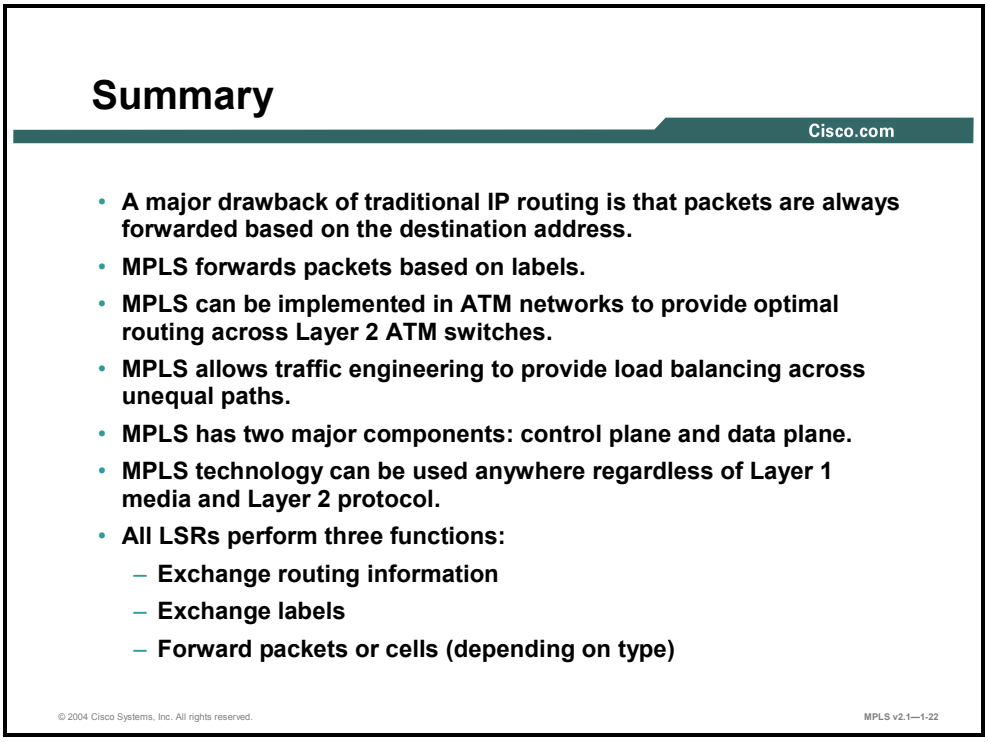

Lesson 2

Introducing MPLS Labels and Label Stack

Overview

This lesson explains the four fields that make up an MPLS label. This lesson also explains how label stacking is used and how labels are forwarded in frame-mode and cell-mode environments.

To fully understand MPLS, it is necessary to have a clear understanding of the format of an MPLS label and a 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

This lesson describes MPLS labels and an MPLS label stack, 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 where MPLS labels are inserted in an IP packet
- Describe the format and fields of an MPLS label
- Describe the features of an MPLS label stack
- Describe how MPLS forwards packets

Where Are MPLS Labels Inserted?

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

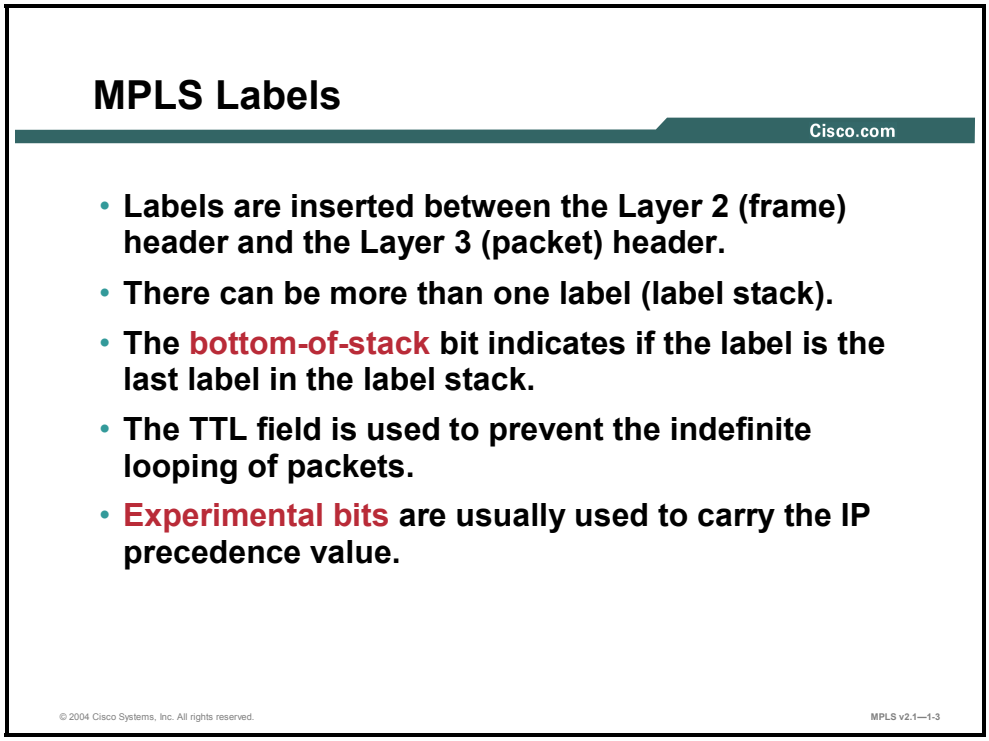

MPLS uses a 32-bit label that is inserted between the Layer 2 and Layer 3 headers. An MPLS label contains the following four fields:

- The actual label
- Experimental field
- Bottom-of-stack bit
- TTL field

These fields are explained in detail in this lesson.

What Is the MPLS Label Format?

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

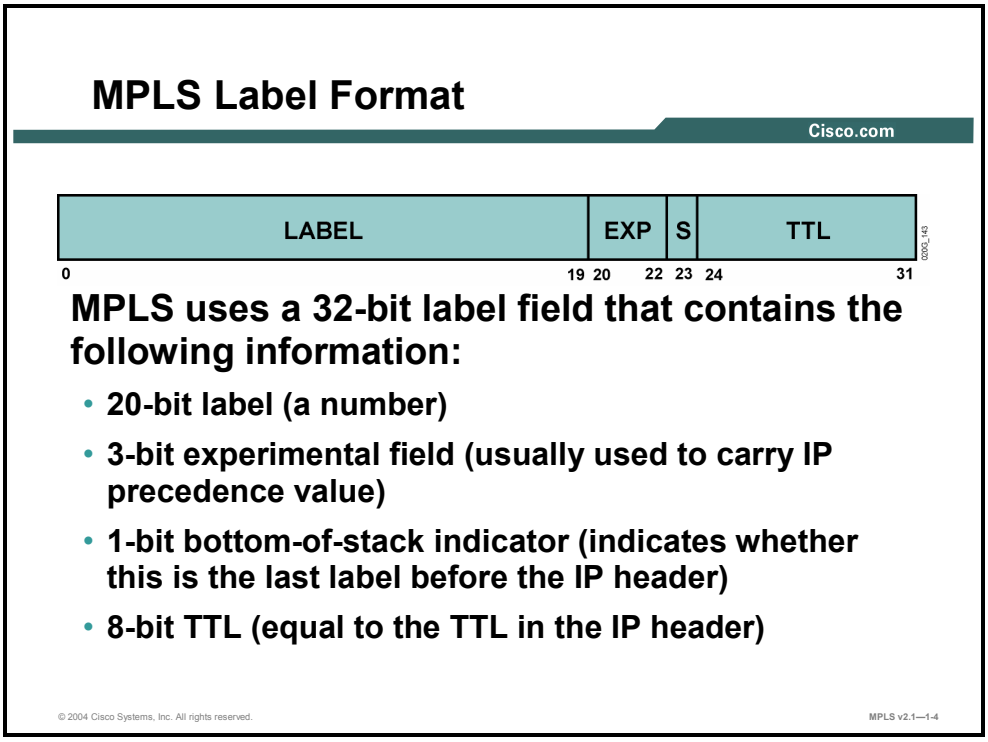

A label contains the fields listed in this table.

Label Fields

Field	Description
`20-bit label`	The actual label. Values 0 to 15 are reserved.
`3-bit experimental field`	Used to define a CoS (IP precedence).
`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`	Has the same purpose as the TTL (time-to-live) field in the IP header.

What Is an MPLS Label Stack?

This topic describes the features of an MPLS label stack.

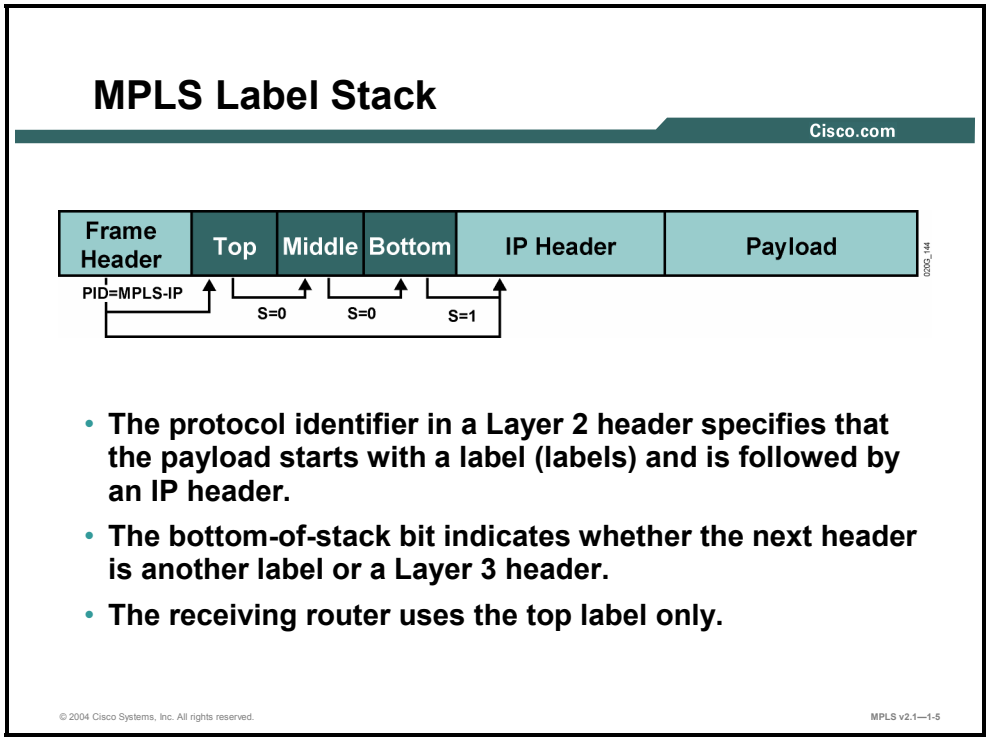

A label does not contain any information about the Layer 3 protocol being carried in a packet. A new protocol identifier is used for every MPLS-enabled Layer 3 protocol.

The following Ethertype values are used to identify Layer 3 protocols with most Layer 2 encapsulations:

- **Unlabeled IP unicast:** Process ID (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.

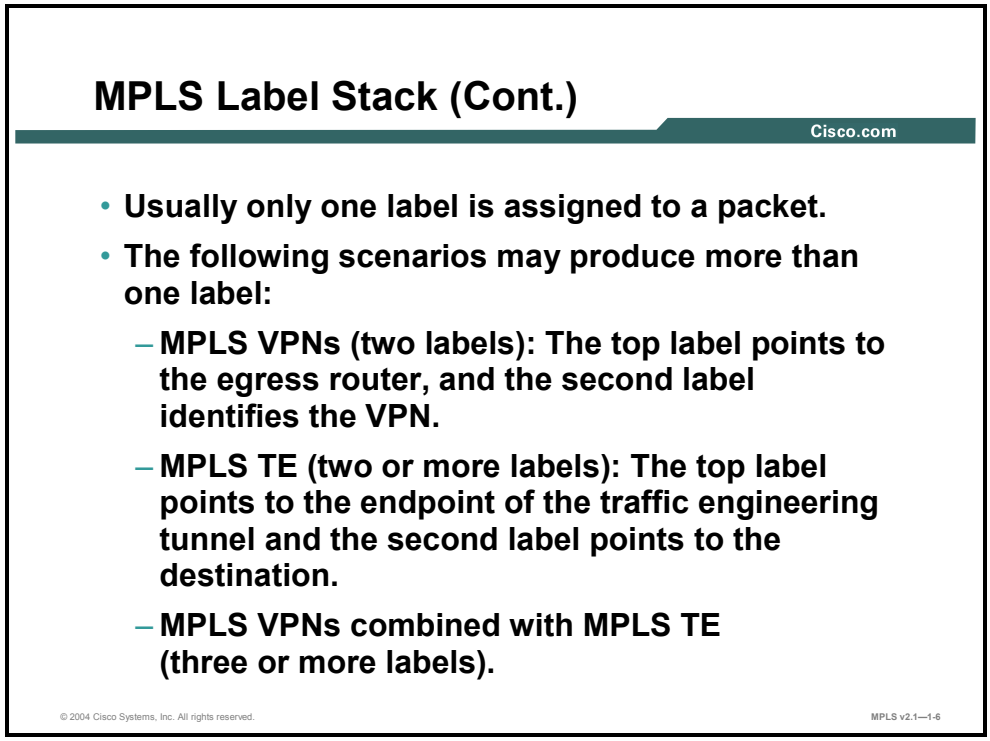

As previously noted, MPLS supports multiple labels in one packet. Simple MPLS uses just one label in each packet. The following applications may add labels to packets:

- **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 TDP or LDP.
- **MPLS TE:** MPLS TE uses 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 with some other features might result in three or more labels being inserted into one packet.

What Is MPLS Forwarding?

This topic describes how MPLS forwards packets.

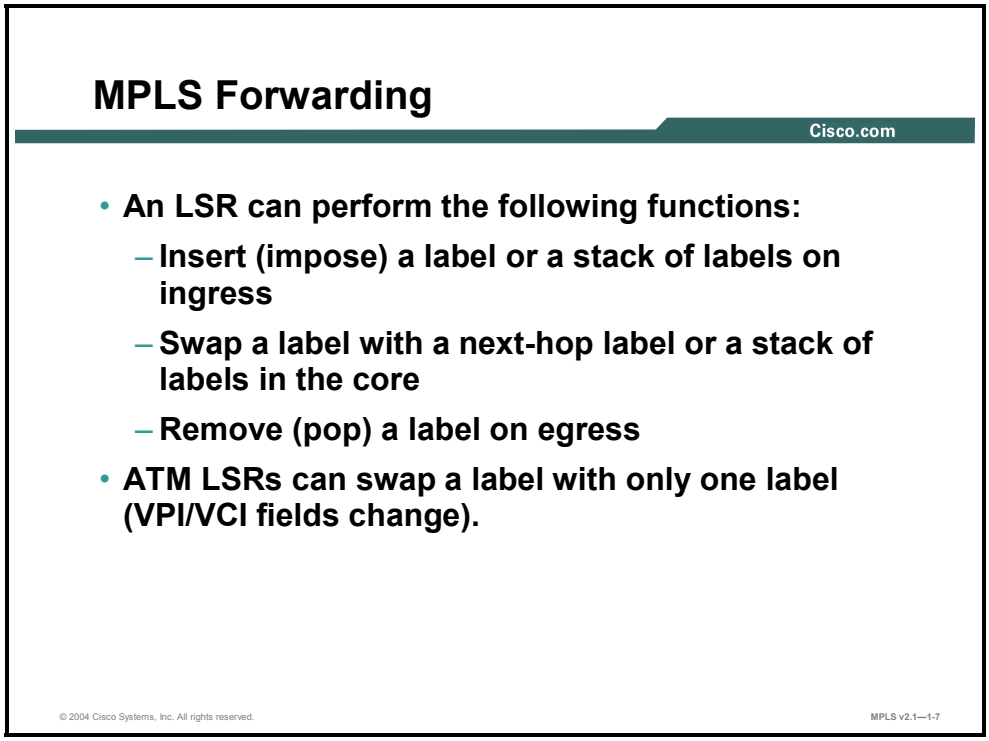

An IP packet going through an MPLS domain experiences the following:

- A label or a stack of labels is inserted (imposed) on an edge LSR.
- The top label is swapped with a next-hop label or a stack of labels on an LSR.
- The top label is removed on the LSP tunnel endpoint (usually one hop before the egress edge LSR or on the egress edge LSR itself).

ATM LSRs support the swapping of only one label (normal ATM operation).

Example: MPLS Network—Frame-Mode MPLS

This figure shows an MPLS network using frame-mode MPLS.

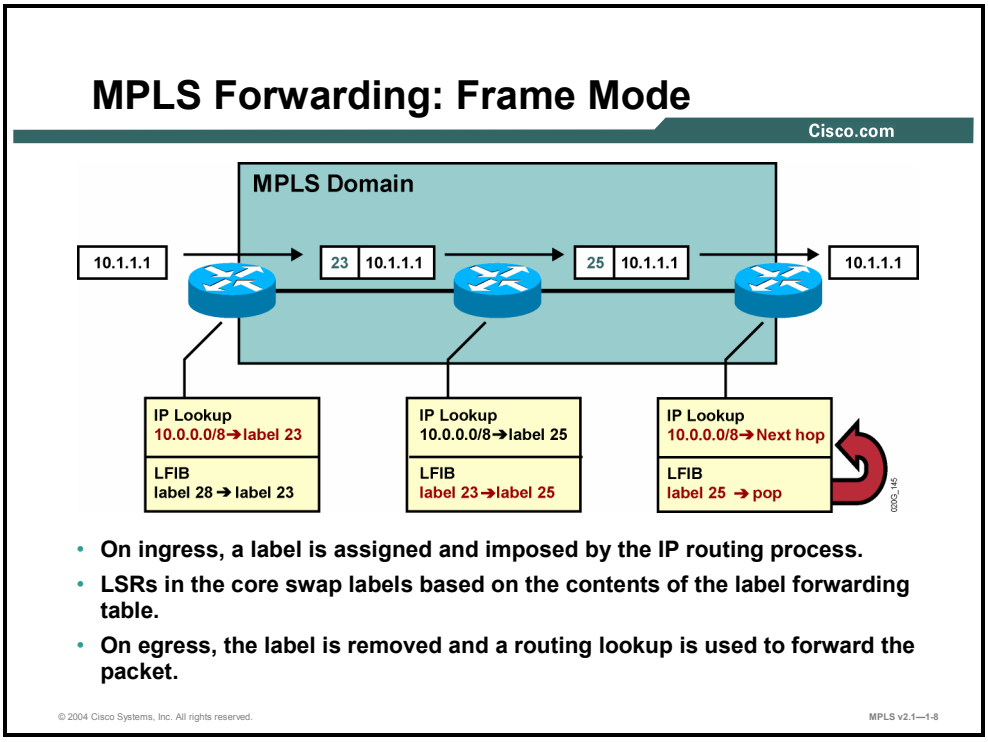

All LSRs are capable of forwarding IP packets or labeled packets. The ingress edge LSR performs a routing lookup and assigns a label.

The middle router simply swaps the label.

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

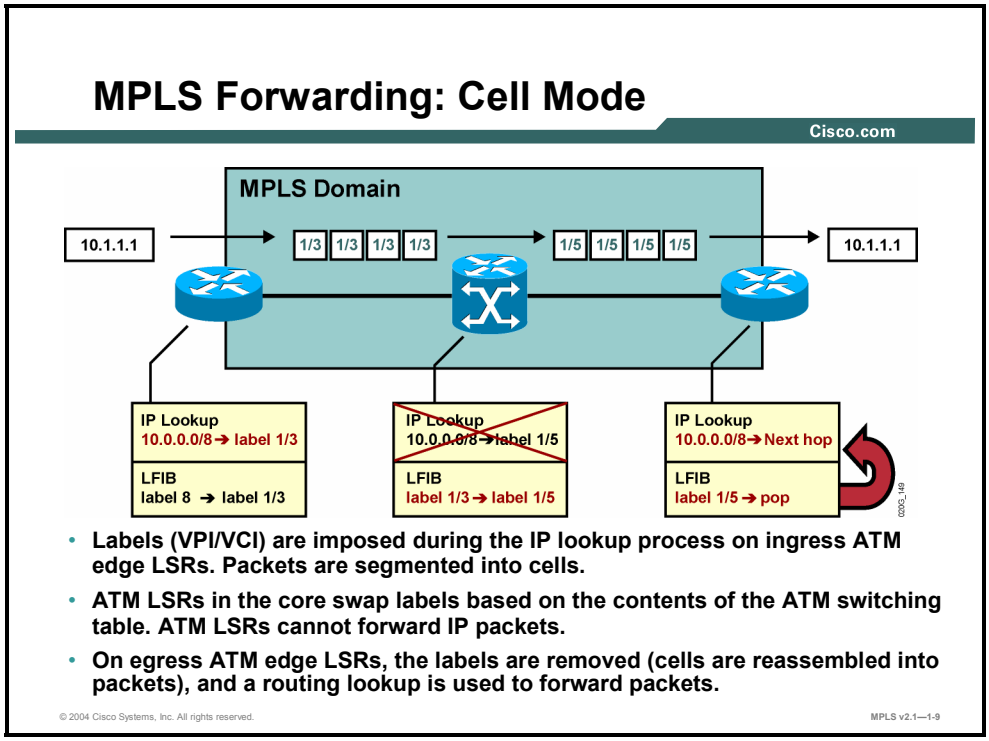

Cell-mode MPLS is similar to frame-mode MPLS. The difference is that ATM LSRs (ATM switches) cannot forward IP packets.

Summary

This topic summarizes the key points discussed in this lesson.

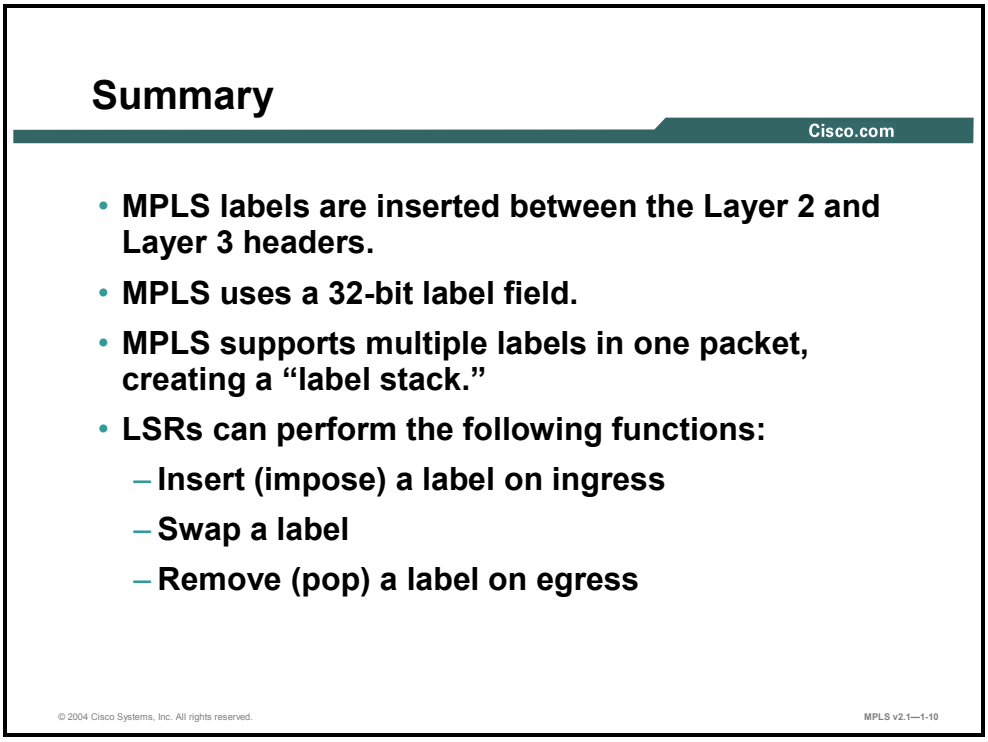

Lesson 3

Identifying MPLS Applications

Overview

This lesson looks at some of the different types of applications with which you can use 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

This lesson describes the different MPLS applications where 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 unicast IP routing
- Describe the features of multicast IP routing
- Describe MPLS use in TE environments
- Describe MPLS use in QoS environments
- Describe MPLS use in VPNs
- 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.

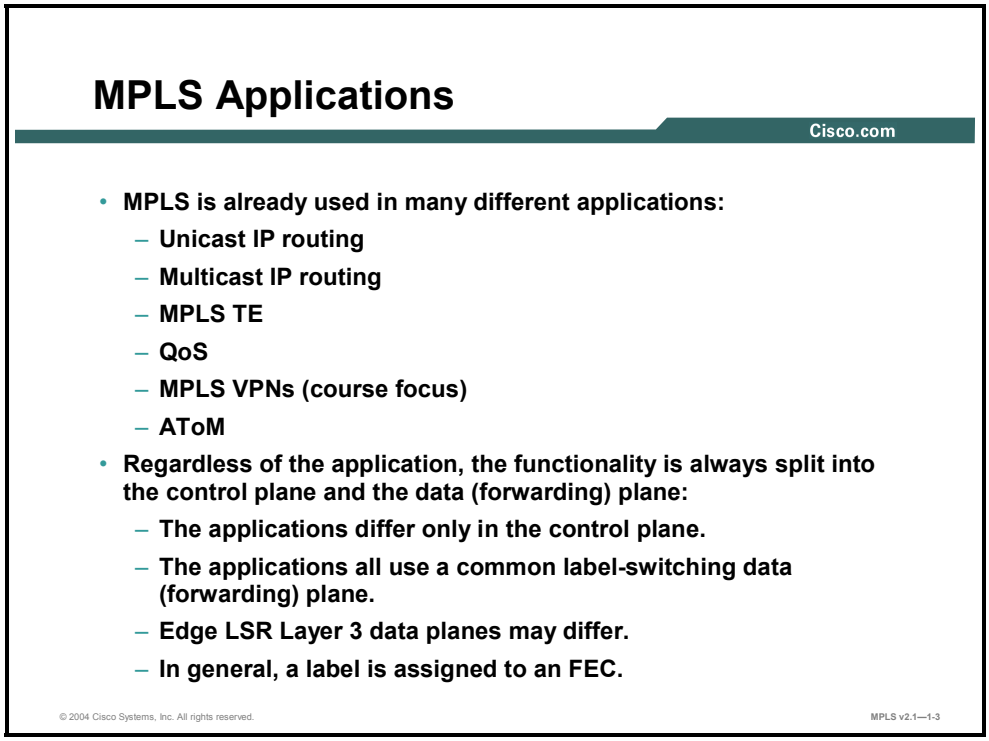

MPLS can be used in different applications, as described here:

- 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.
- Any Transport over MPLS (AToM) is a solution for transporting Layer 2 packets over an IP or MPLS backbone.

The data plane (forwarding plane) is the same regardless of the application. The control plane, however, needs appropriate mechanisms to exchange routing information and labels.

The term “forwarding equivalence class” (FEC) is used to describe the packets that are forwarded based upon a common characteristic (that is, destination address, QoS class, and so on).

What Is Unicast IP Routing?

This topic describes the features of unicast IP routing.

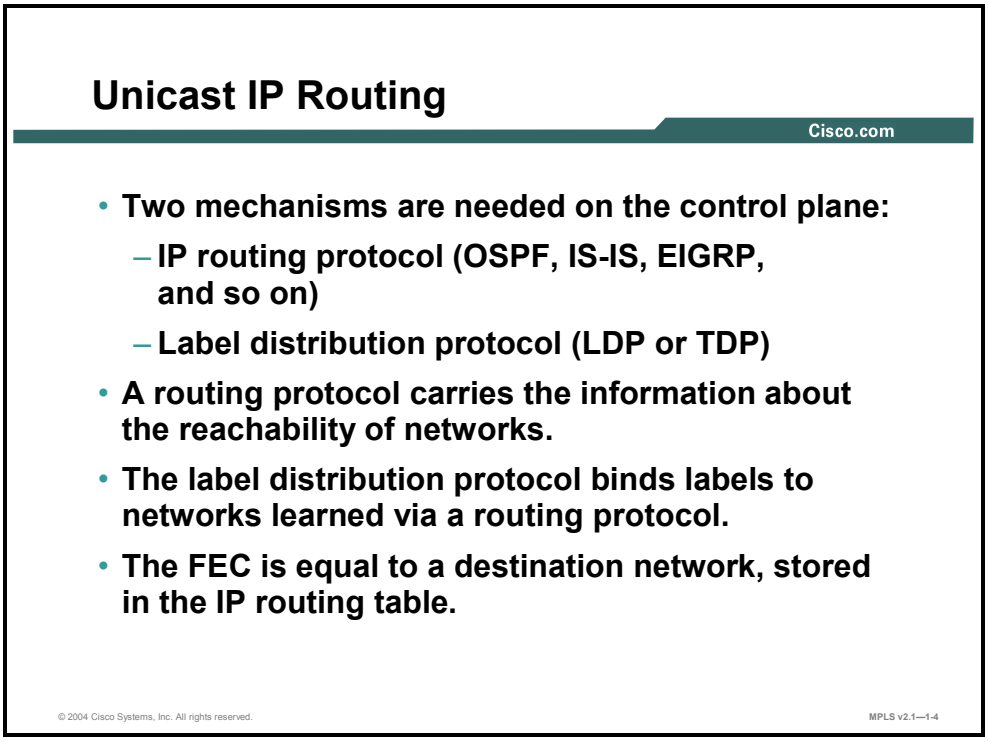

A unicast IP routing setup usually requires the following two components:

- IP routing protocol (for example, OSPF, EIGRP, IS-IS)
- Label distribution protocol (TDP or LDP)

These two components are enough to create a full mesh of LSP tunnels.

A label is assigned to every destination network found in the IP forwarding table. That is why an FEC corresponds to an IP destination network.

What Is Multicast IP Routing?

This topic describes the features of multicast IP routing.

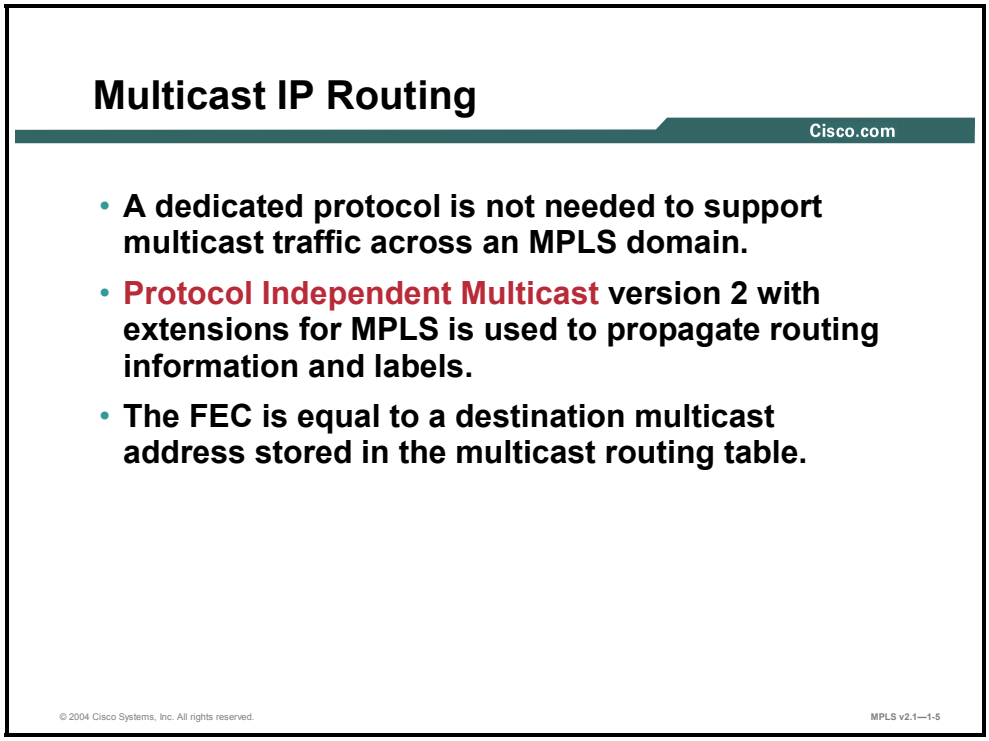

Multicast IP routing can also use MPLS. Cisco Protocol Independent Multicast (PIM) Version 2 with extensions for MPLS is used to propagate routing information and labels.

The FEC is equal to a destination multicast address.

Using MPLS Traffic Engineering

This topic describes MPLS use in TE environments.

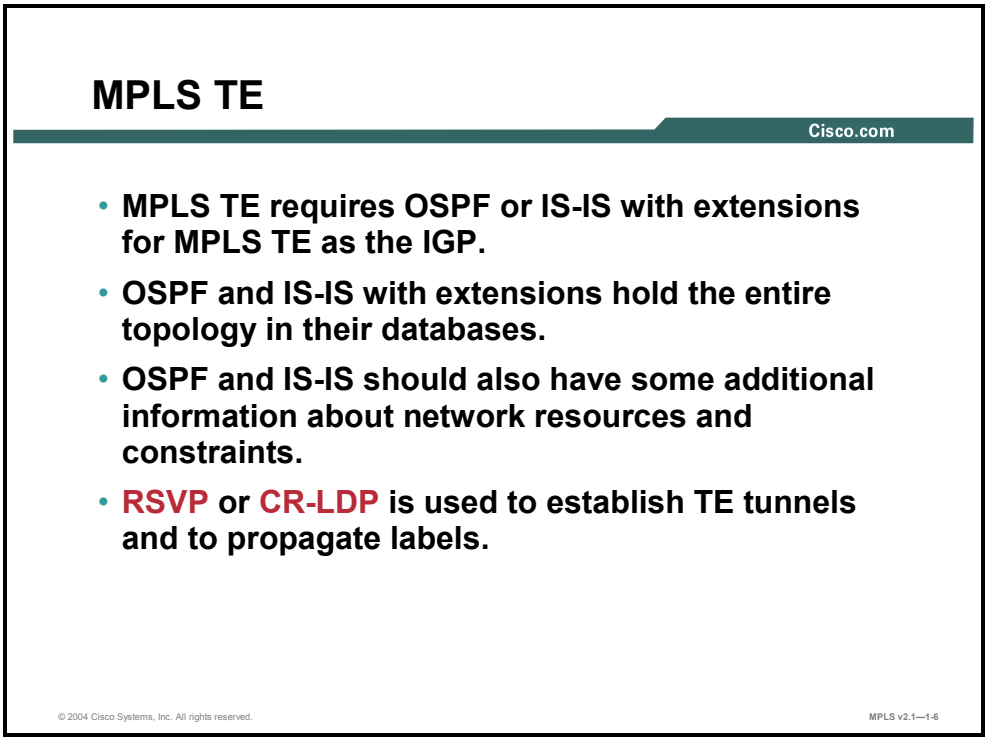

MPLS TE has the following special requirements:

- Every LSR must see the entire topology of the network (only OSPF and 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.
- Either RSVP or Constraint Route-LDP (CR-LDP) 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 Quality of Service?

This topic describes MPLS use in QoS environments.

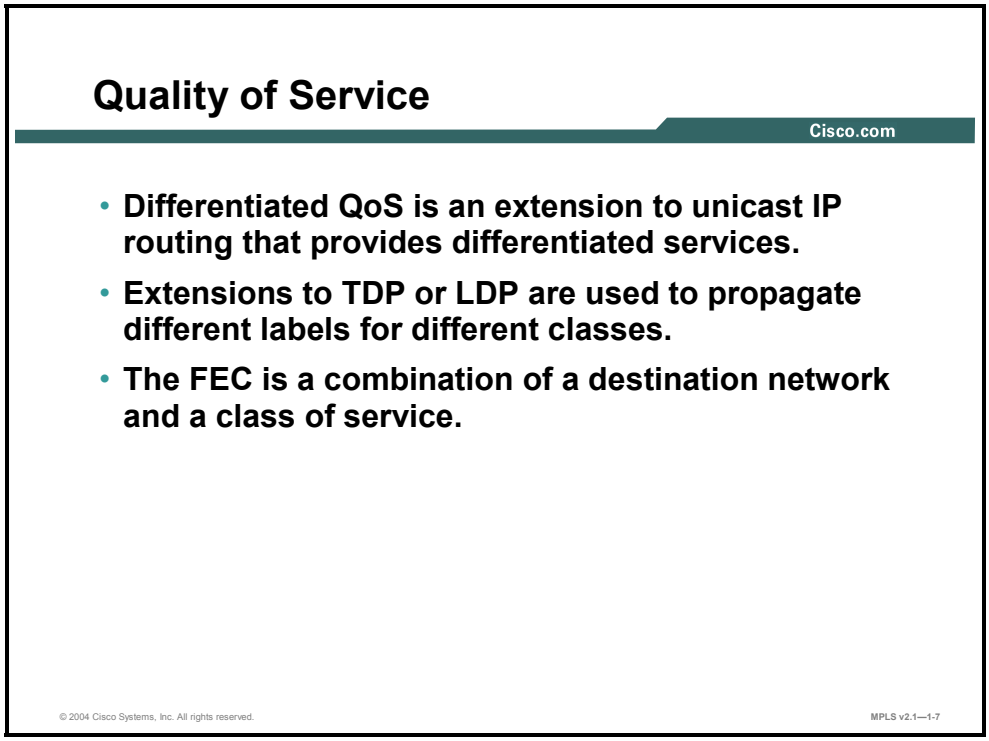

Differentiated QoS is achieved by using MPLS experimental bits or by creating separate LSP tunnels for different classes. Extensions to TDP or LDP are used to create multiple LSP tunnels for the same destination (one for each class).

The FEC is equal to a combination of a destination network and a CoS.

What Are Virtual Private Networks?

This topic describes MPLS use in VPNs.

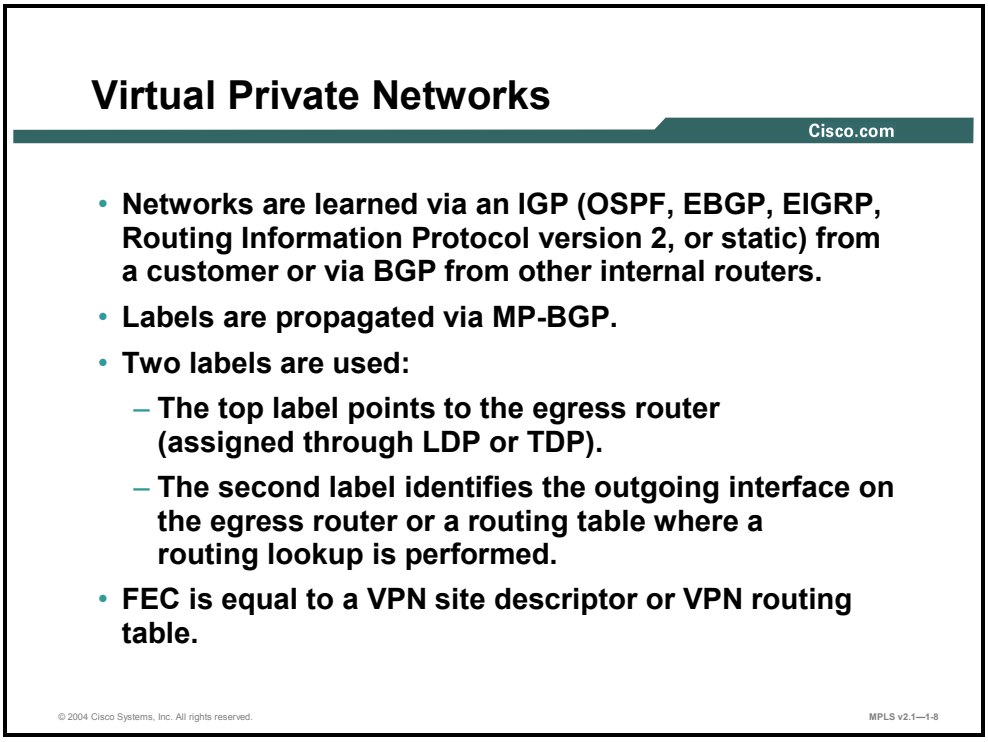

MPLS VPNs use an additional label to determine the VPN and the corresponding VPN destination network. MP-BGP is used to propagate VPN routing information and labels across the MPLS domain. TDP or LDP is needed to link edge LSRs with a single LSP tunnel.

The FEC is equal to a VPN destination network.

What Are the Interactions Between MPLS Applications?

This topic identifies the interactions that occur between MPLS applications.

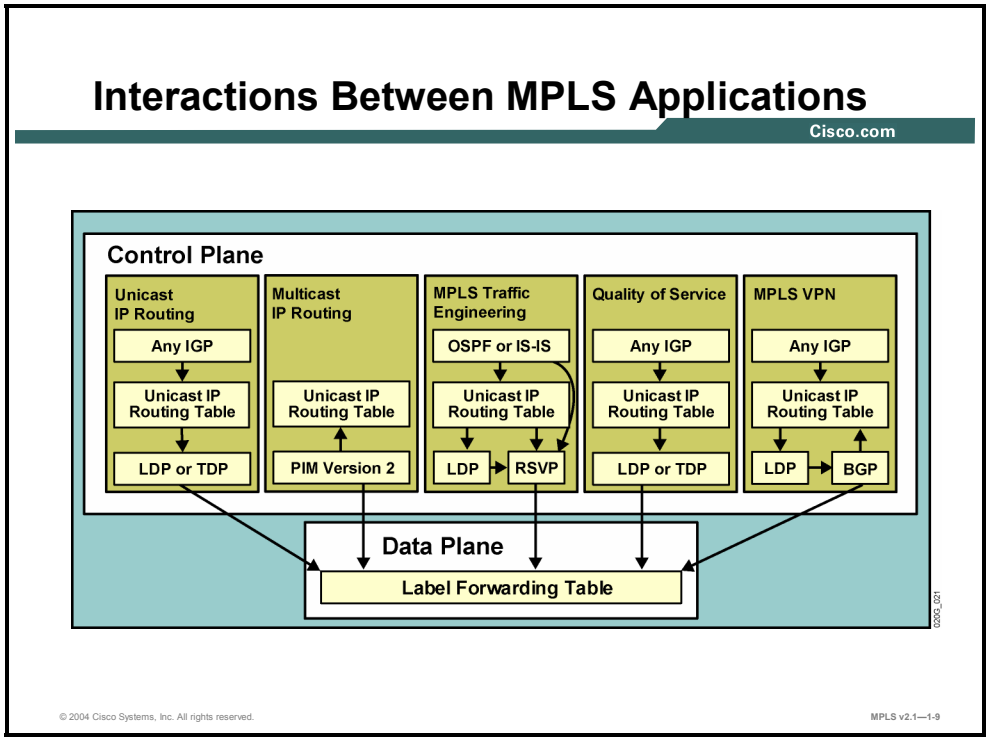

Each application may use a different routing protocol and a different label exchange protocol, but all applications use a single label-forwarding engine.

Example: Interactions Between MPLS Applications

The figure shows the complete architecture when all applications are used.

Summary

This topic summarizes the key points discussed in this lesson.

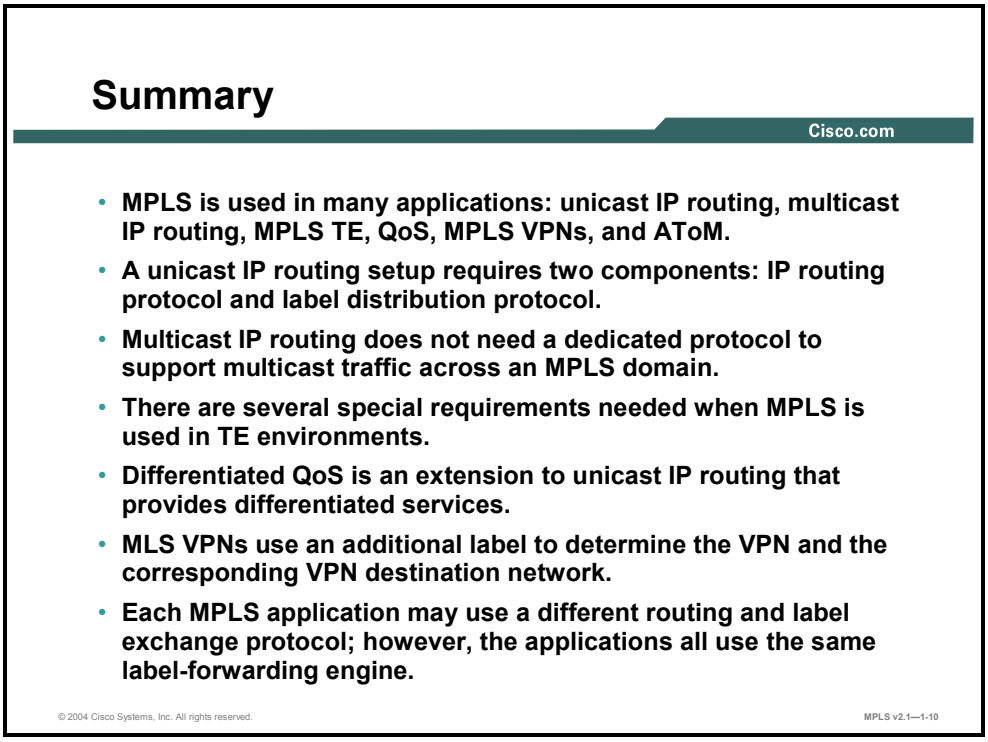

Module Summary

This topic summarizes the key points discussed in this module.

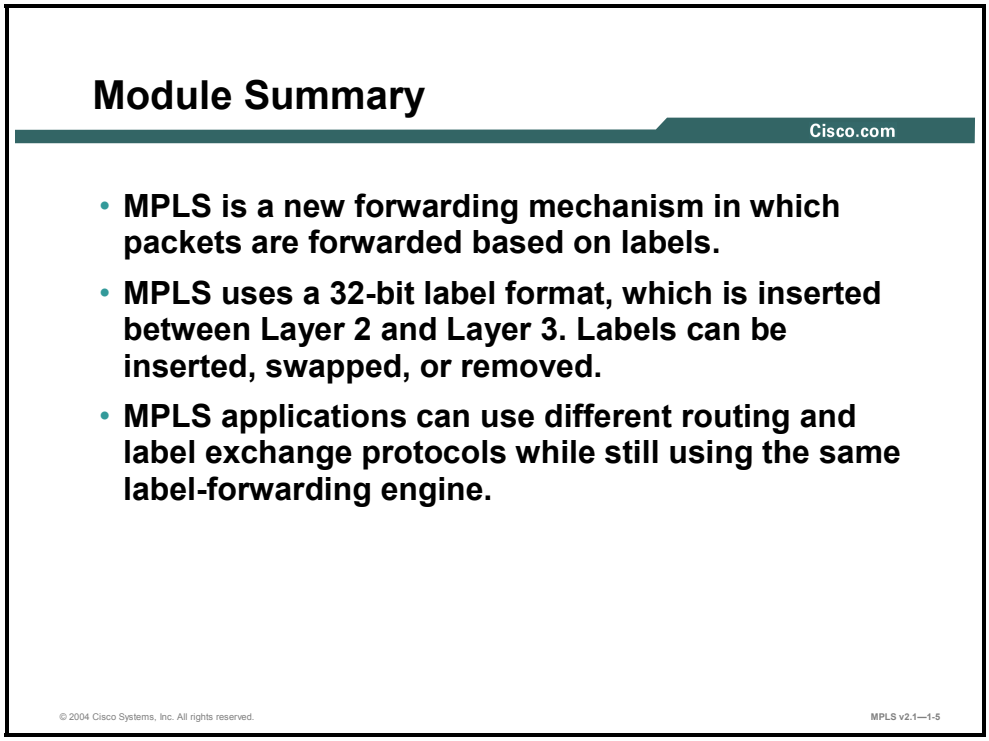

MPLS forwards packets based on labels. MPLS can be implemented in ATM networks to provide optimal routing across Layer 2 ATM switches. MPLS uses the concept of a “label stack” where multiple labels are supported in one packet. You can use MPLS in many applications. When many MPLS applications are being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

References

For additional information, refer to these resources:

- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3032, *MPLS Label Stack Encoding*

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 drawbacks 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 of the following 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) As a result of implementing MPLS in ATM networks, which of the following is true? (Source: Introducing Basic MPLS Concepts)

A) Layer 2 devices run a Layer 3 routing protocol.
B) Virtual circuits must be configured manually.
C) MPLS cannot run in an ATM network.
D) ATM switches needed to be made Layer 3- *and* Layer 4-aware.

Q4) In MPLS TE, which two of the following 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.

Q5) The label distribution protocol (either TDP or LDP) is the responsibility of the _____. (Source: Introducing Basic MPLS Concepts)

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

Q6) 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

Q7) Which two of the following statements are true? (Choose two.) (Source: Introducing Basic MPLS Concepts)

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

Q8) MPLS labels can correspond to which of the following? (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

Q9) Which one of the following terms 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

Q10) Which two of the following statements are true? (Choose two.) (Source: Introducing Basic MPLS Concepts)

A) MPLS labels are inserted between the Layer 2 header and the Layer 3 header.
B) MPLS labels are inserted after the Layer 3 header.
C) In ATM networks, MPLS uses the VPI/VCI fields as the label.
D) MPLS will not work in ATM networks.

Q11) Which two of the following statements are true? (Choose two.) (Source: Introducing MPLS Labels and Label Stack)

A) MPLS labels are 32 bits.
B) MPLS labels are 64 bits.
C) MPLS labels are inserted before the Layer 2 header.
D) MPLS labels are inserted after the Layer 2 header.

Q12) How long is the actual MPLS label contained in the MPLS label field? (Source: Introducing MPLS Labels and Label Stack)

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

Q13) Which two of the following statements are true? (Choose two.) (Source: Introducing MPLS Labels and Label Stack)

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.

Q14) Which two of the following are normal functions of an LSR? (Choose two.) (Source: Introducing MPLS Labels and Label Stack)

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

Q15) Cisco routers automatically assign the IP precedence value to which field in the MPLS label? (Source: Introducing MPLS Labels and Label Stack)

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.

Q16) Which of the following is NOT a valid Ethertype used to identify Layer 3 protocols with most Layer 2 encapsulations? (Source: Introducing MPLS Labels and Label Stack)

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

Q17) The label distribution protocol is found on which plane? (Source: Identifying MPLS Applications)

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

Q18) Which two of the following 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.

Q19) When MPLS is used for QoS, which of the following statements 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.

Q20) In MPLS VPN networks, which one of the following statements 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-BPG.
D) Two labels are used; the top label identifies the VPN, and the bottom label identifies the egress router.

Q21) Which two of the following 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.

Q22) 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) A

Q4) B, C

Q5) D

Q6) B

Q7) A, C

Q8) D

Q9) C

Q10) A, C

Q11) A, D

Q12) D

Q13) A, C

Q14) A, D

Q15) B

Q16) C

Q17) C

Q18) A, B

Q19) C

Q20) C

Q21) A, D

Q22) 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 the LIB, FIB, and LFIB tables are populated with label information
- Describe how convergence occurs in a frame-mode MPLS network
- Describe typical label distribution over LC-ATM interfaces and VC merge
- Describe MPLS label allocation, distribution, and retention modes
- Describe how LDP neighbors are discovered

Lesson 1

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 PHP and how the 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 label 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 label switch paths
- Describe the function of PHP
- Describe the impact that IP aggregation has on label-switched paths
- Describe how labels are allocated and distributed in a frame-mode MPLS network
- Describe how MPLS labels are distributed and allocated 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.

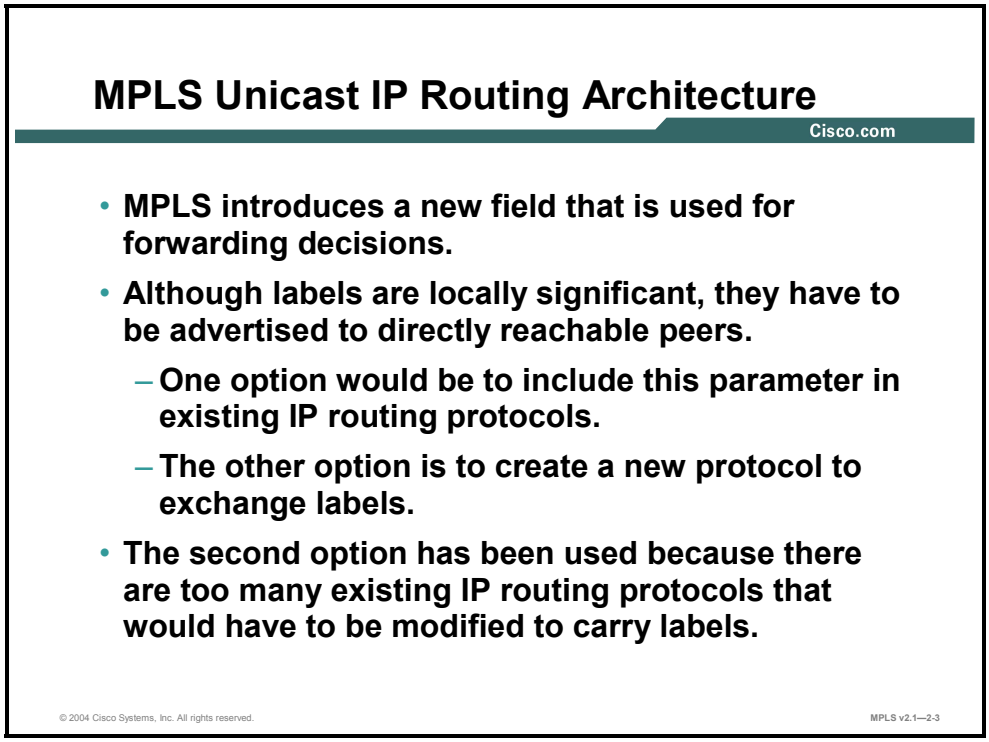

One application of MPLS is unicast IP routing. A label is assigned to destination IP networks and is later used to label packets sent toward those destinations.

Note	In MPLS terminology, a forwarding equivalence class (FEC) 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.

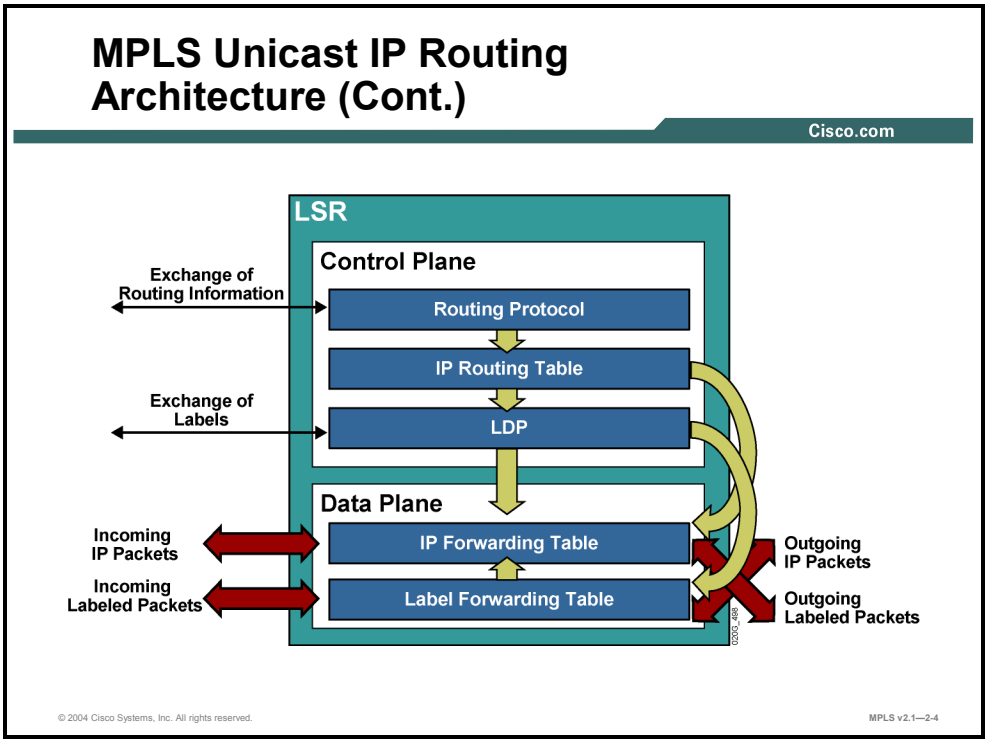

Example: Building Blocks for IP Forwarding

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) or the Cisco proprietary protocol Tag Distribution Protocol (TDP) in the control plane exchanges labels and stores them in the 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.

The following 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.

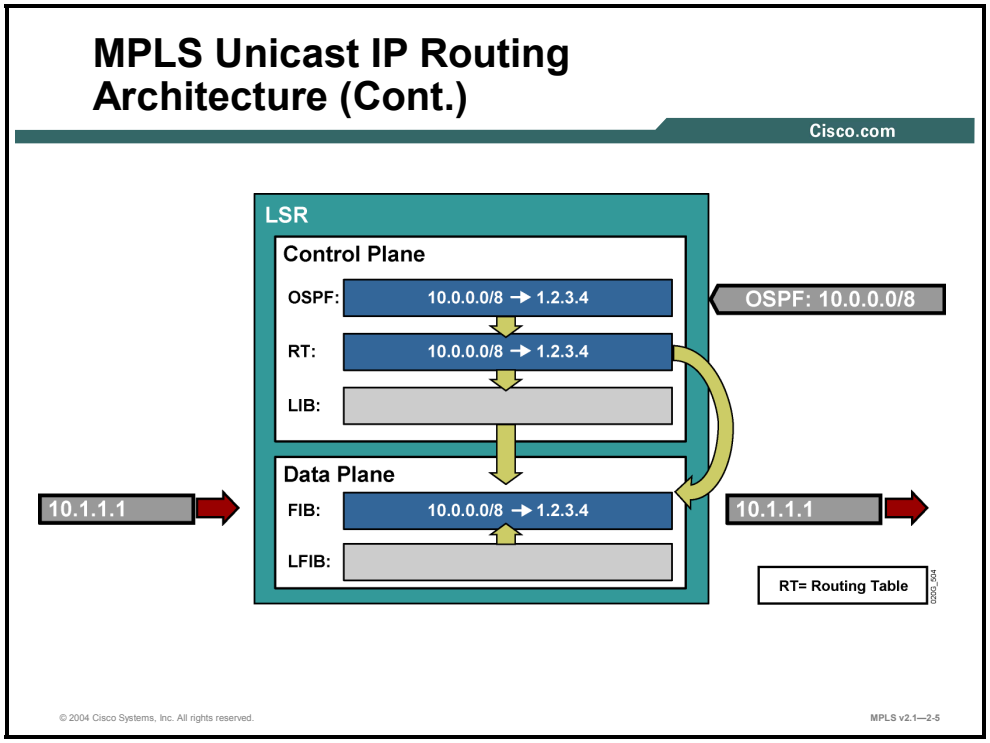

Example: Using the FIB Table to Forward Packets

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

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.

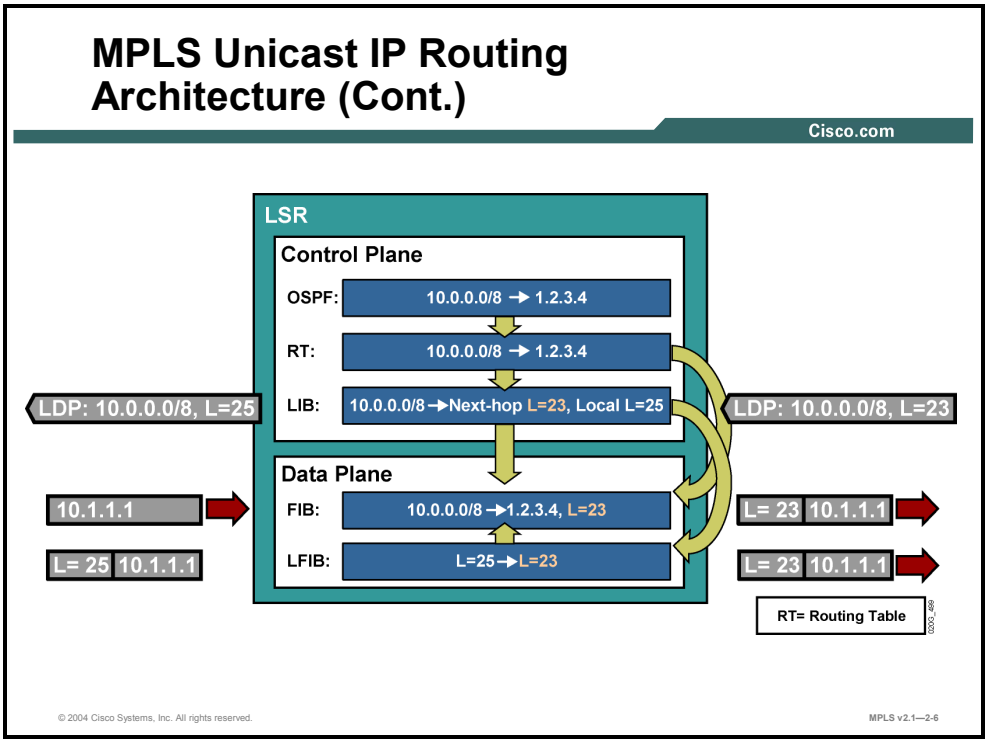

Example: Using LDP

The figure shows a router where OSPF is used to exchange IP routing information and LDP is used to exchange labels.

An incoming IP packet is forwarded by using the FIB table, where a next-hop label dictates that the outgoing packet should be labeled with label 23.

An incoming labeled packet is forwarded by using the LFIB table, where the incoming (locally significant) label 25 is swapped with the next-hop label 23.

What Are Label-Switched Paths?

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

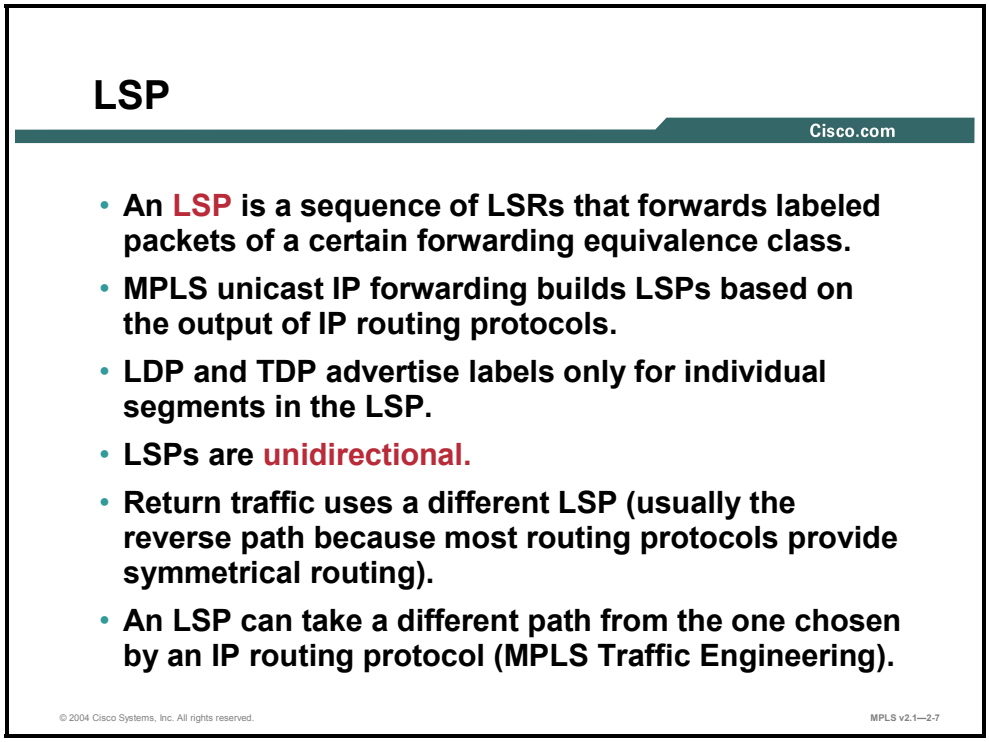

An LSP is a sequence of LSRs that forwards labeled packets for a particular FEC. Each LSR swaps the top label in a packet traversing the LSP. An LSP is similar to Frame Relay or ATM virtual circuits. In cell-mode MPLS, an LSP *is* a virtual circuit.

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 later in this course.

An Interior Gateway Protocol (IGP) is used to populate the routing tables in all routers in an MPLS domain. LDP or TDP 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 packets form the LSP in the opposite order.

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

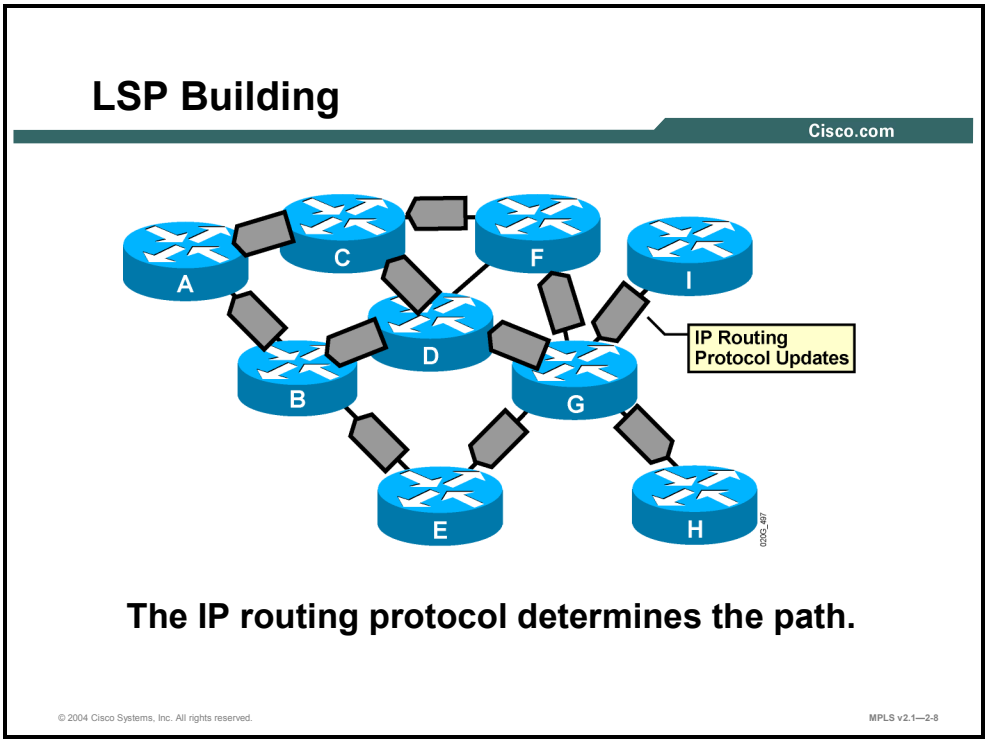

Example: IGP Propagates Routing Information

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 or TDP, which propagate labels for those networks and routers, adds this information to the FIB and LFIB tables.

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

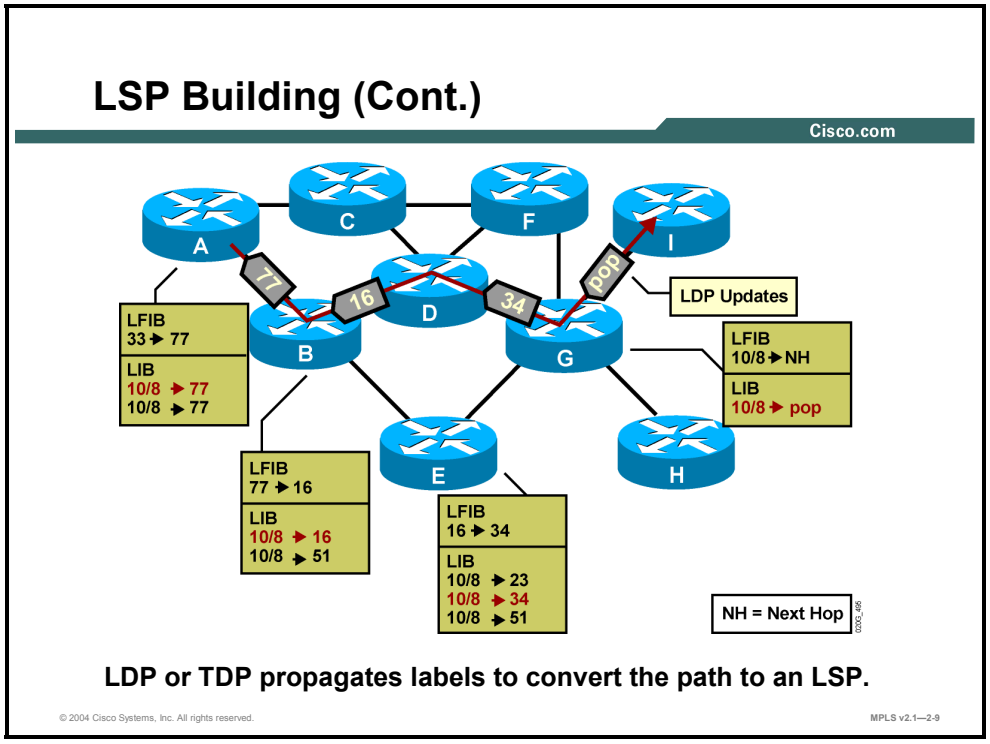

Example: LFIB and LIB Tables

The figure shows the contents of LFIB and LIB tables. Frame-mode MPLS uses a liberal retention mode, which is evident from the contents of the LIB 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.

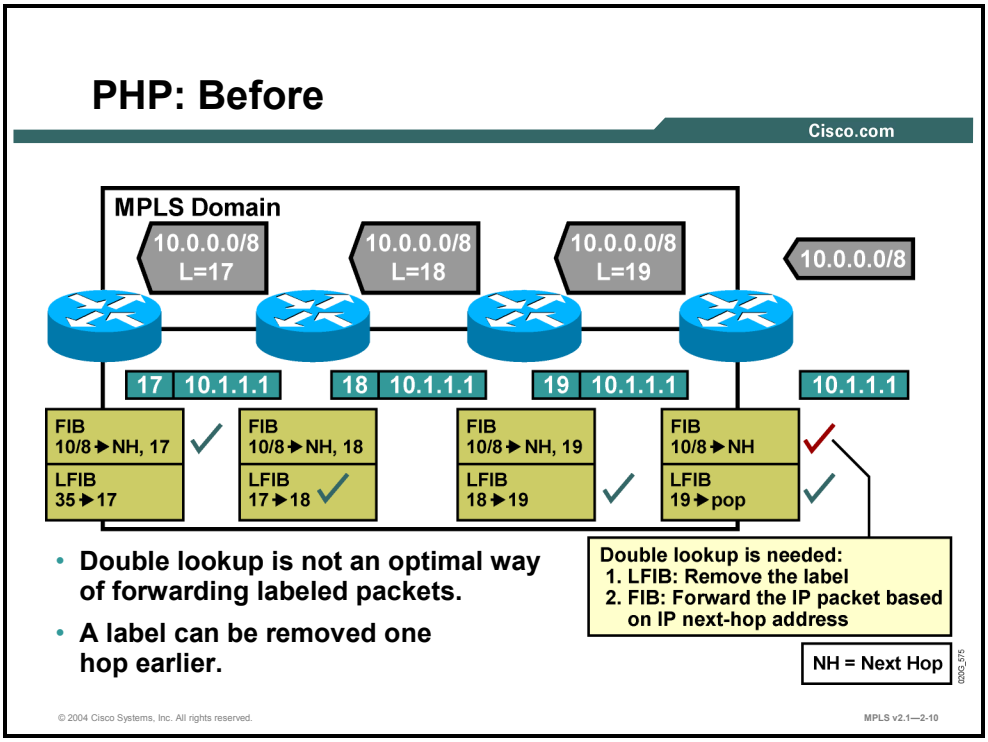

Example: PHP—Before

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.

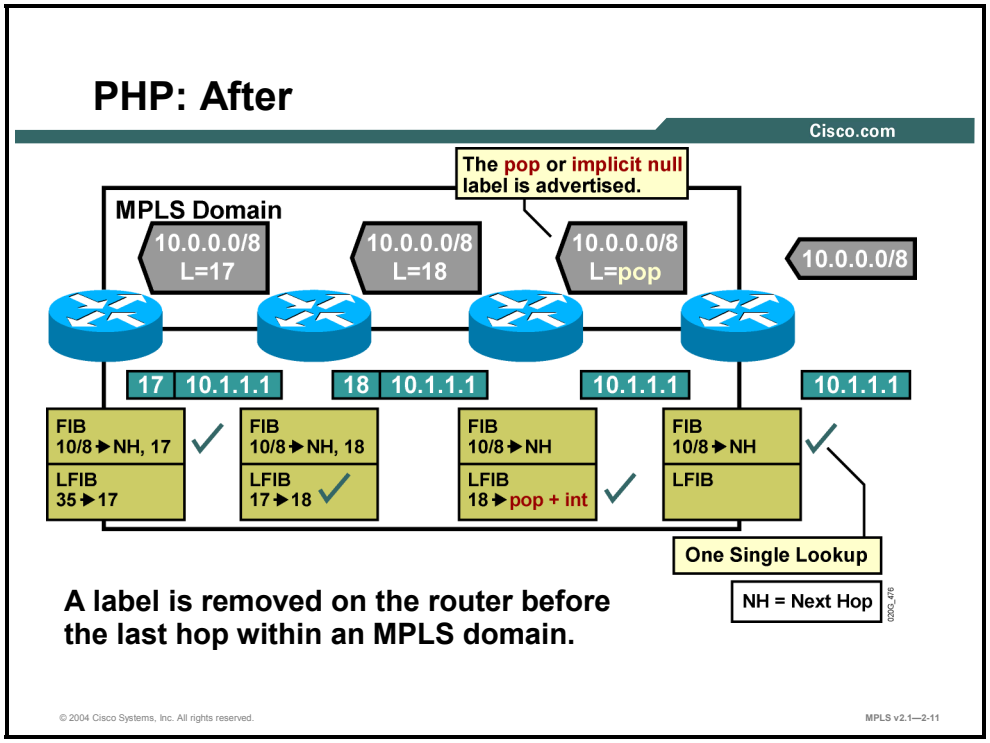

Example: PHP—After

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.

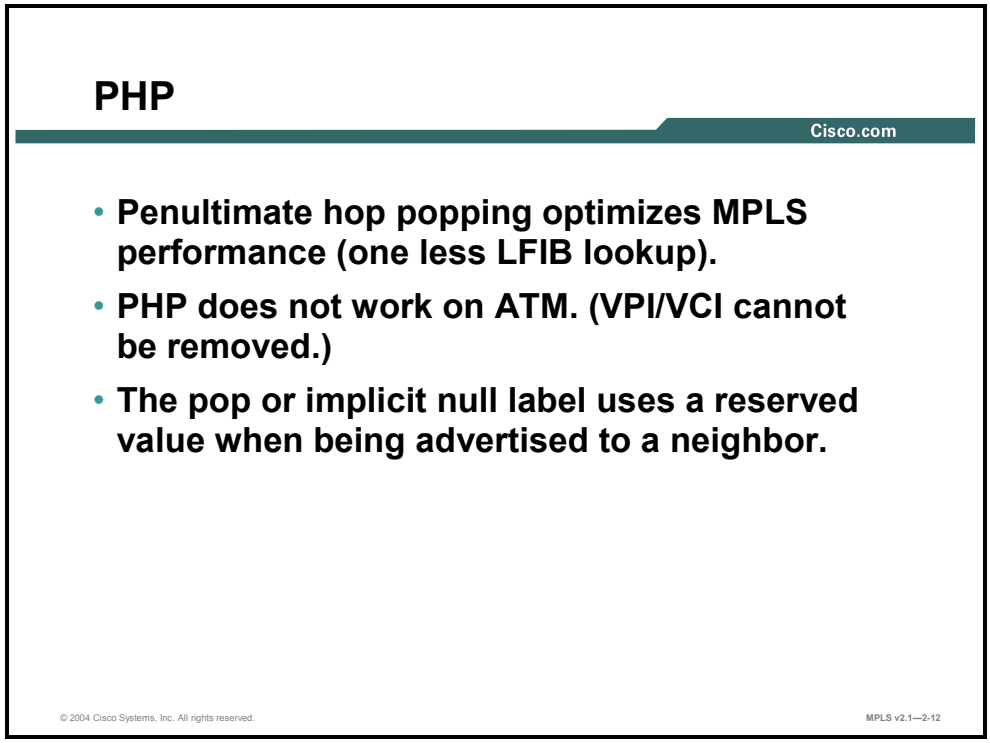

PHP optimizes MPLS performance by reducing the number of table lookups on the egress router.

PHP is not supported on ATM devices because a label is part of the ATM cell payload and cannot be removed by the ATM switching hardware.

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

What Is the Impact of IP Aggregation on Label-Switched Paths?

This topic describes the impact that IP aggregation has on label-switched paths.

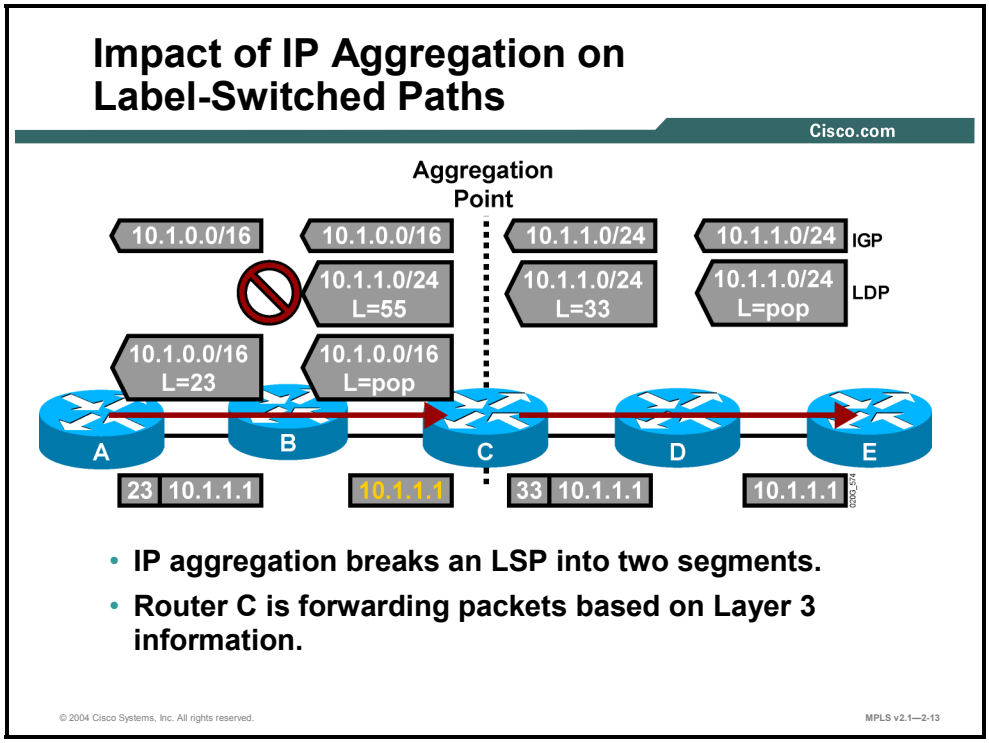

Example: MPLS IP Aggregation Problem

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 or TDP propagate 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.

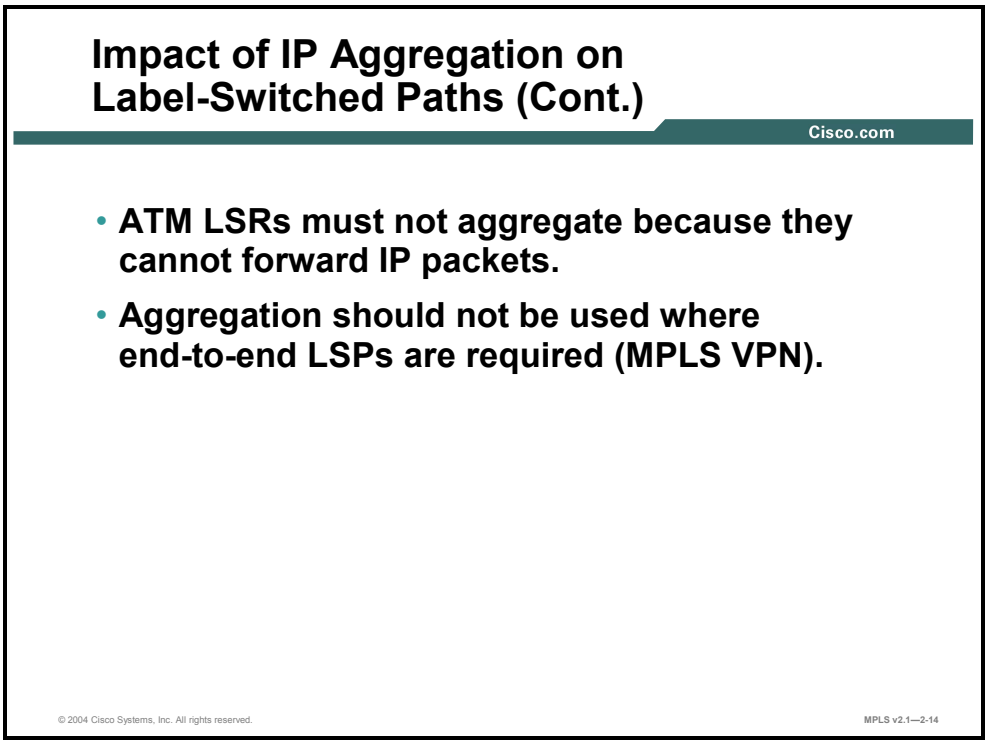

When cell-mode MPLS is used, ATM switches are IP-aware; ATM switches run an IP routing protocol, and LDP or TDP, and are generally seen as IP routers. In reality, however, ATM switches are capable of forwarding only cells, not IP packets.

Aggregation (or summarization) should not be used on ATM LSRs. This is because aggregation breaks LSPs in two, which means that ATM switches would have to perform Layer 3 lookups.

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 the following:

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

Allocating Labels in a Frame-Mode MPLS Network

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

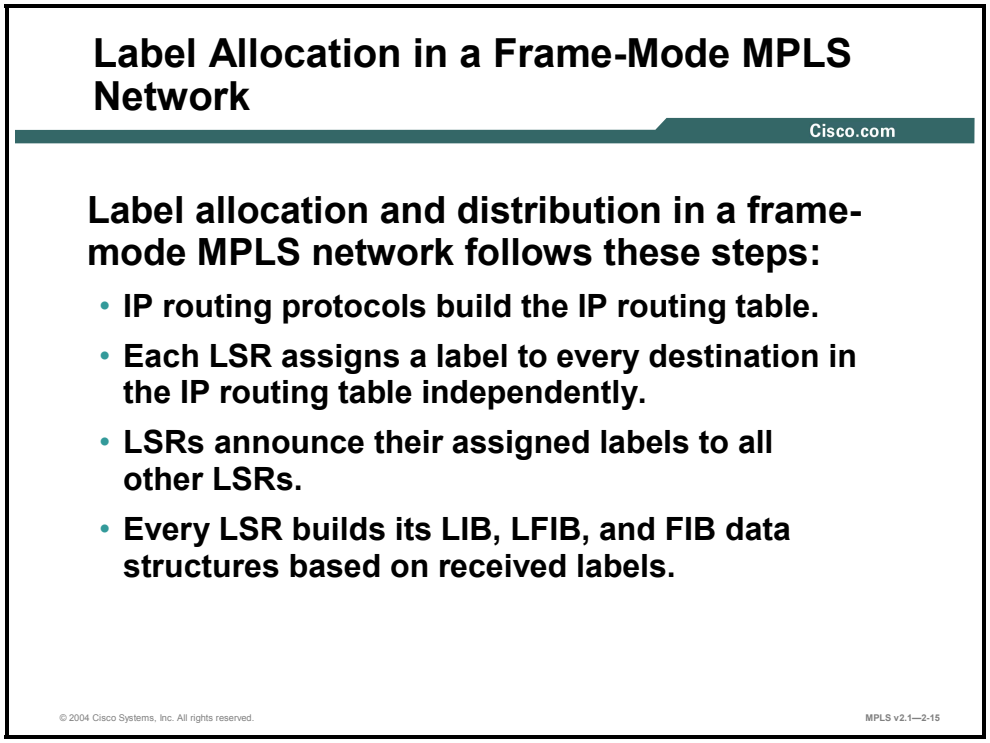

Unicast IP routing and MPLS functionality can be divided into the following steps:

- 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)

The following 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. 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. 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.

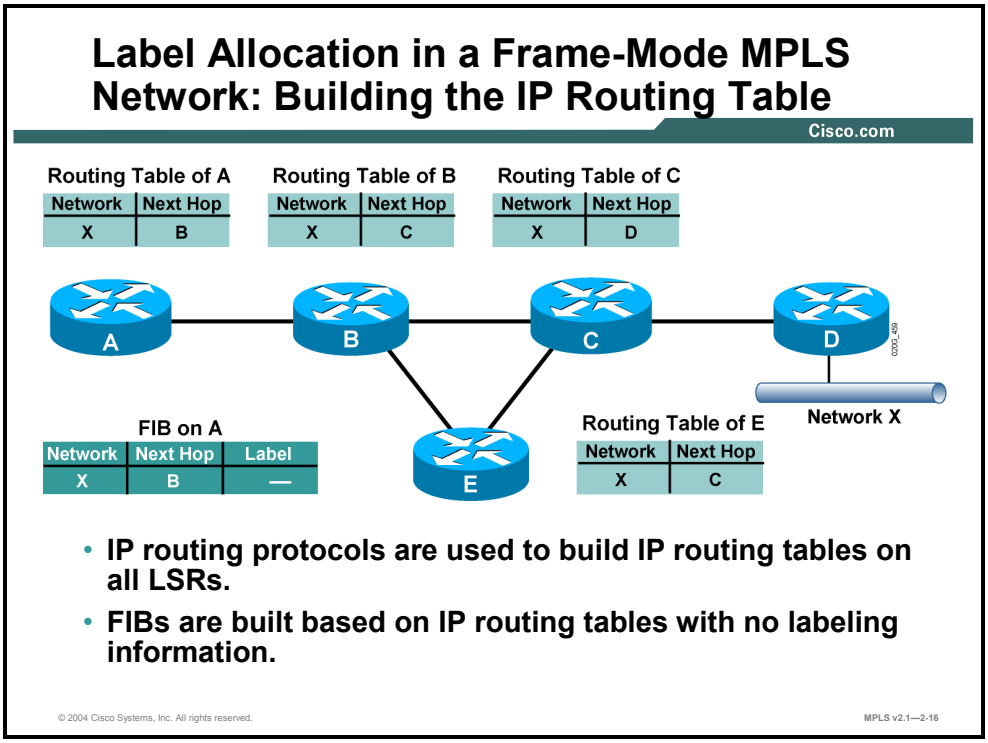

Example: Label Allocation

The figure illustrates how 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).

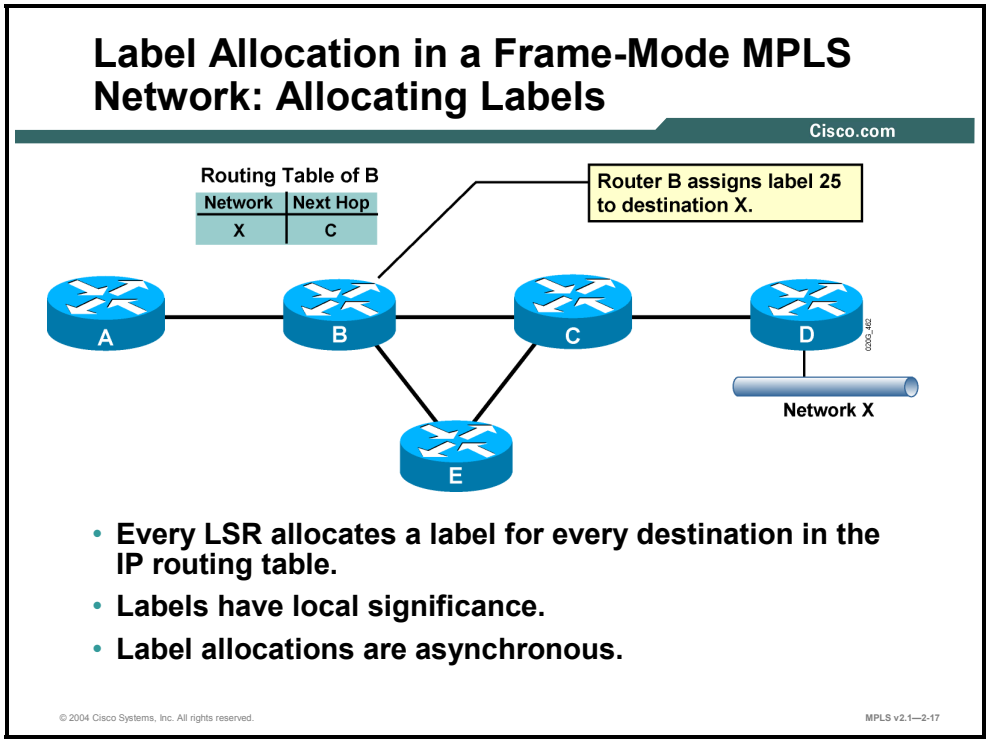

The figure shows how router B generates a locally significant and locally unique label 25 assigned to IP network X. Router B generates this label regardless of other routers (asynchronous allocation of labels).

Note	Labels 0 to 15 are reserved.

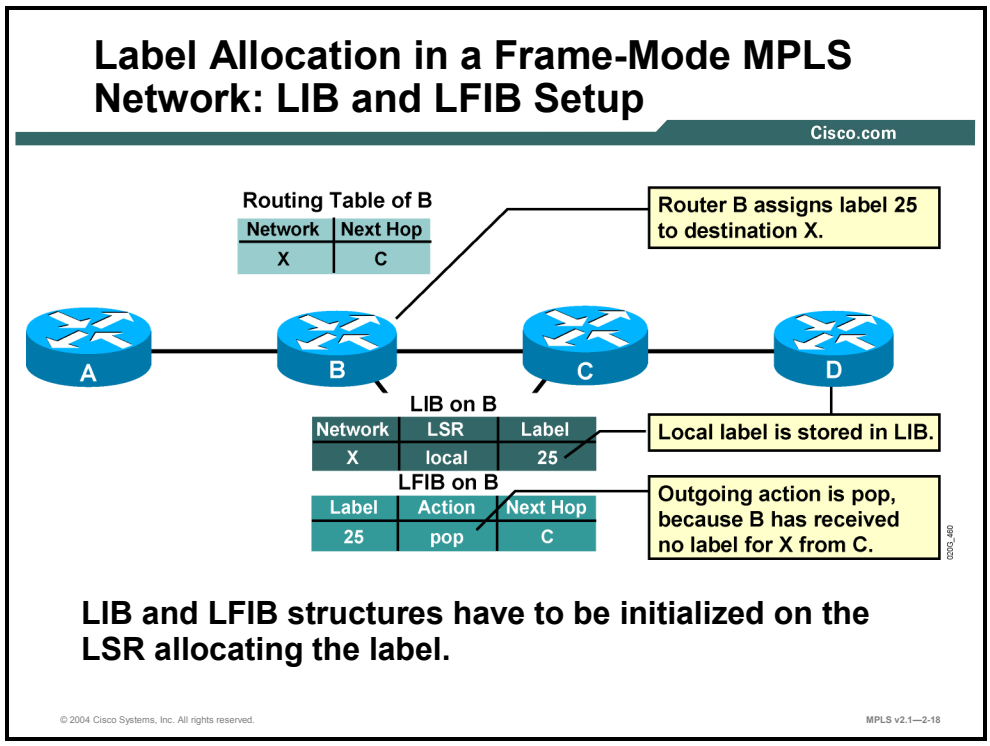

When a label is assigned to an IP prefix, it is stored in the following two tables:

- 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.

Distributing and Advertising Labels

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

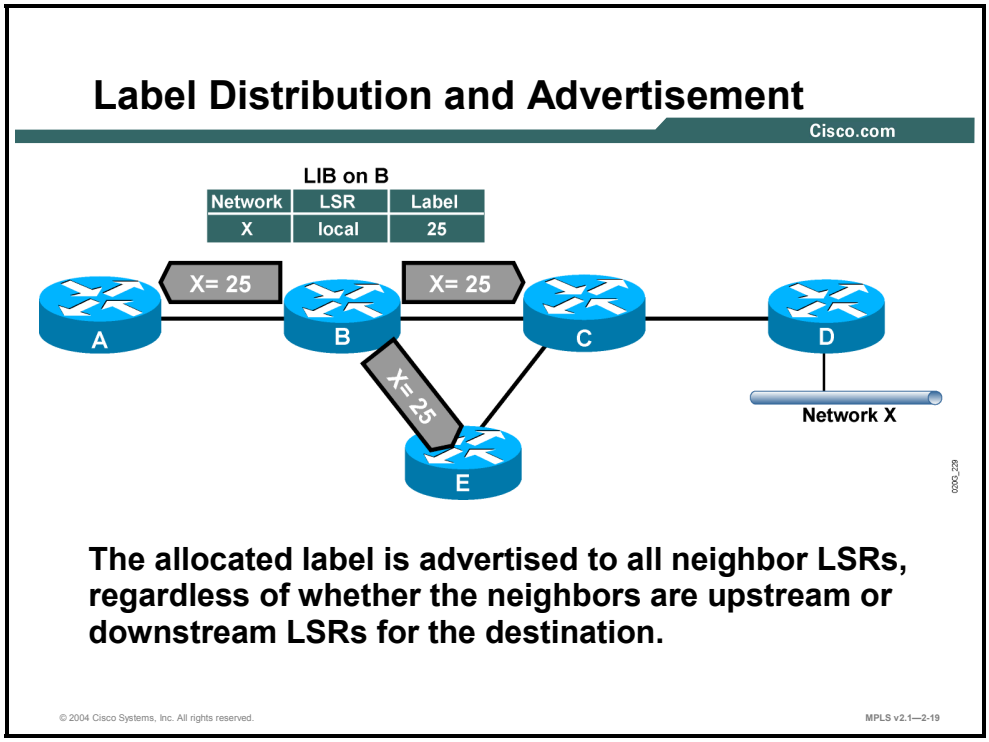

Example: Label Distribution and Advertisement

The figure illustrates the next step after a local label has been generated. 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.

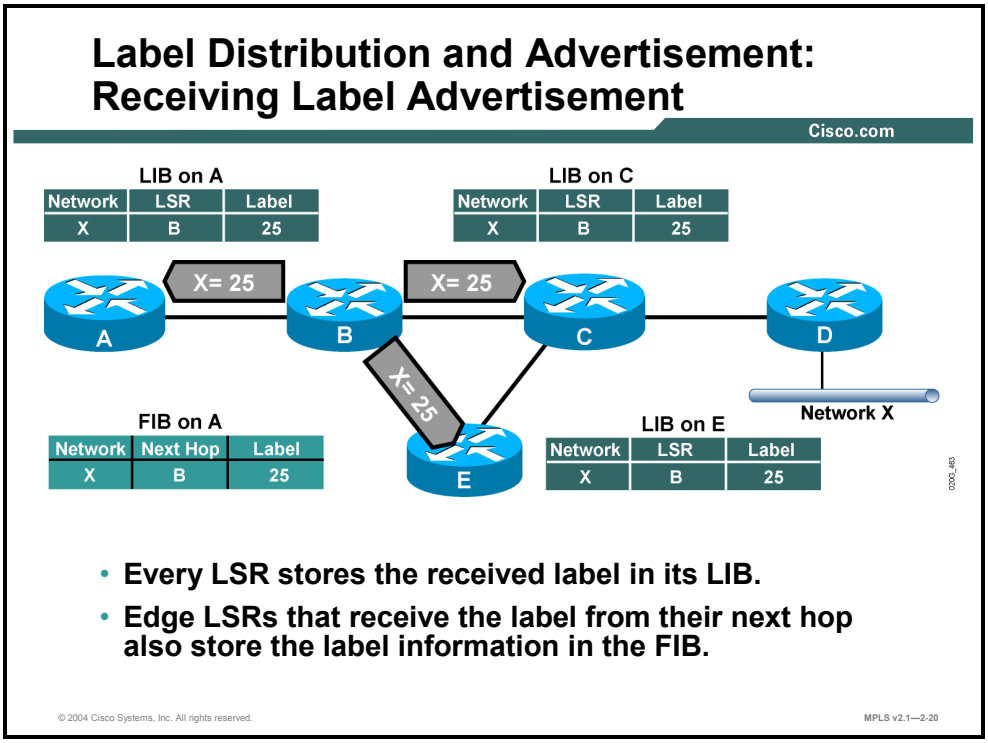

Upon receiving an LDP update, router A can fill in the missing piece in its LIB, LFIB, and FIB tables, as listed here:

- Label 25 is stored in the LIB table as the label for network X received from LSR B.
- Label 25 is 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).

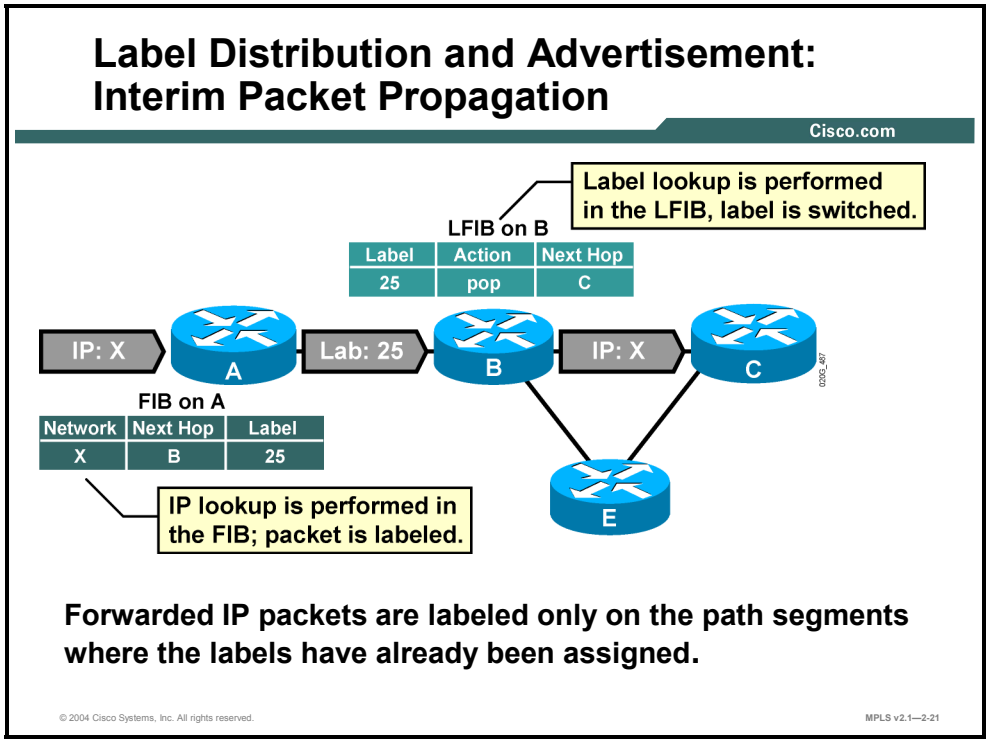

Example: Interim Packet Propagation Through an MPLS Network

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 (the action in the LFIB is “pop”).

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.

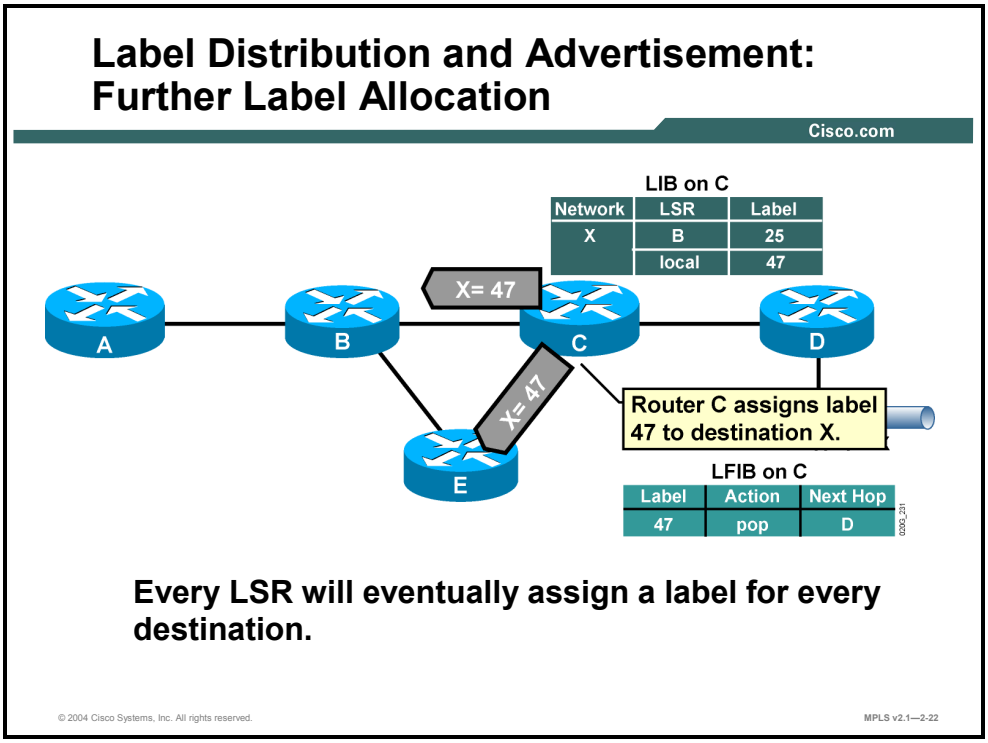

Because all routers in an MPLS domain asynchronously do the same as routers A and B, an LSP tunnel is generated, spanning from router A to router D.

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.

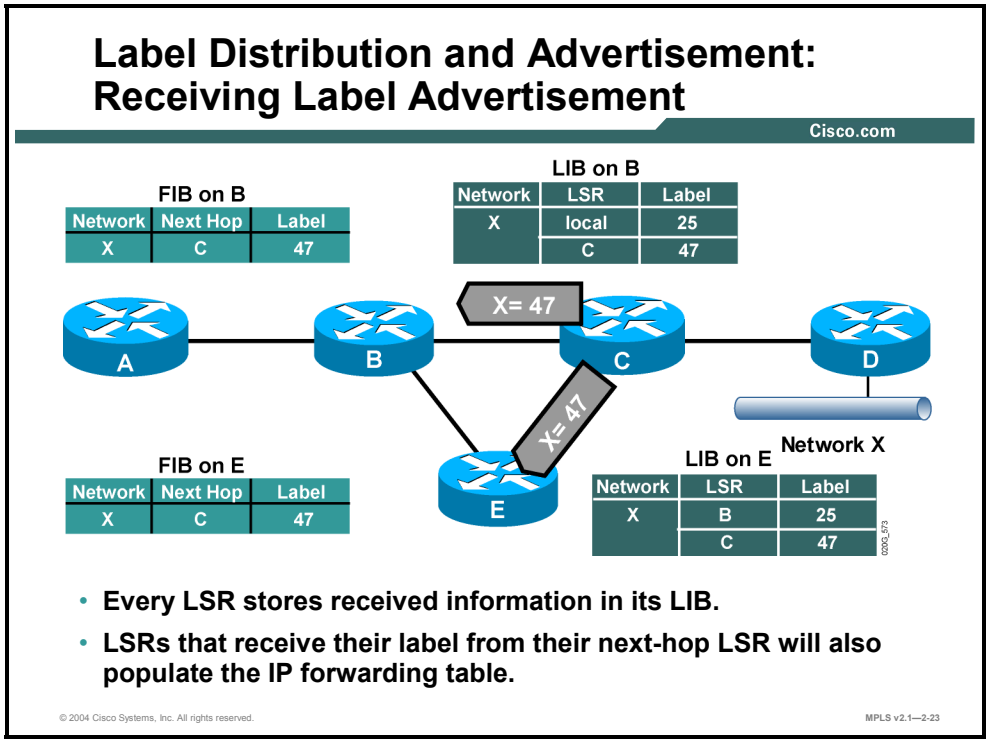

Router B can now map the entry for network X in its FIB, and the local label 25 in its LFIB, to the next-hop label 47 received from the downstream neighbor router C.

Populating LFIB

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

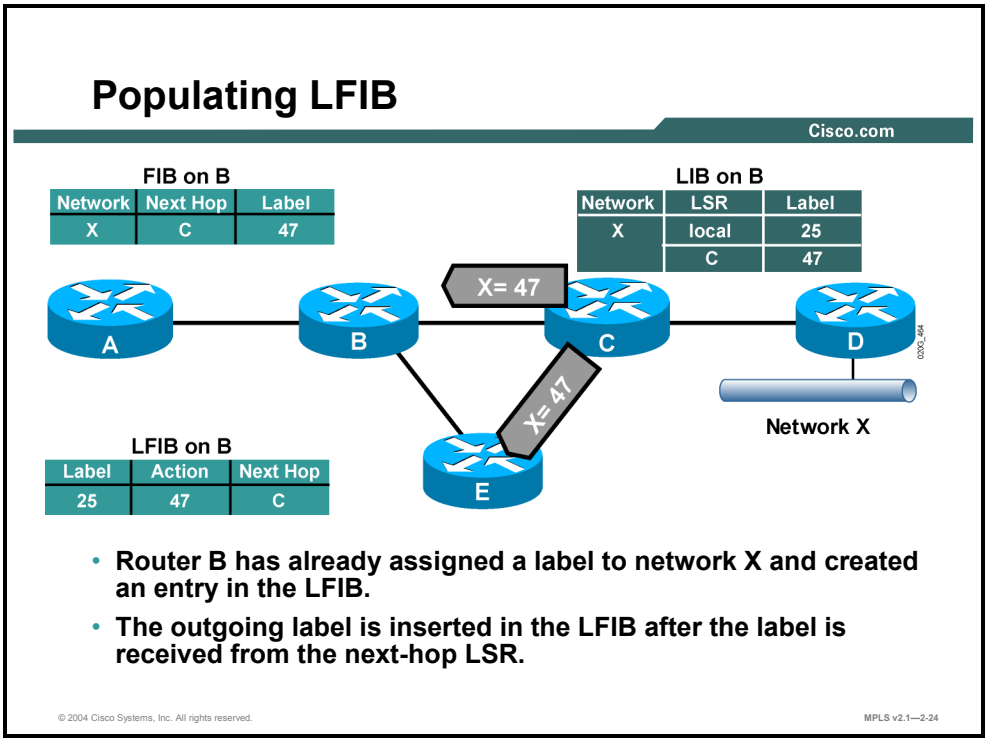

Example: Populating LFIB

After router C advertises label 47 to adjacent routers, the LSP tunnel for network X has two hops. The steps 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.

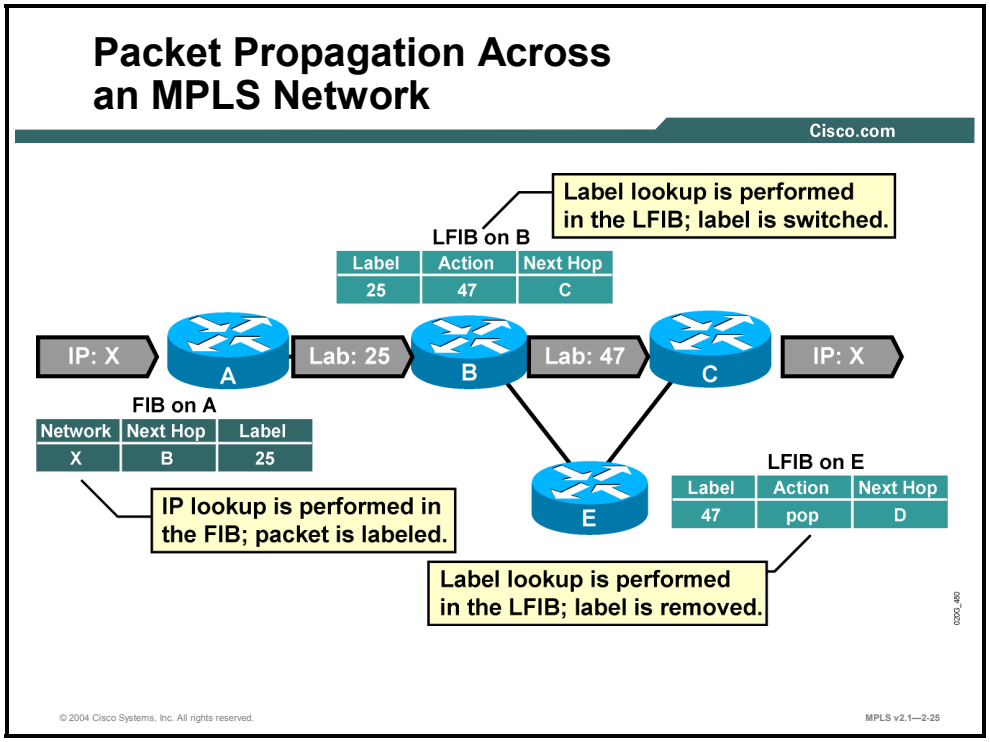

Example: Packet Propagation Through an MPLS Network

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.

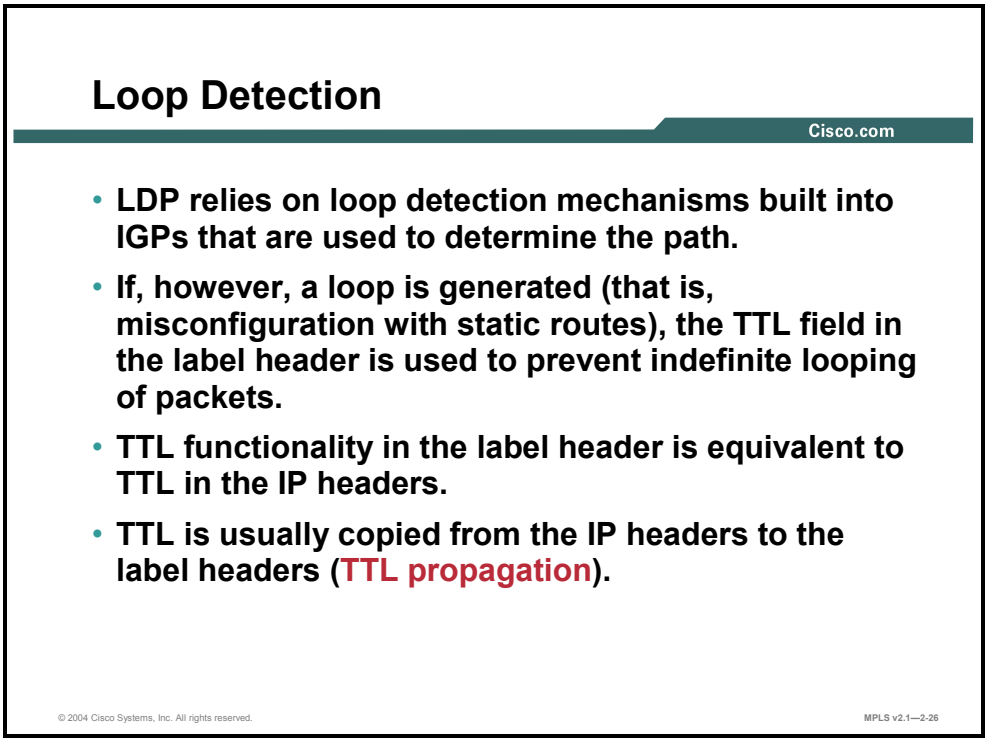

Loop detection in an MPLS-enabled network relies on more than one mechanism.

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.

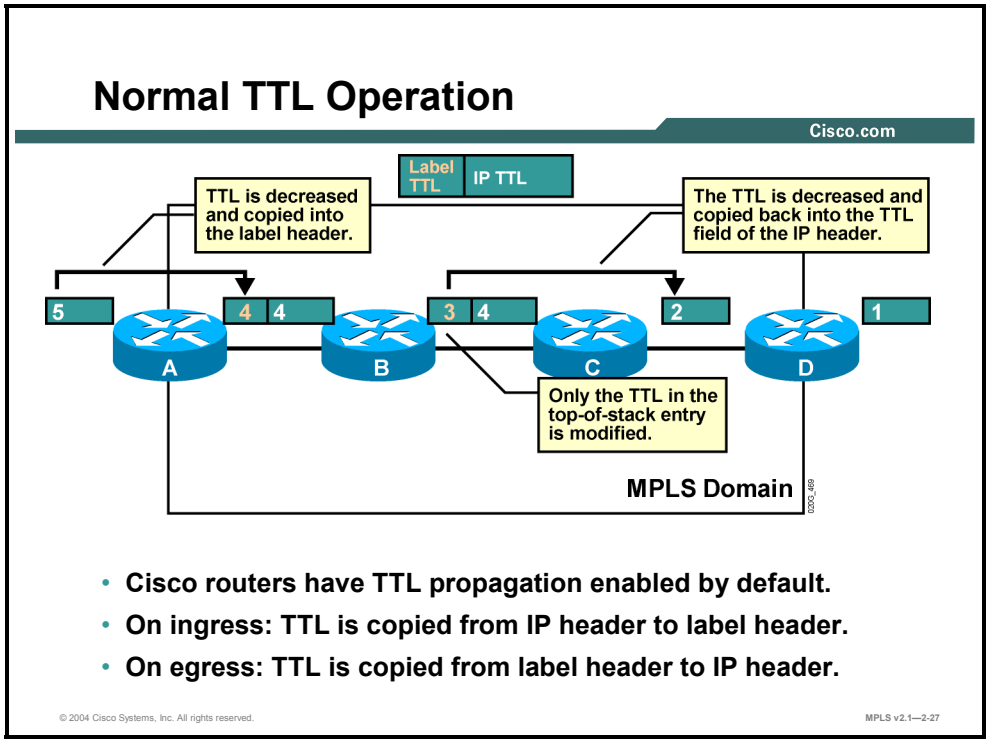

Example: Normal TTL Operation

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.

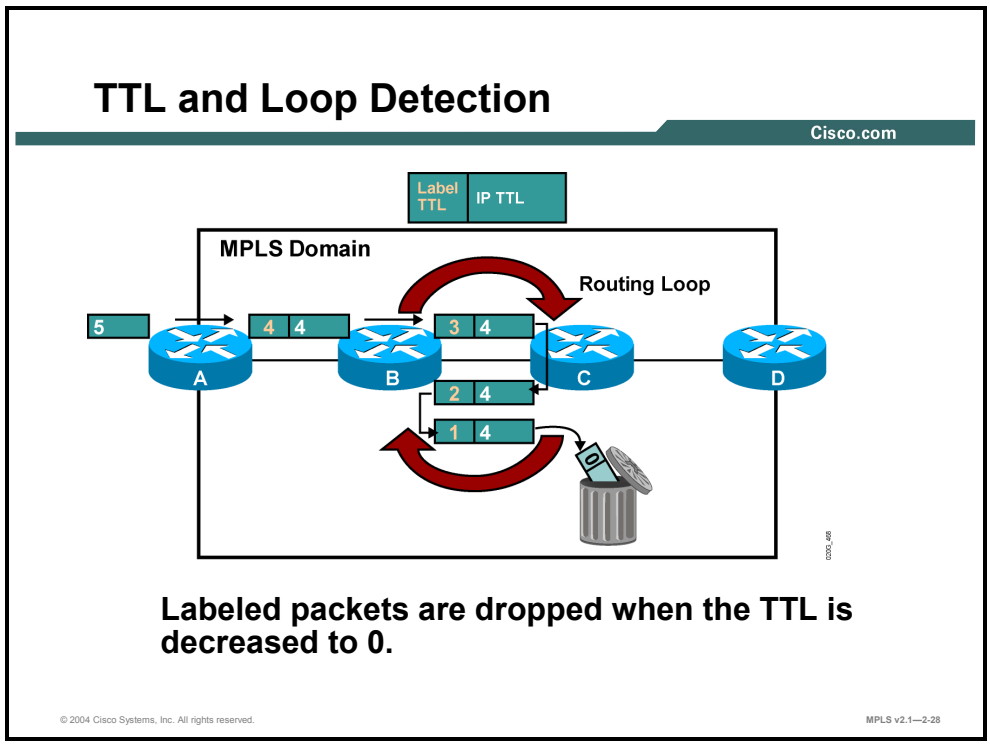

Example: TTL and Loop Detection

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.

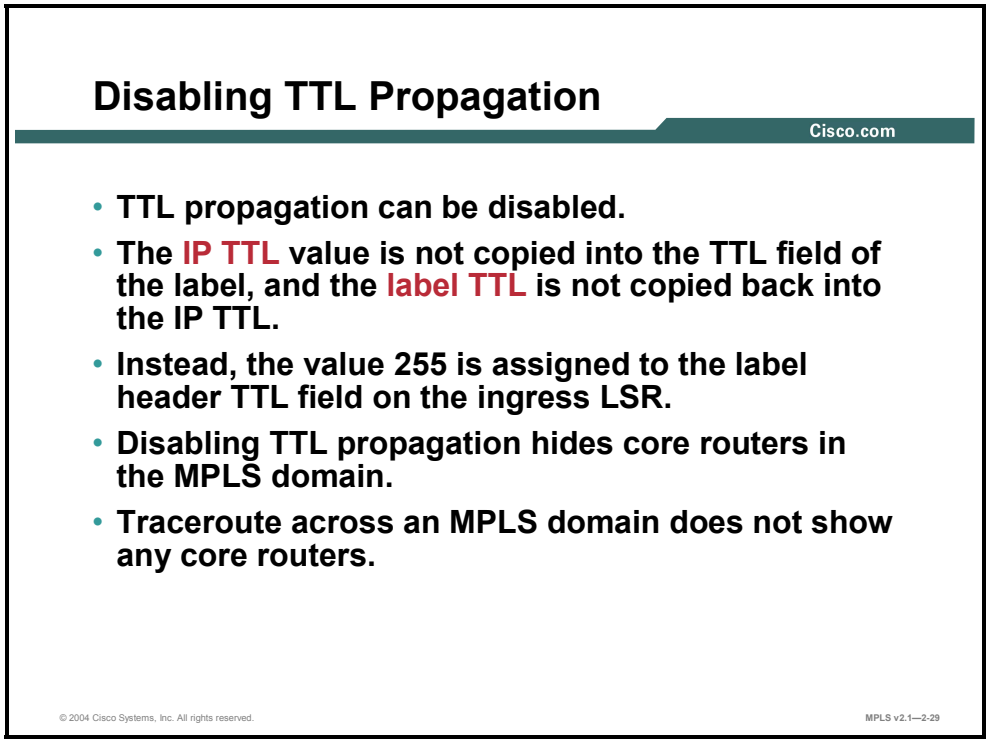

TTL propagation can be disabled to hide the core routers from the end users. Disabling TTL propagation causes routers to set the value 255 into the TTL field of the label when an IP packet is labeled.

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.

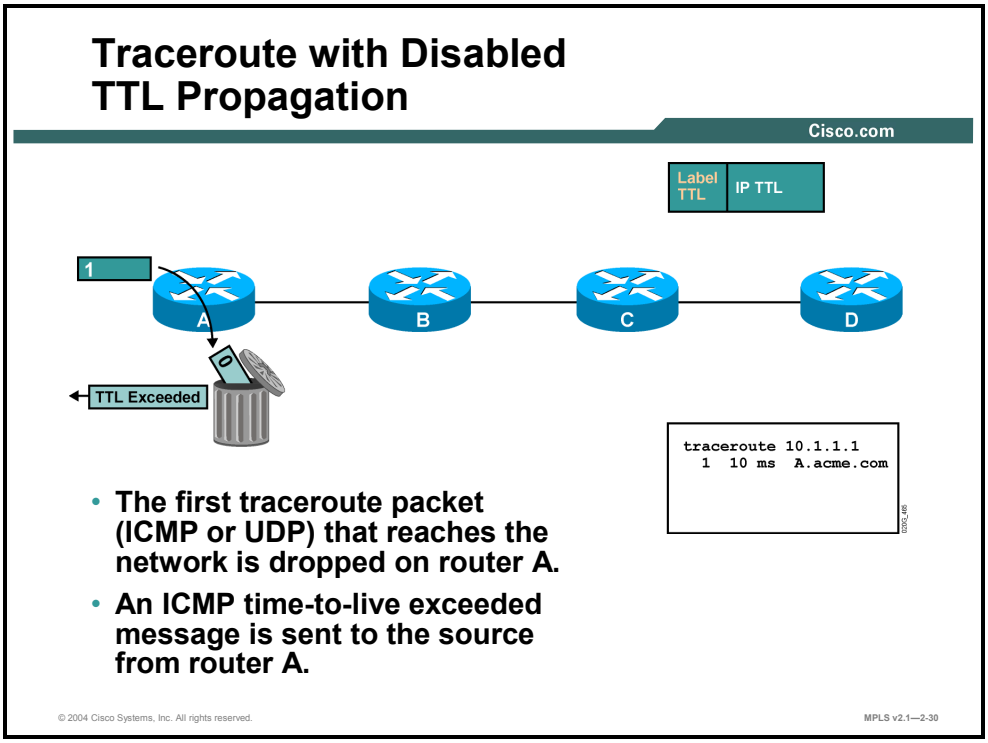

Example: Traceroute with Disabled TTL Propagation

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.

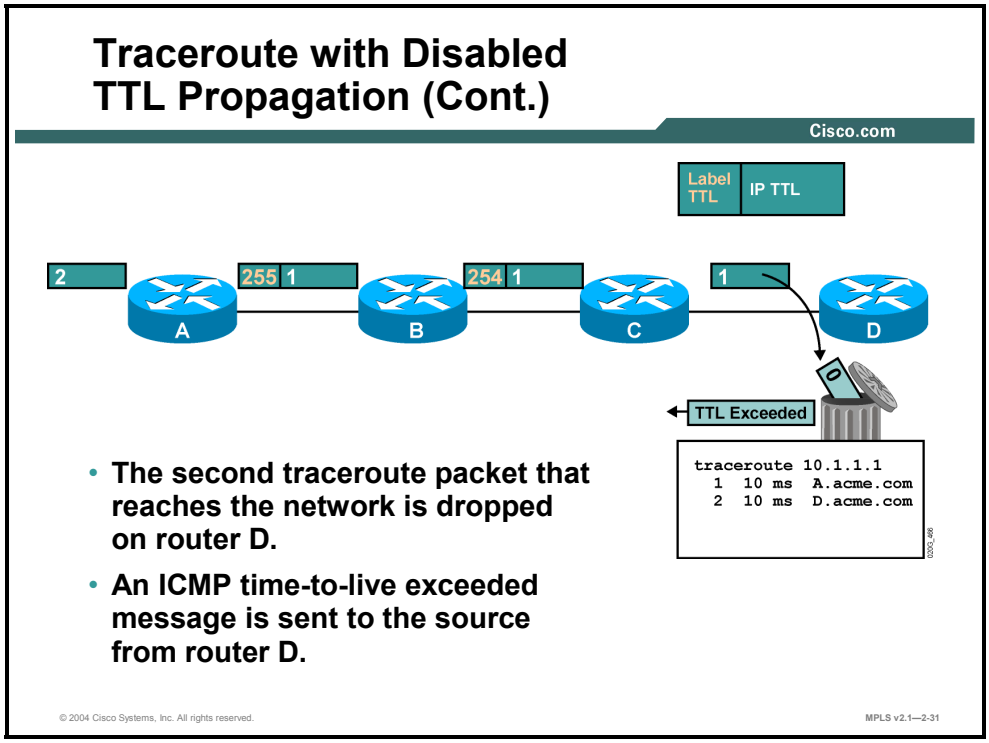

The traceroute application increases the initial TTL for every packet that it sends. The second packet, therefore, would be able to reach one hop farther (router B in the example). However, the TTL value is not copied into the TTL field of the label. Instead, router A sets the TTL field of the label 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.

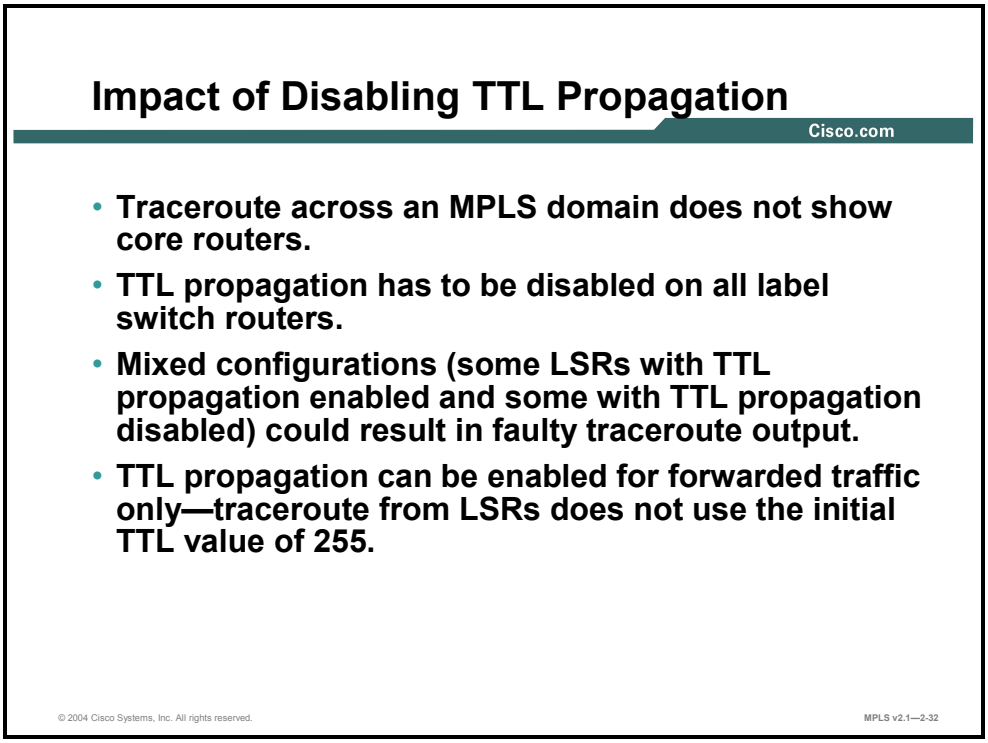

Cisco routers have TTL propagation enabled by default.

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.

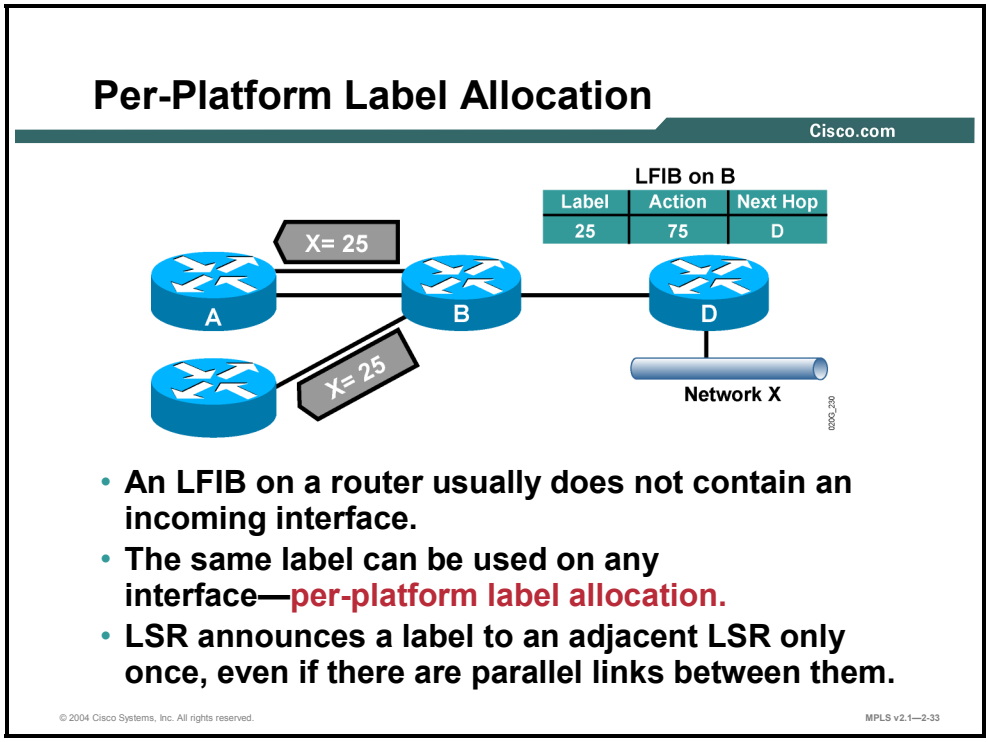

Here are the two possible approaches for assigning labels to networks:

- **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.

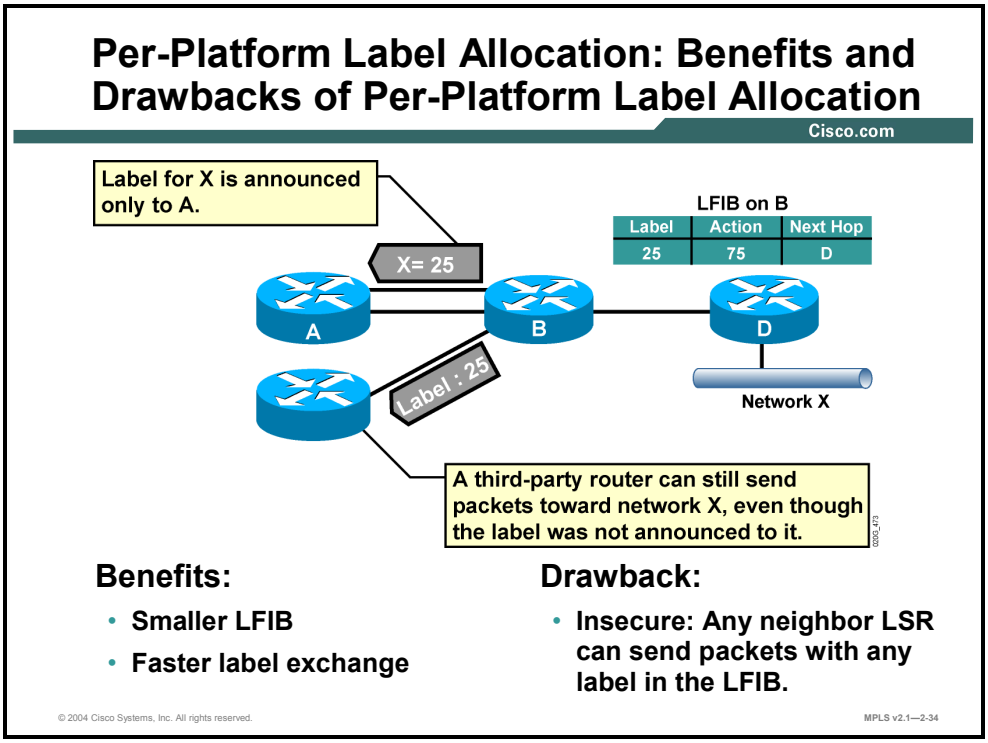

A potential drawback of per-platform label allocation is illustrated in the figure, which shows how an adjacent router can send a labeled packet with a label that has not been previously advertised to this router (label spoofing). If label switching has not been enabled on that interface, the packet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

Summary

Cisco.com

- 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.)

Cisco.com

- 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 2

Introducing Convergence in Frame-Mode MPLS

Overview

This lesson presents LDP convergence issues and describes how routing protocols and 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 after a link failure has been resolved

What Is the MPLS Steady-State Operation?

This topic describes an MPLS network steady-state operation.

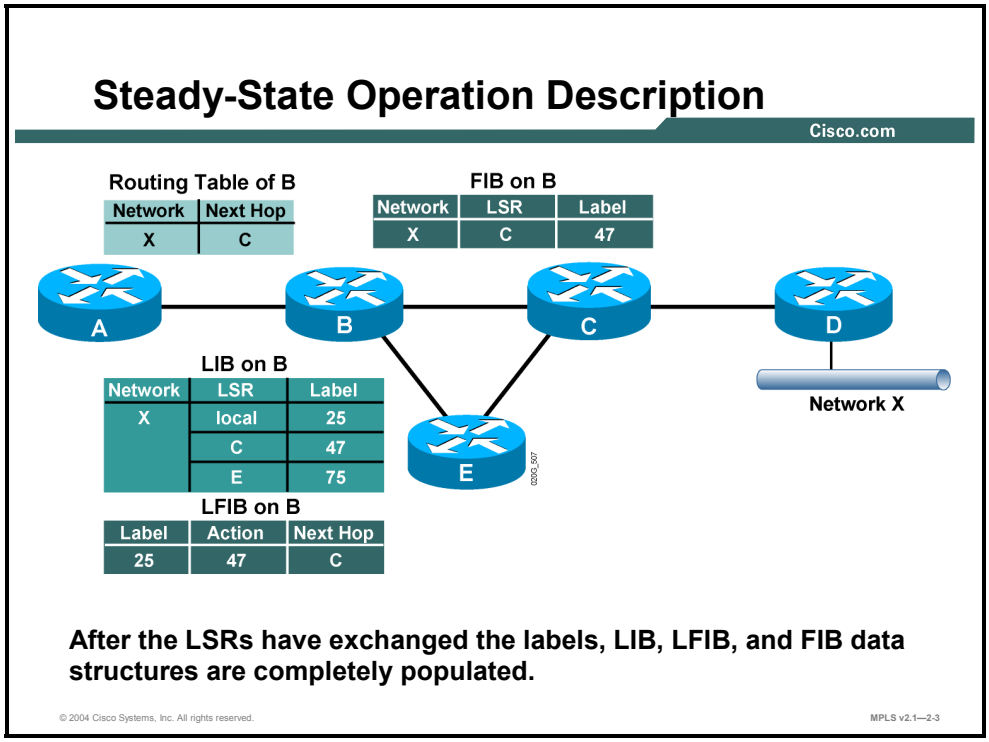

MPLS is fully functional when the IGP and LDP (or TDP) have populated all the tables, as listed here:

- Main IP routing table
- LIB table
- FIB table
- LFIB table

Although it takes longer for LDP to exchange labels (compared with IGP), a network can use the FIB table in the meantime, so 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.

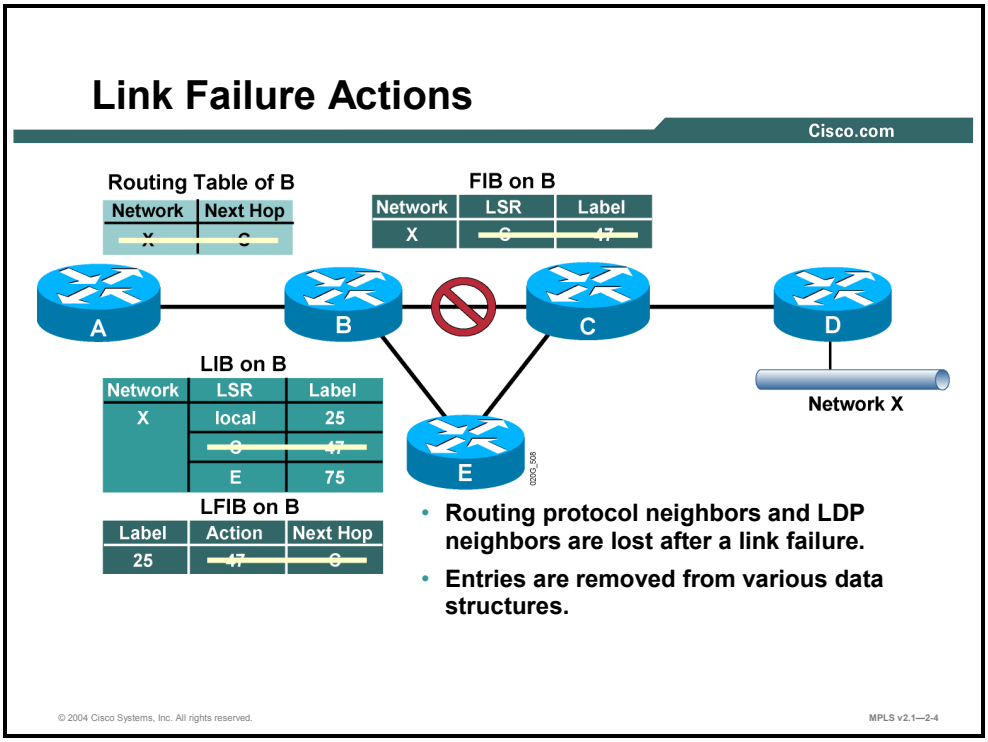

Example: Link Failure Actions

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.

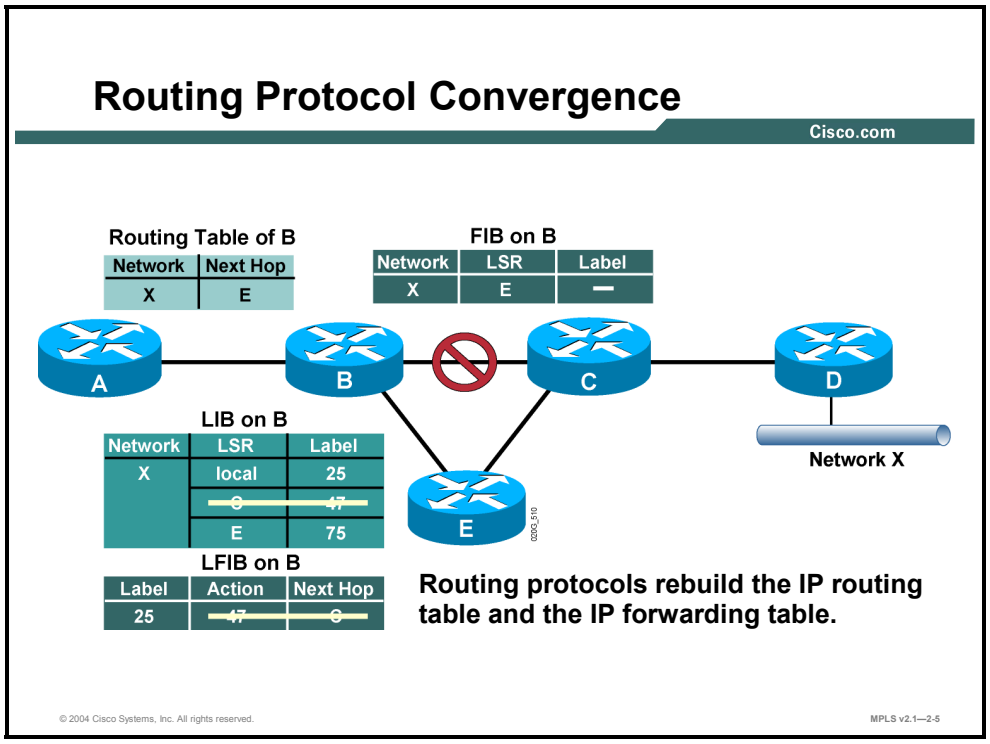

Example: Routing Protocol Convergence

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.

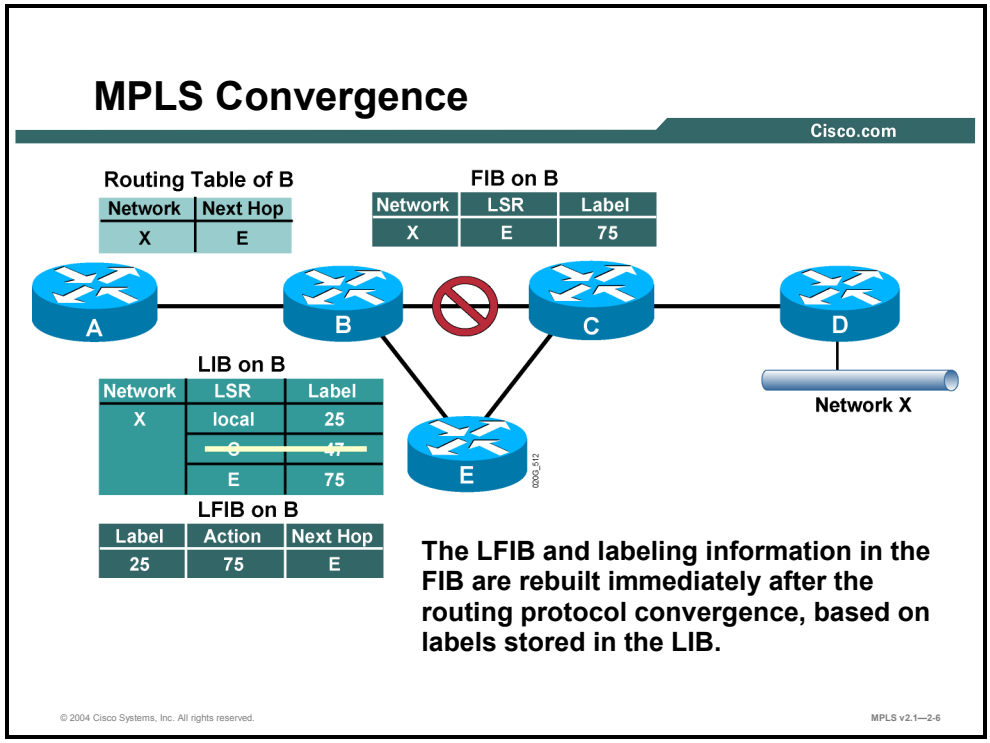

After the IGP determines that there is another path available, a new entry is created in the FIB table.

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.

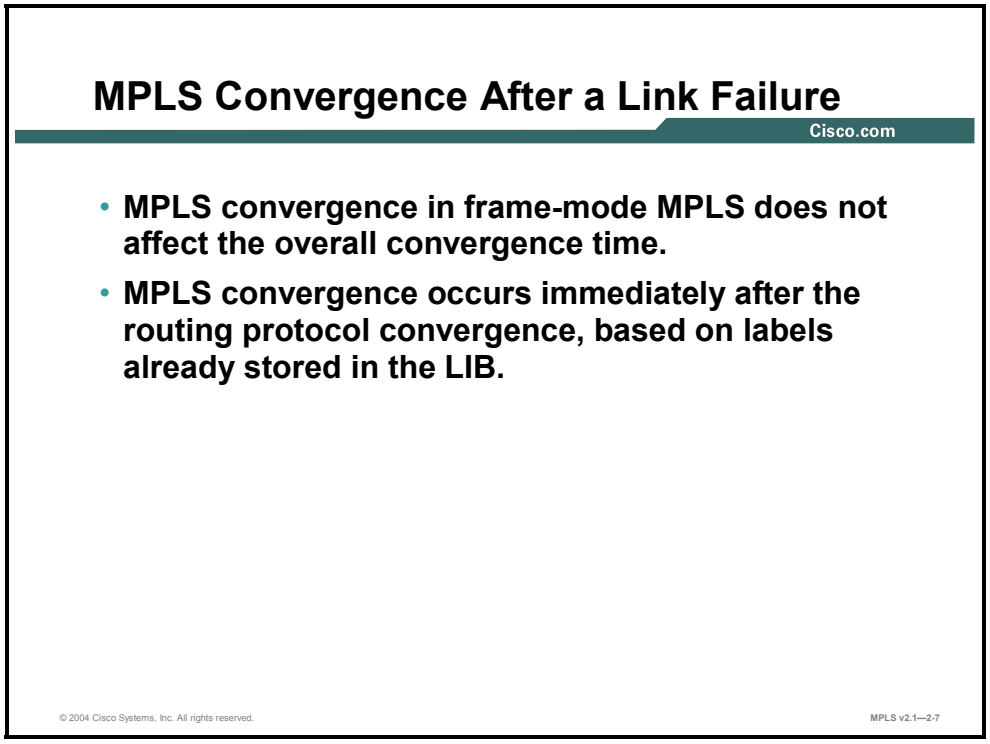

The overall convergence in an MPLS network is not affected by LDP convergence when there is a link failure.

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 Happens in Link Recovery?

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

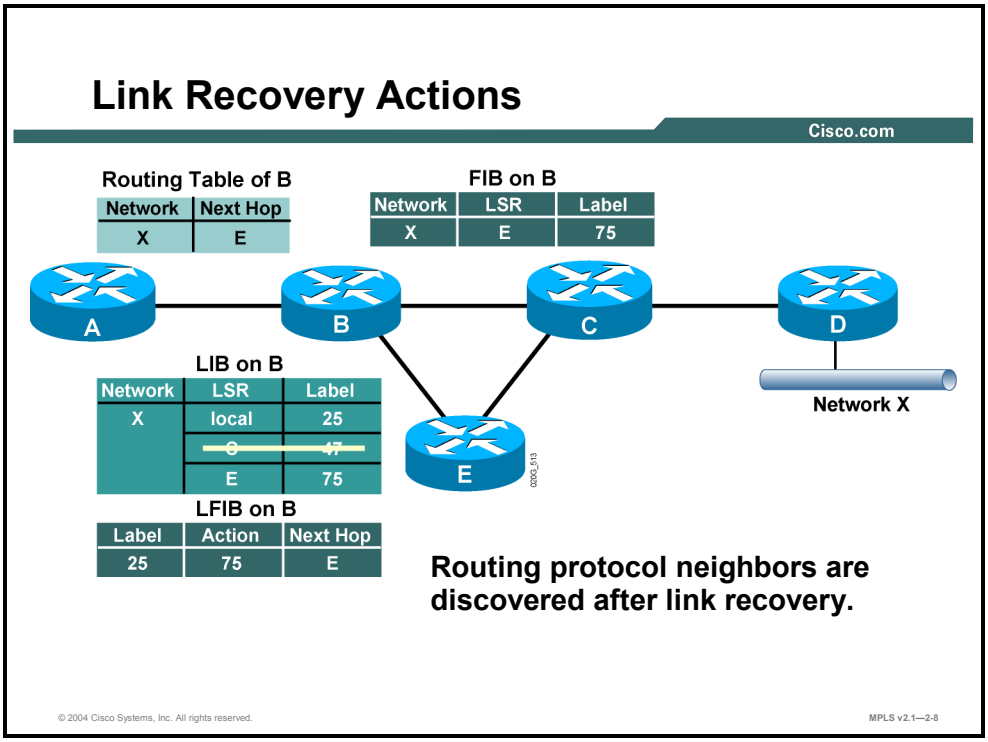

Example: Link Recovery Actions

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

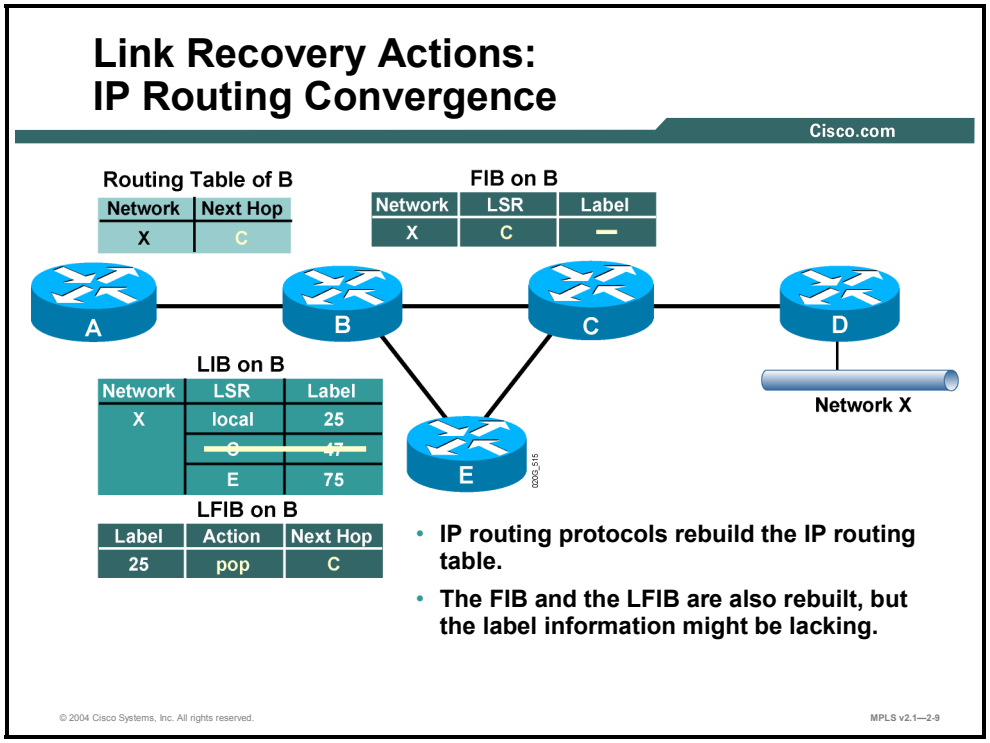

The IGP determines that the link is available again and changes the next-hop address for network X to point to router C. However, when router B also tries to set the next-hop label for network X, it has to wait for the LDP session between routers B and C to be reestablished.

A pop action is used in the LFIB 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 VPN networks. This is because the VPN traffic cannot be forwarded before the LDP session is fully operational.

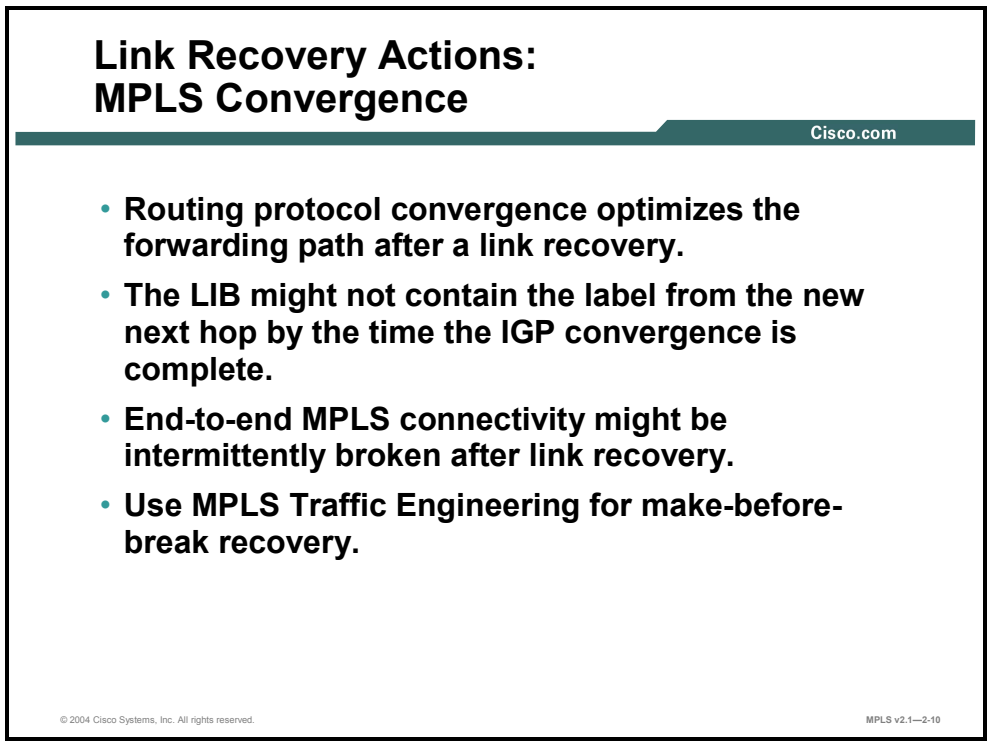

Link recovery requires that an LDP session be established (reestablished), which adds to the convergence time of LDP.

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

MPLS TE can be used to prevent longer downtime when a link fails or is recovering.

Summary

This topic summarizes the key points discussed in this lesson.

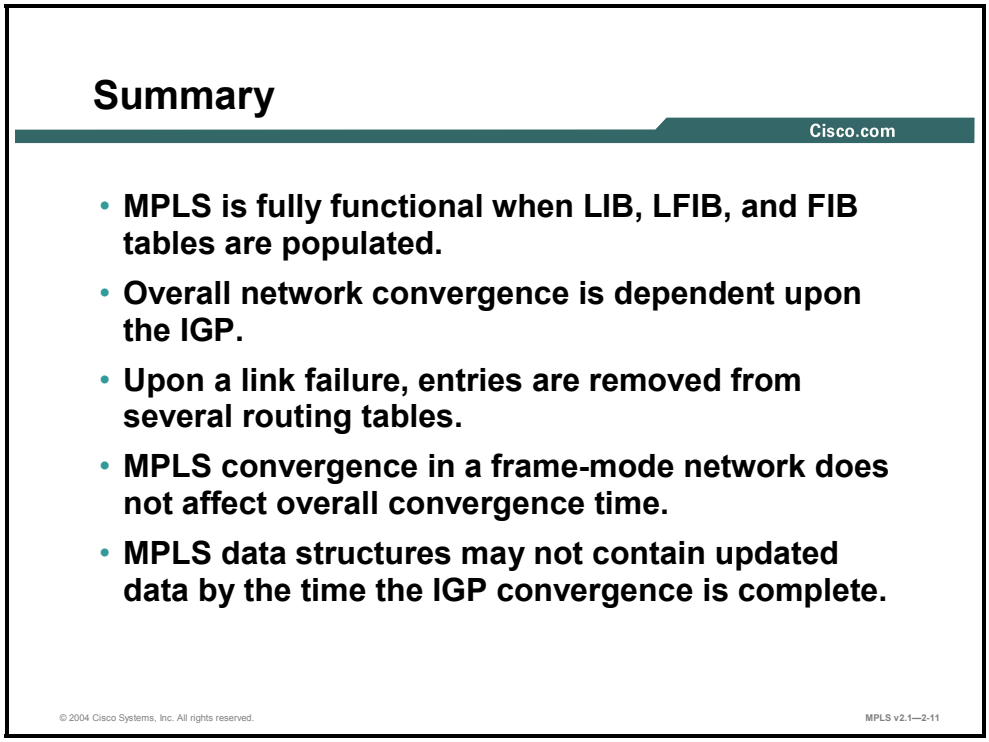

Lesson 3

Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge

Overview

This lesson describes how tables are built and how labels are processed in cell-mode MPLS networks. The lesson also introduces a concept called virtual circuit merge (VC merge). It is important to understand the differences between label distribution in frame-mode MPLS networks and cell-mode MPLS networks. This lesson explores some of the key differences when a cell-mode network is deployed.

Objectives

Upon completing this lesson, you will be able to describe typical label distribution over label controlled-ATM (LC-ATM) interfaces and VC merge. This ability includes being able to meet these objectives:

- Identify issues that can arise in cell-mode MPLS network deployments
- Describe how the IP routing table is populated in a cell-mode MPLS network
- Describe how the IP forwarding table is populated in a cell-mode MPLS network
- Describe how labels are requested in cell-mode MPLS networks
- Describe how labels are allocated in cell-mode MPLS networks
- Identify the issues that can occur with the interleaving of cells in cell-mode MPLS networks
- Describe the characteristics of VC merge
- Describe how loop detection is managed in cell-mode MPLS networks
- Describe the characteristics of per-interface label allocation

What Are Cell-Mode MPLS Network Issues?

This topic describes the issues that can arise in cell-mode MPLS network deployments.

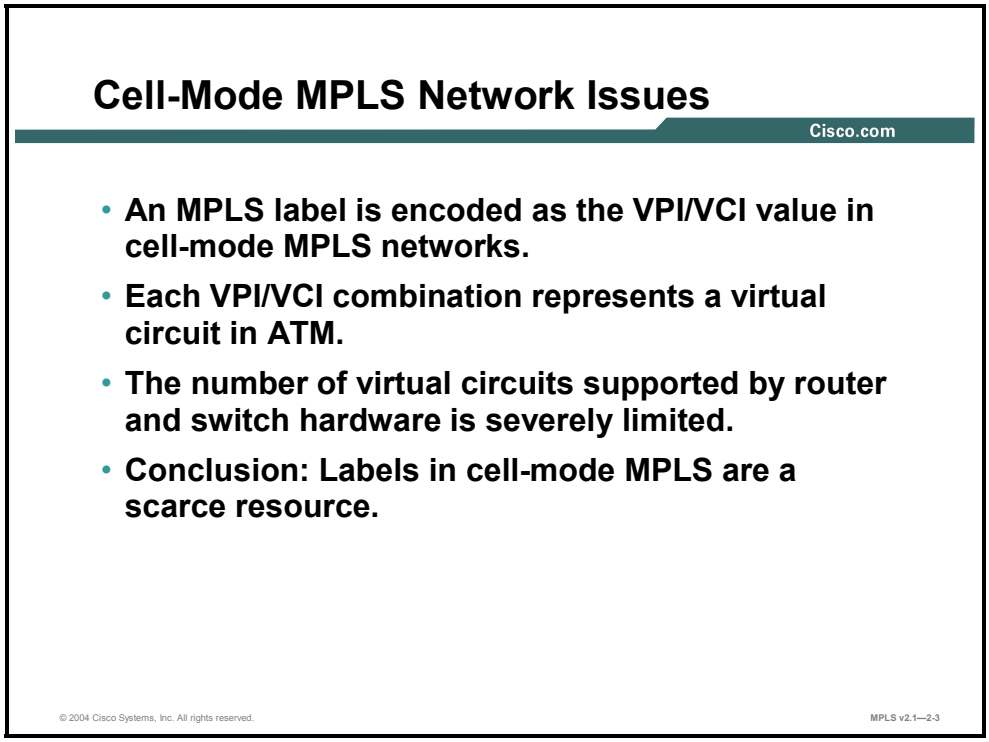

Cell-mode MPLS is significantly different from frame-mode MPLS because of some ATM-specific requirements. Some of the differences are as follows:

- ATM uses cells, not frames. A single packet may be encapsulated into multiple cells. Cells are a fixed length, which means that normal labels cannot be used because they would increase the size of a cell. The virtual path identifier/virtual channel identifier (VPI/VCI) field in the ATM header is used as the MPLS label. An LSP tunnel is therefore called a virtual circuit in ATM terminology.
- ATM switches and routers usually have a limited number of virtual circuits that they can use. MPLS establishes a full mesh of LSP tunnels (virtual circuits), which can result in an extremely large number of tunnels.

Additional mechanisms must be used because of the limitations of ATM hardware.

Building the IP Routing Table

This topic describes how the IP routing table is populated in cell-mode MPLS networks.

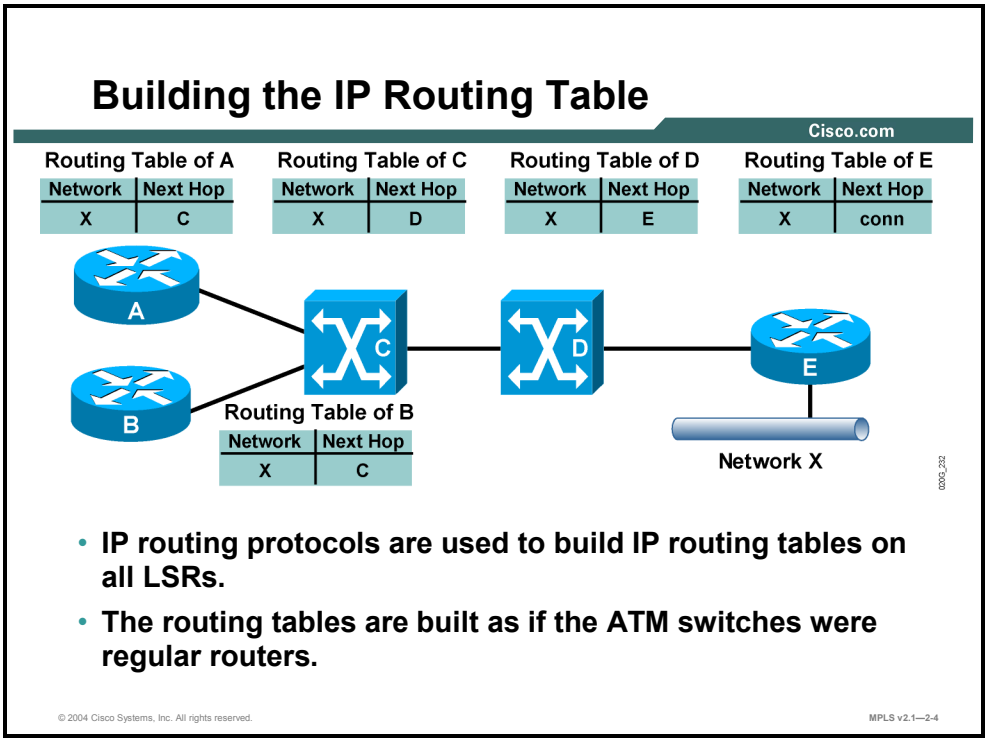

Example: Building the IP Routing Table

The figure shows how IP- and MPLS-aware ATM switches exchange IP routing information with routers.

On the control plane, each ATM switch acts as an IP router, and the routing tables are built as if the ATM switches were routers.

Because the ATM switch acts as an IP router, it is seen as an extra IP hop in the network.

Building the IP Forwarding Table

This topic describes how the IP forwarding table is populated in cell-mode MPLS networks.

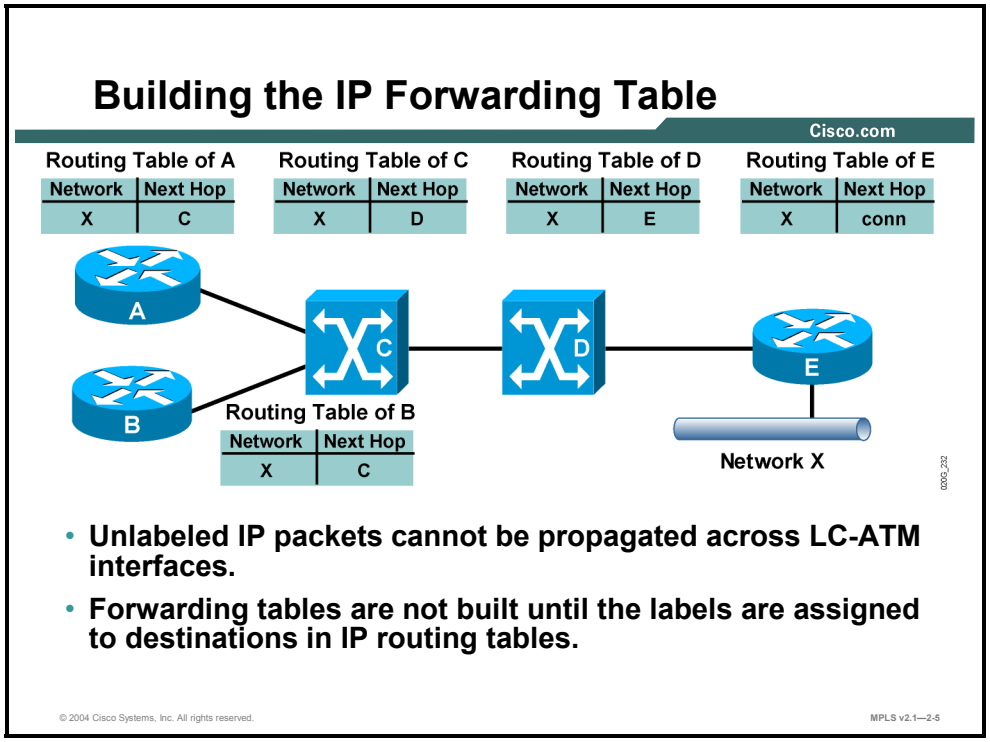

Because ATM switches cannot forward IP packets, labels cannot be asynchronously assigned and distributed.

Instead, the router initiates an ordered sequence of requests on the upstream side of the ATM network.

It is not until the request is answered, with the label and assigned to destinations in the IP routing table, that the forwarding table is populated.

Requesting a Label

This topic describes how labels are requested in cell-mode MPLS networks.

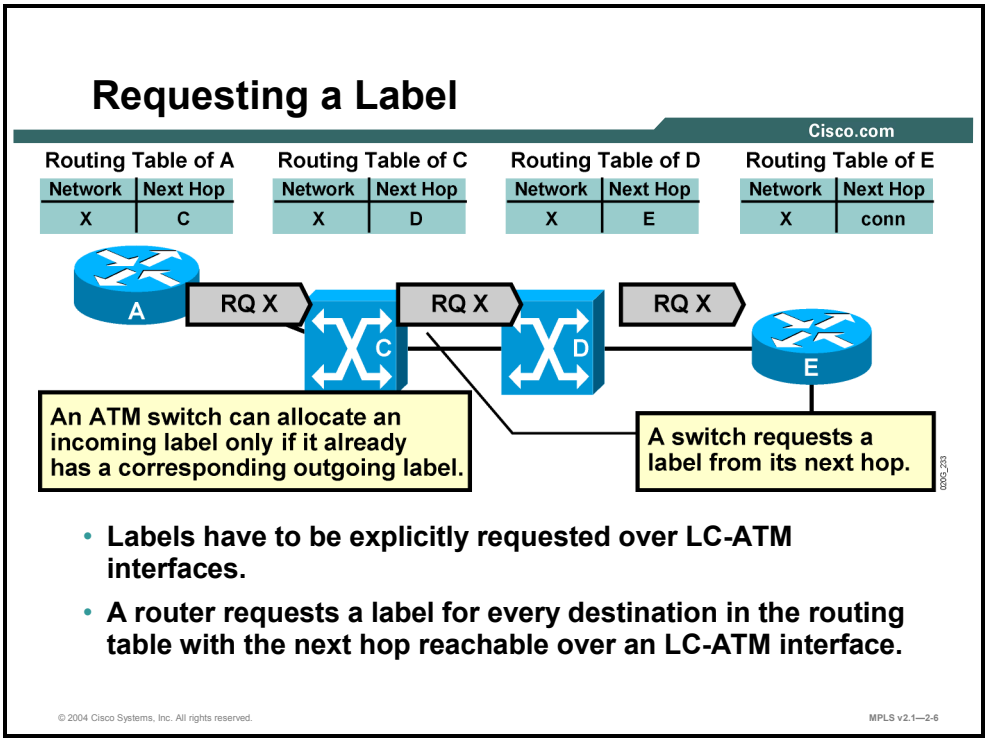

Example: Requesting a Label

In the figure, a request is sent from router A to the ATM switch C. Because the ATM switch cannot perform IP lookups, the switch is not allowed to reply with the local label unless it already has the next-hop label. If switch C does not have the next-hop label, switch C must forward the request to the next downstream neighbor, ATM switch D.

If switch D does not have the next-hop label, switch D must forward the request to the next downstream neighbor.

When the request reaches router E, a reply can be sent because the cell-mode part of the network ends on router E (which, therefore, acts as an ATM edge LSR).

Allocating a Label

This topic describes how labels are allocated in cell-mode MPLS networks.

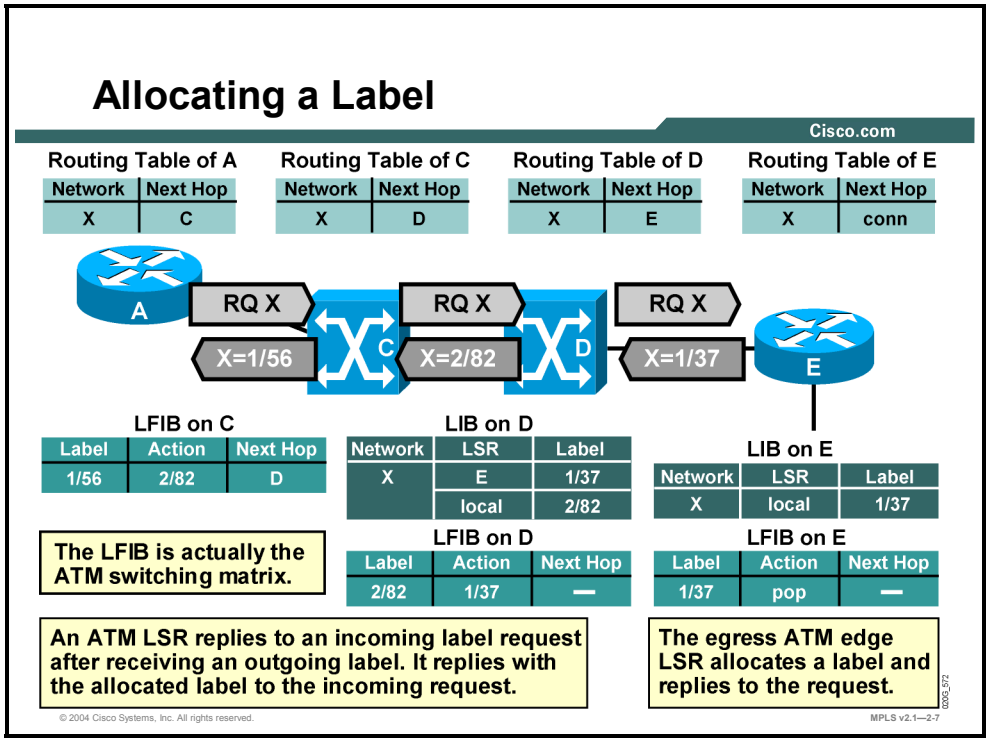

Example: Allocating a Table

In the figure, router E replies with its local label 1/37. The ATM switch D can now generate and use its local label 2/82. Switch C receives the next-hop label from switch D and forwards its own local label 1/56 to router A.

As seen in the figure, an ordered sequence of downstream requests is followed by an ordered sequence of upstream replies. This type of operation is called downstream-on-demand allocation of labels.

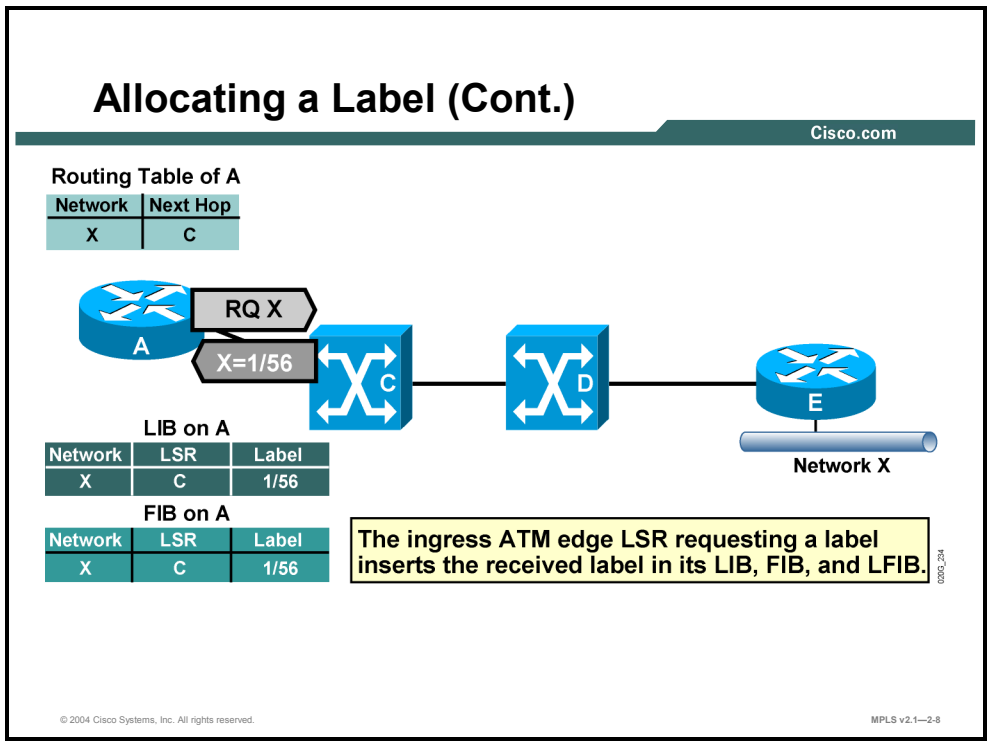

The processing of LDP replies on router A (also an ATM edge LSR) is similar to processing in frame-mode MPLS; the received label is stored in the LIB, FIB, and LFIB tables.

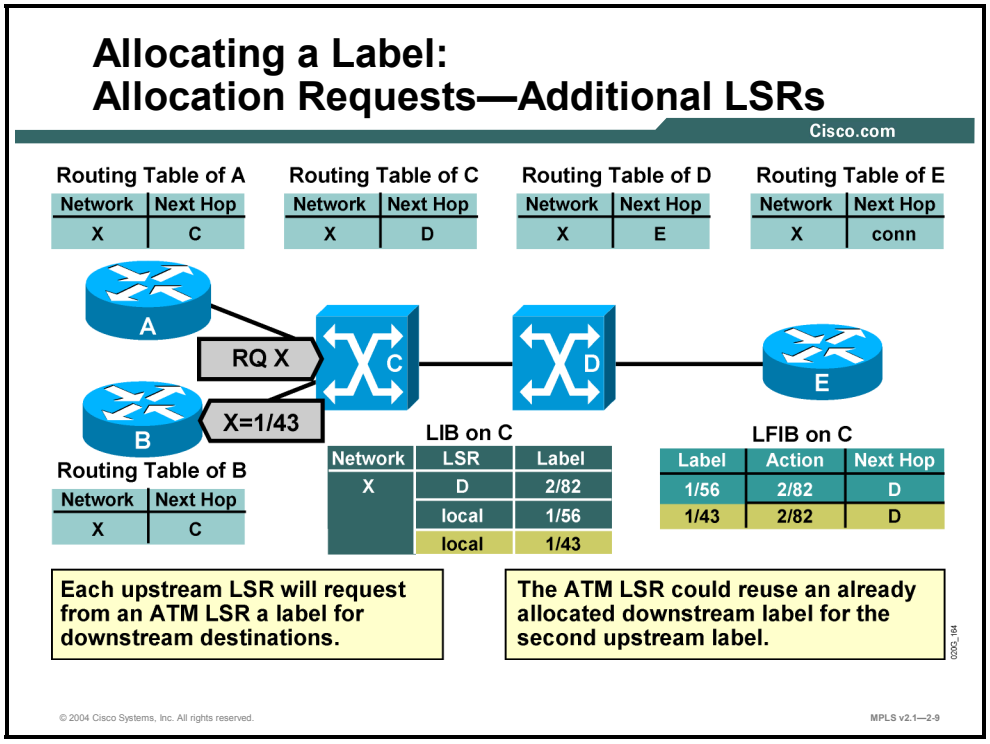

Example: Additional LSRs

The figure shows how another router, router B, requests a label for the same destination that router A has previously requested. The ATM switch C already has a next-hop label for network X and, therefore, can immediately reply to router B.

The figure also shows that the switch used a different local label, 1/43, from the label sent to router A, 1/56. This is because ATM switches use per-interface VPI/VCI values and can now also use per-interface label space.

What Are Cell Interleave Issues?

This topic identifies the issues that can occur with the interleaving of cells in cell-mode MPLS networks.

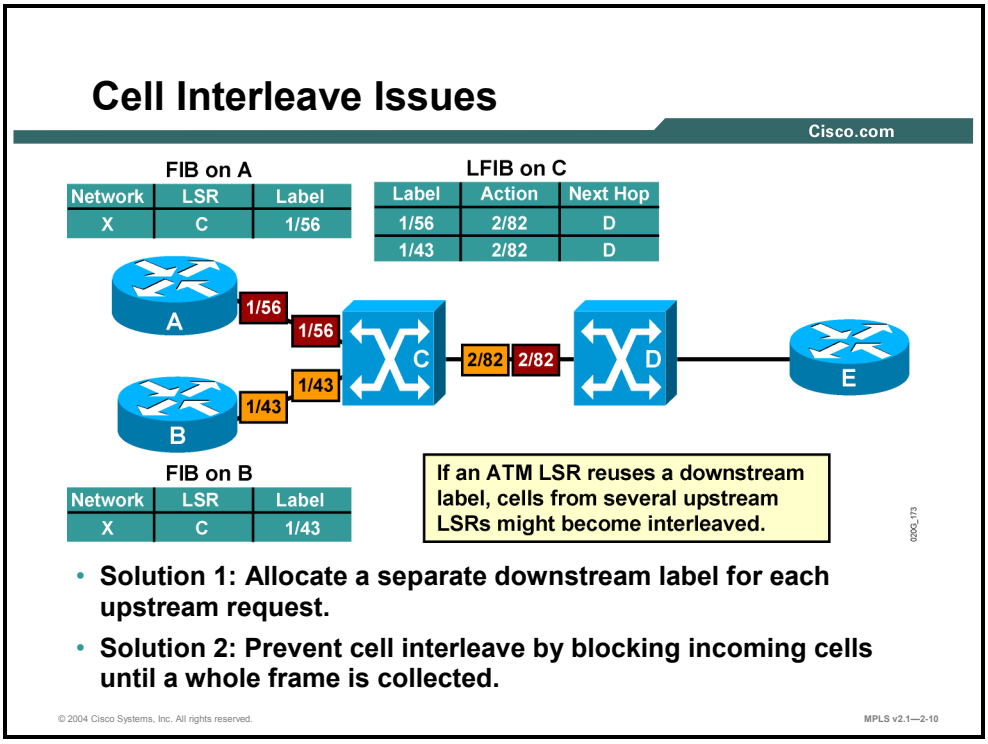

Routers A and B request for the same network X has resulted in an unusual situation. Two virtual circuits from routers A and B (1/56 and 1/43) merge into one VC (2/82).

Standard ATM virtual switching hardware does not support this situation and, as a result, segmented packets from the two sources may become interleaved between the ATM switches C and D.

The receiving router, E, is then unable to correctly reassemble those cells into two packets.

Here are the two possible solutions to this problem:

- Allocate a new downstream label for each request. This solution would result in a greater number of labels.
- Buffer the cells of the second packet until all cells of the first packet are forwarded. This solution results in an increased delay of packets because of buffering.

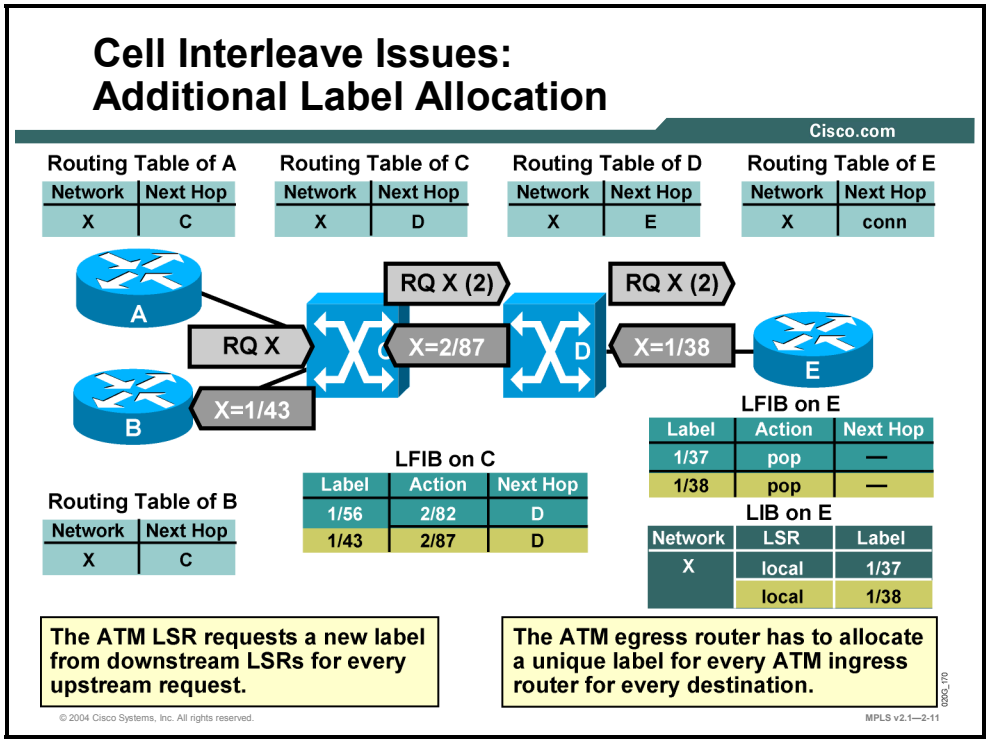

Example: Additional Label Allocation

This figure illustrates the first option, where an additional LSP tunnel is created for the same destination network X for every upstream ATM edge LSR.

ATM switch C now has two next-hop labels for network X, one for source router A and the other for source router B.

What Is VC Merge?

This topic describes the characteristics of VC merge.

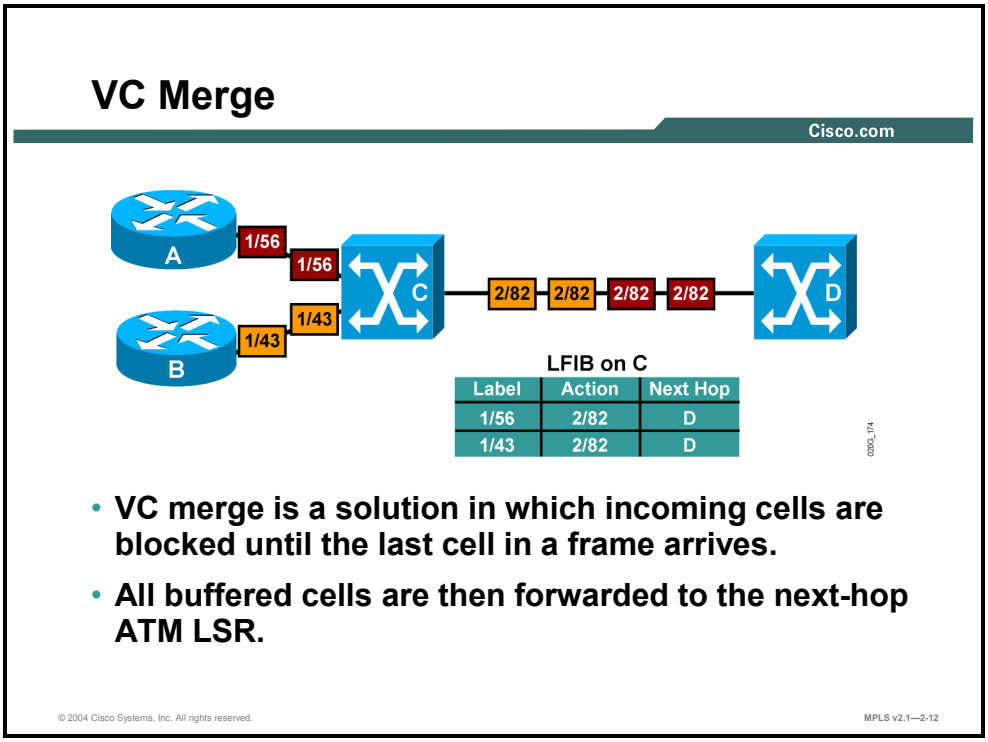

Example: VC Merge

The figure illustrates the second option, where the ATM switch C buffers cells coming from router B until the last cell of the packet coming from router A is forwarded.

This option reduces the number of labels (virtual circuits) needed in the ATM network, but increases the average delay across the network.

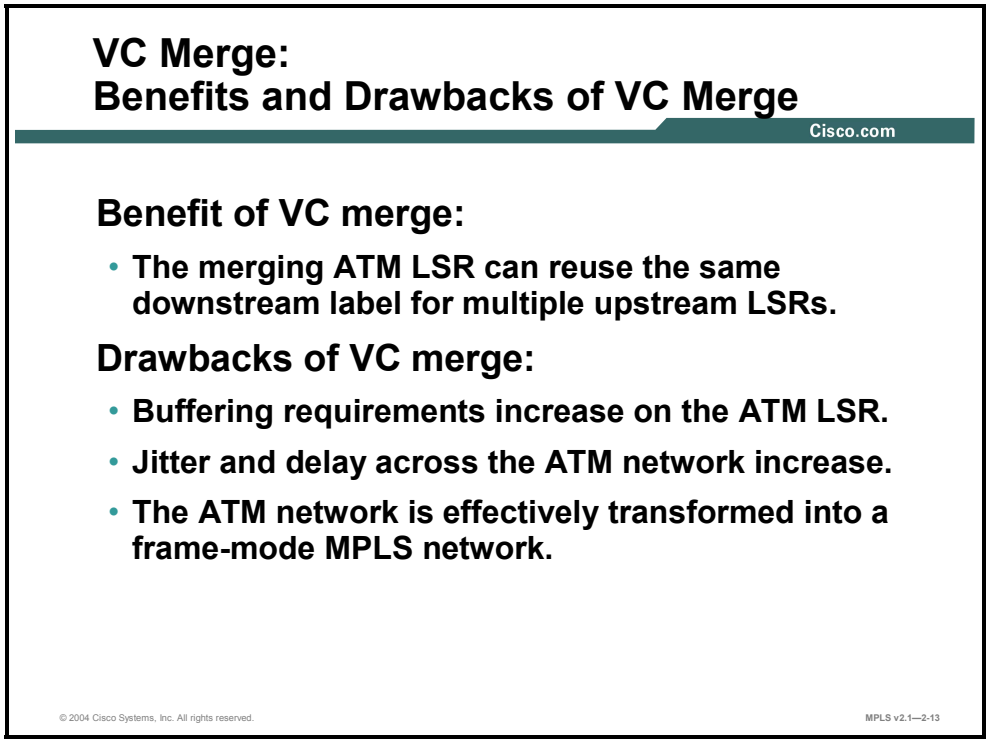

The major benefit of VC merge is that it minimizes the number of labels (VPI/VCI values) needed in the ATM part of the network. As identified in the topic What Are Cell-Mode MPLS Network Issues, labels are a scarce resource in cell-mode MPLS networks.

The major drawbacks to VC merge are as follows:

- Buffering requirements increase on the ATM LSR.
- There is an increase in delay and jitter in the ATM network.
- ATM networks under heavy load become more like frame-based networks.

Detecting Loops in Cell-Mode MPLS Networks

This topic describes how loop detection is managed in cell-mode MPLS networks.

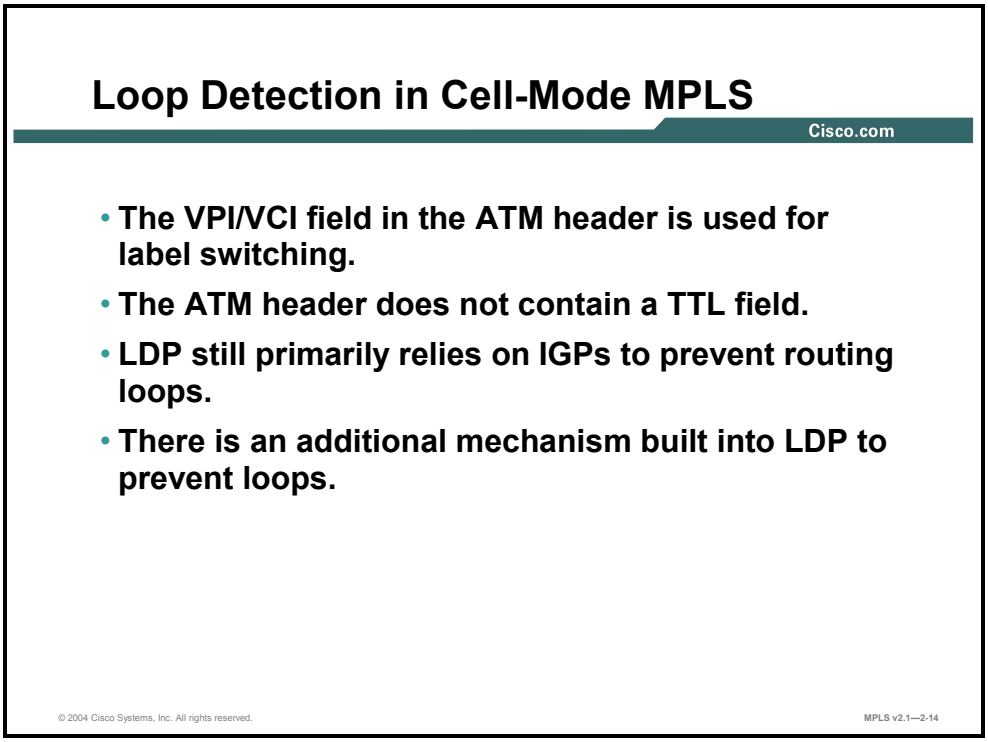

Cell-mode MPLS uses the VPI/VCI fields in the ATM header to encode labels. These two fields do not include a TTL field. Therefore, cell-mode MPLS must use other ways of preventing routing loops.

Again, most loops are prevented by the IGP, used in the network. However, if there is a loop, LDP can identify the LDP requests that were looped.

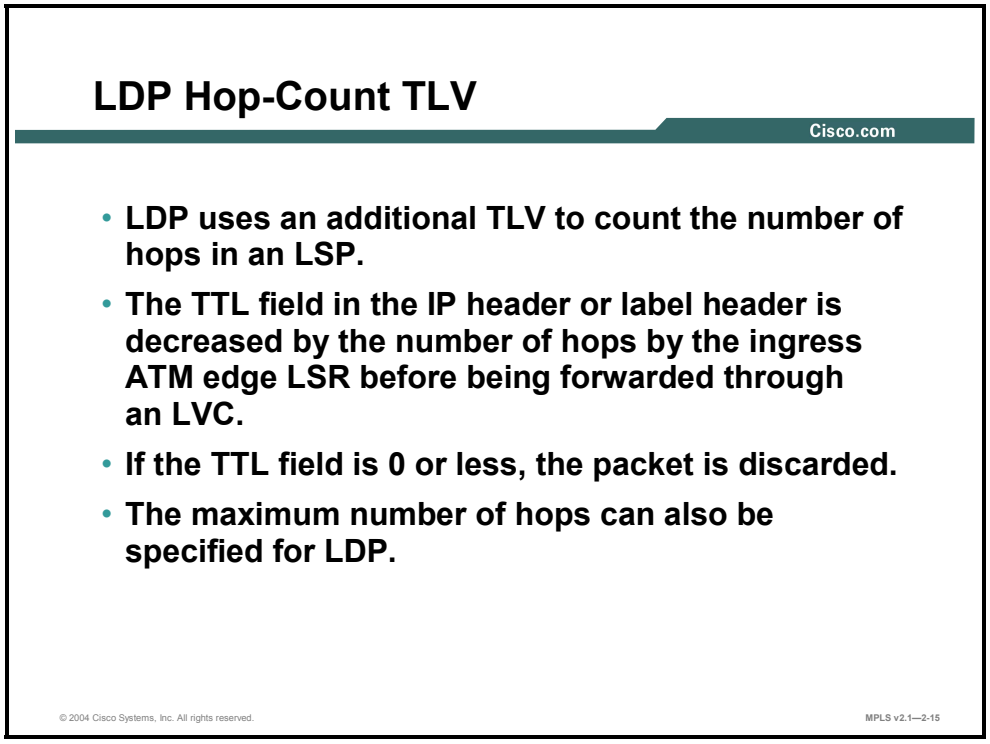

LDP uses a hop-count type, length, value (TLV) attribute to count hops in the ATM part of the MPLS domain.

This hop count can be used to provide correct TTL handling on ATM edge LSRs on behalf of ATM LSRs that cannot process IP packets.

A maximum limit in the number of hops can also be set.

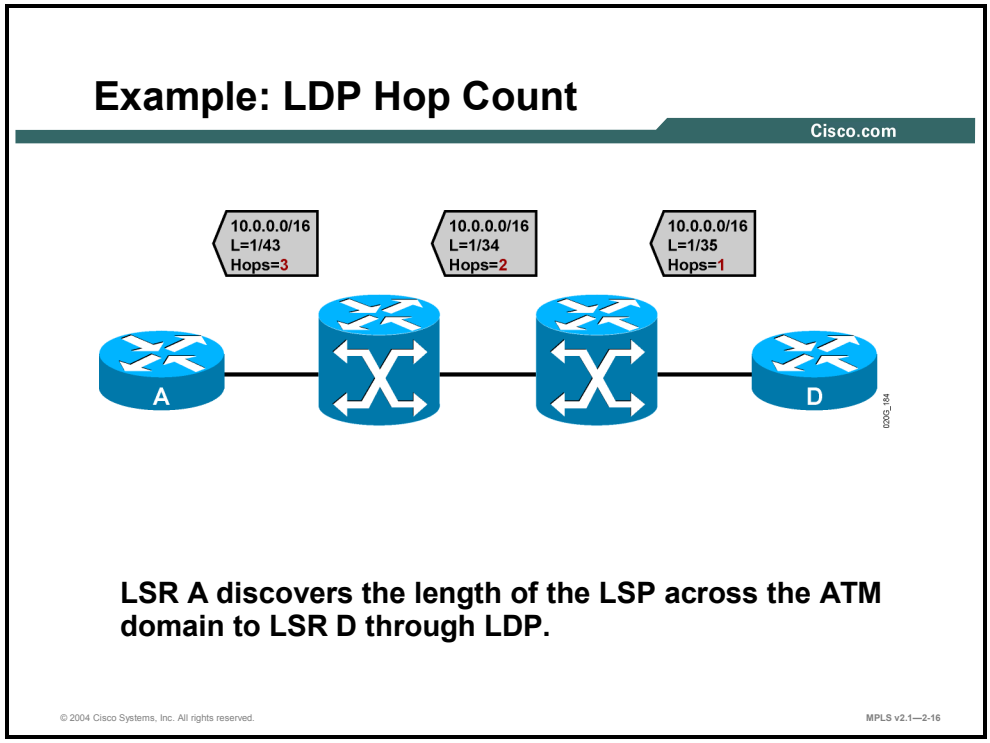

Example: LDP Hop Count

The figure illustrates how LDP, in addition to propagating the IP prefix-to-label mapping, counts hops across an MPLS-enabled ATM network.

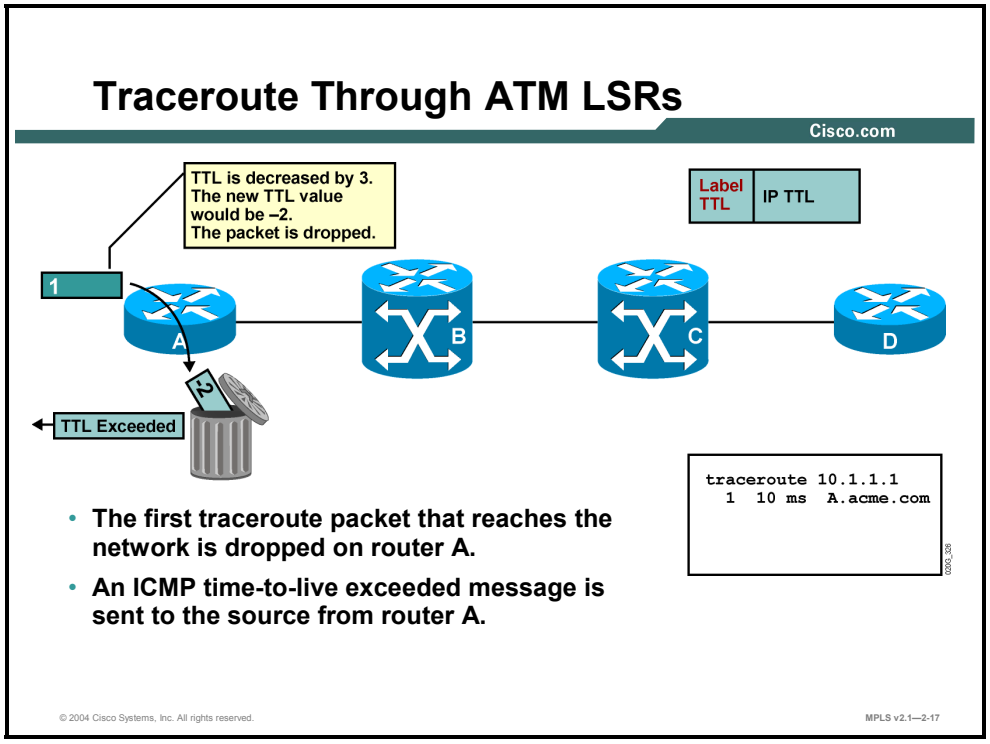

Example: Traceroute Through ATM LSRs

The figures of Traceroute Through ATM LSRs illustrate how traceroute works across an IP-aware ATM network that is not capable of using the TTL field and generating ICMP replies.

This figure illustrates how an edge ATM LSR subtracts the hop-count value instead of simply decreasing the TTL value.

The first packet results in a TTL value of –2 (less than or equal to 0), and the packet is dropped. An ICMP reply is sent to the source.

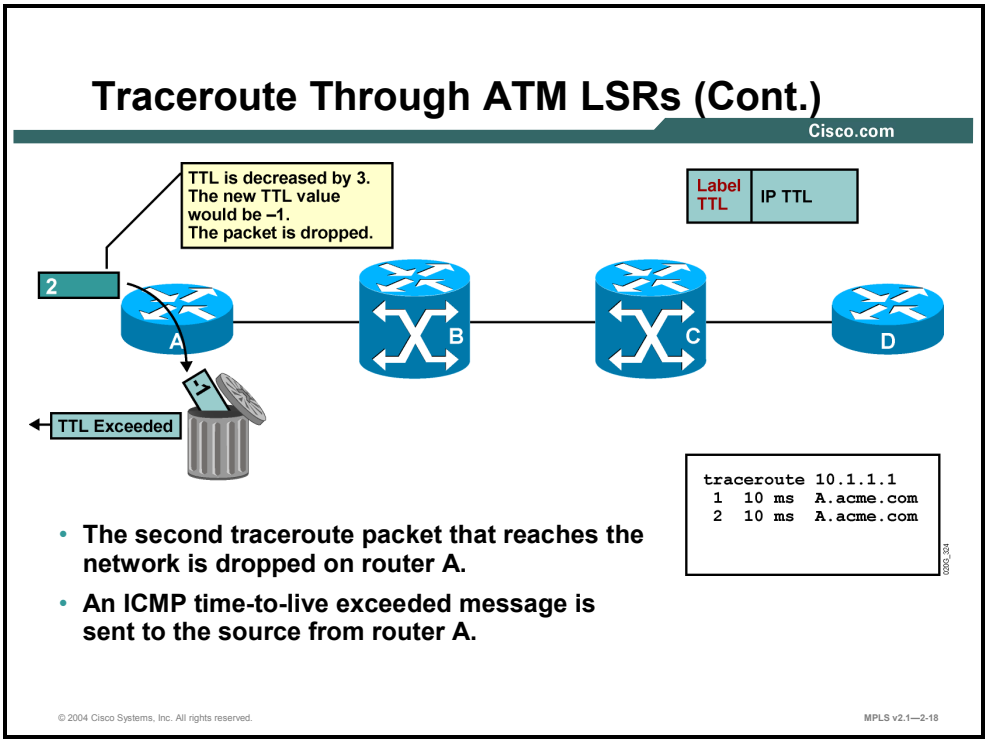

The second packet is also dropped, and another ICMP reply is sent from router A on behalf of ATM switch B, which cannot identify the TTL field and send ICMP replies itself.

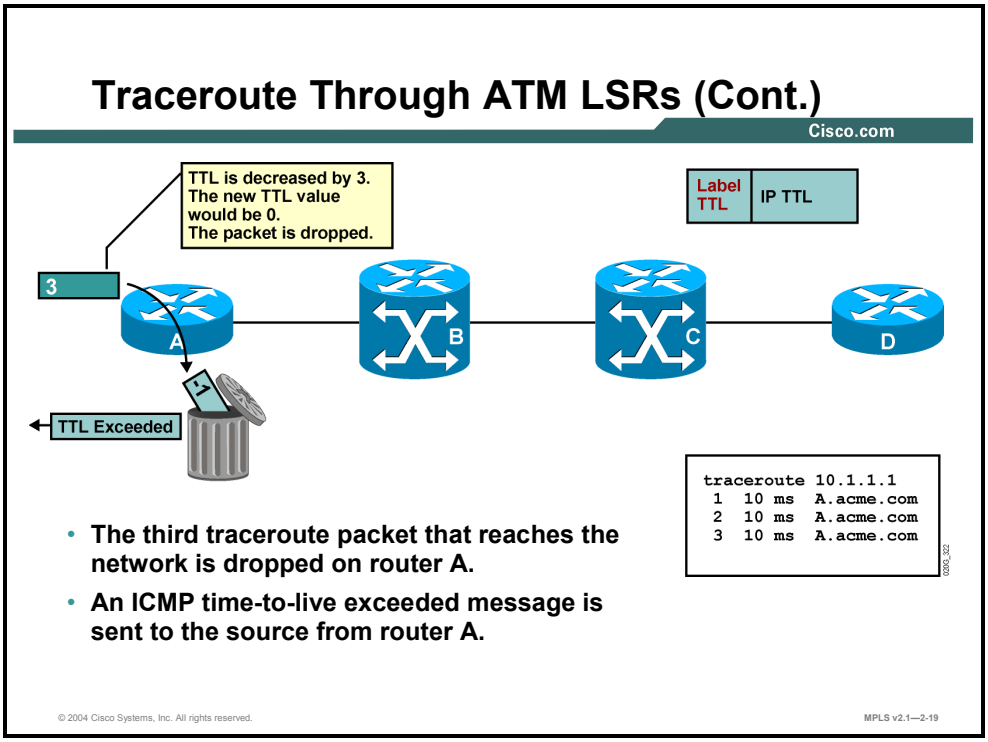

The third packet is also dropped, and the third ICMP reply is sent from router A on behalf of the ATM switch C.

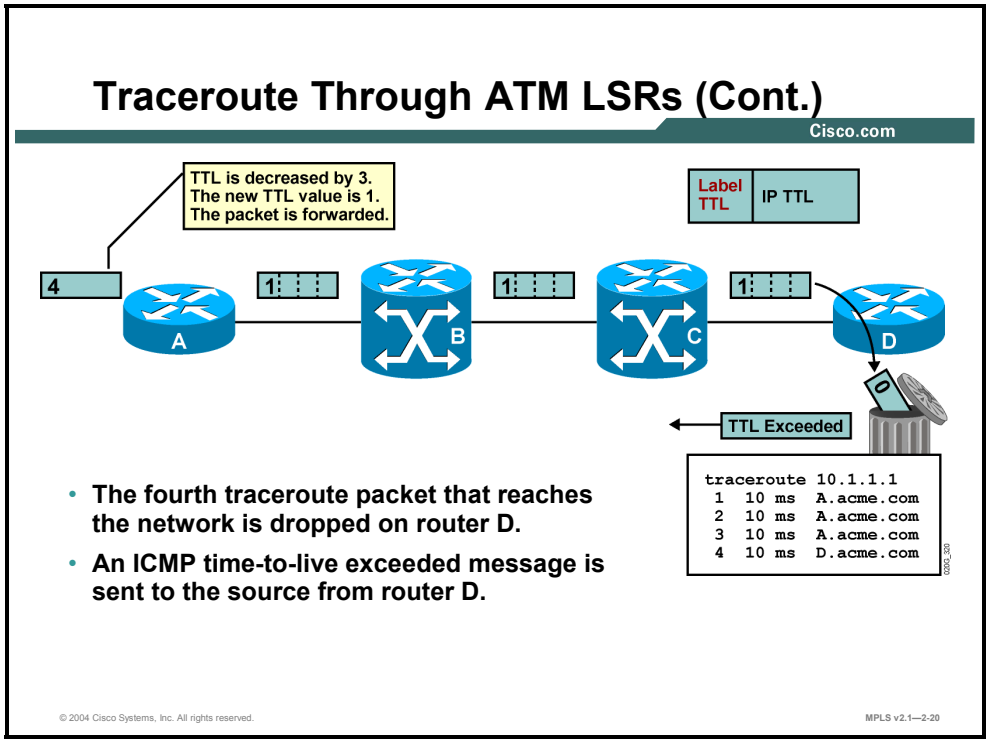

The fourth packet can reach the other edge ATM LSR (a router), which is capable of identifying the TTL field and sending ICMP replies.

The traceroute application receives as many replies as there are hops in the network, even though there are two devices in the path that are not capable of identifying the TTL field.

What Is Per-Interface Label Allocation?

This topic describes the characteristics of per-interface label allocation.

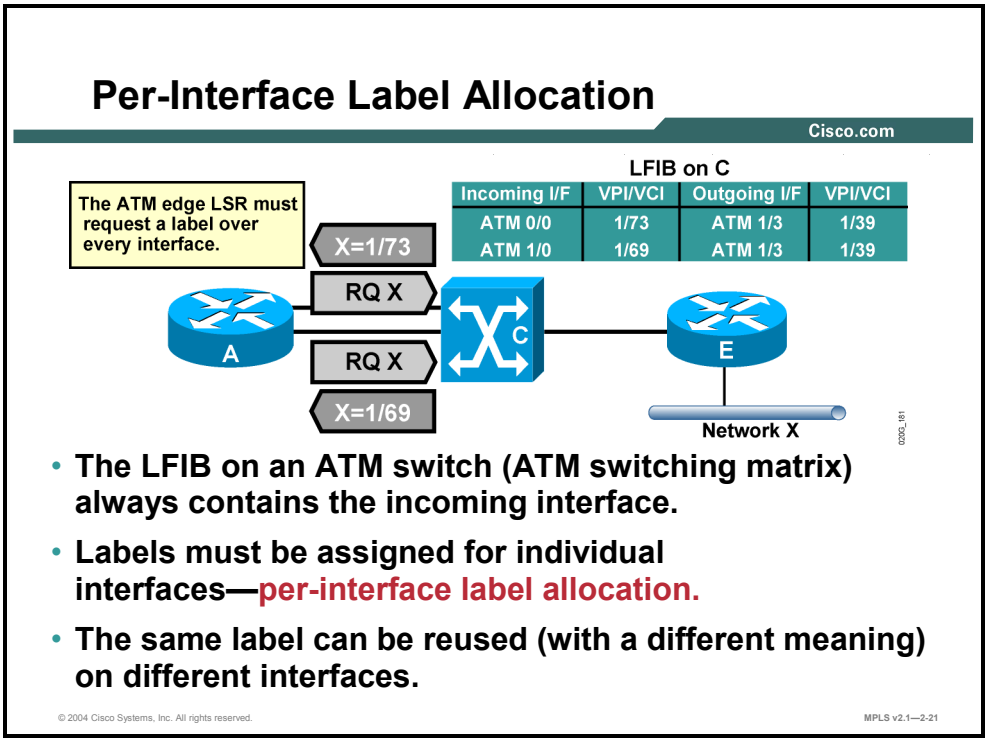

Cell-mode MPLS defaults to using per-interface label space because ATM switches support per-interface VPI/VCI values to encode labels.

Therefore, if a single router has two parallel links to the same ATM switch, two LDP sessions are established and two separate labels are requested.

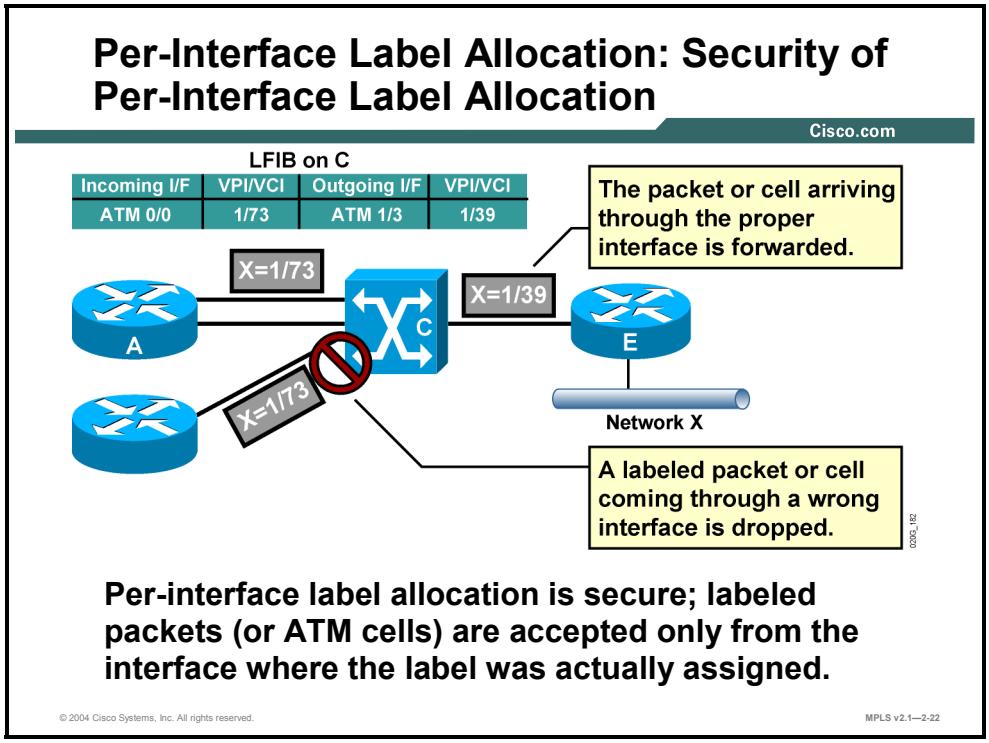

One benefit of per-interface label space is that it prevents label spoofing. In the figure, for example, the bottom router has tried to send a cell with a label that was advertised only to router A. The switch has failed to forward the cell because the cell came in through the wrong interface.

The two main forwarding differences between frame-mode and cell-mode MPLS are as follows:

- Frame-mode MPLS forwards packets based solely on labels.
- Cell-mode MPLS forwards cells based on the incoming interface and the label (VPI/VCI field).

Summary

This topic summarizes the key points discussed in this lesson.

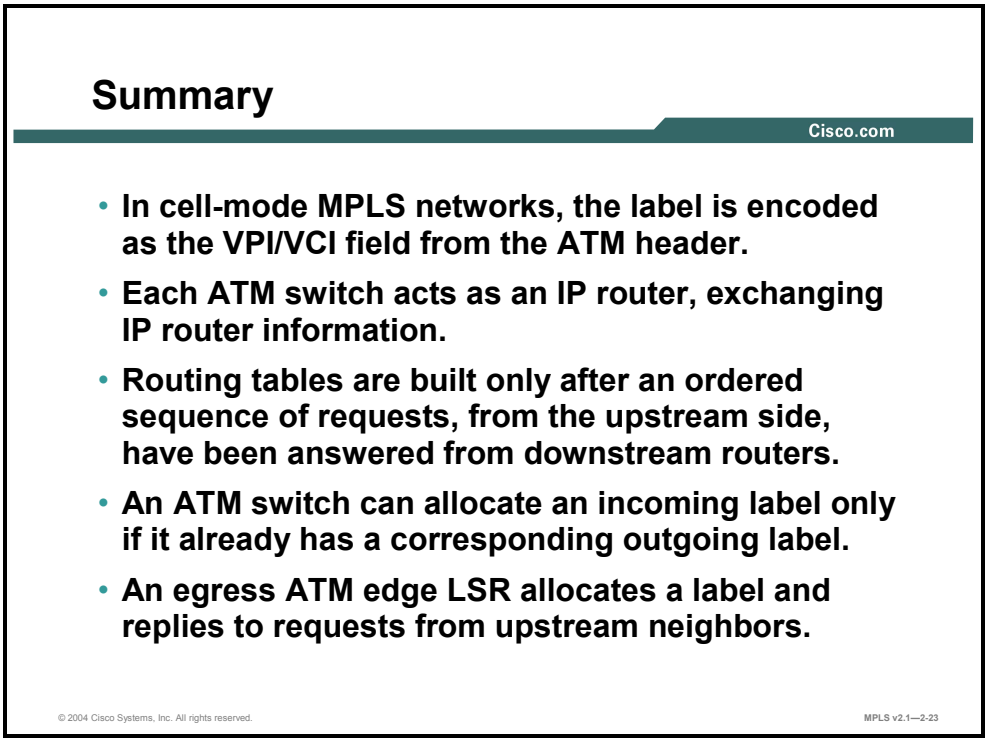

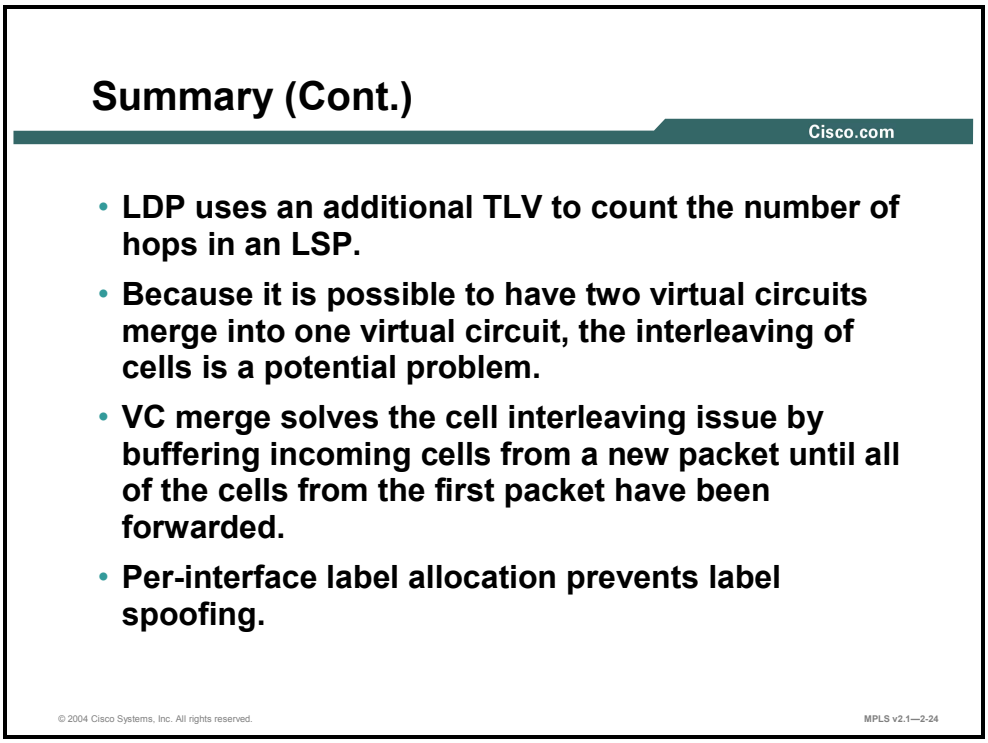

Lesson 4

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 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 features of label space
- Describe the two ways in which labels are distributed to neighbors
- Describe the two ways in which labels are allocated to neighbors
- Describe the two ways in which labels are retained
- Describe the default parameters of Cisco routers when MPLS is implemented

What Are Label Distribution Parameters?

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

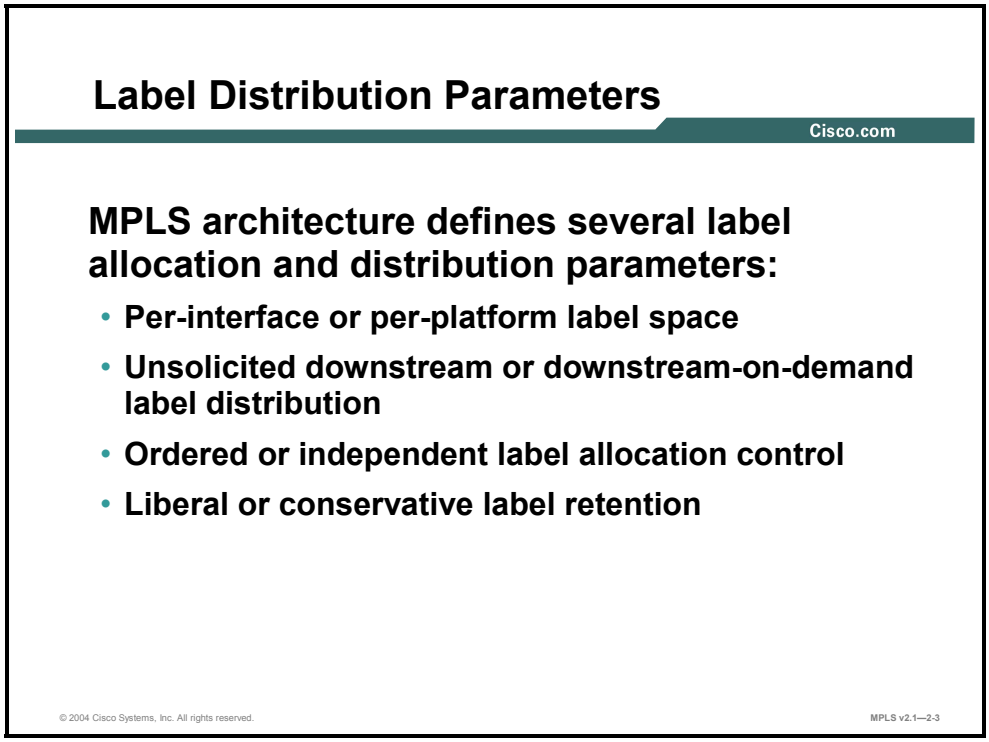

The two label space options are as follows:

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

The two options for label generation and distribution are as follows:

- Unsolicited downstream distribution of labels is used in frame-mode MPLS, where all routers can asynchronously generate local labels and propagate those labels to adjacent routers.
- Downstream-on-demand distribution of labels is used in cell-mode MPLS, where ATM LSRs have to request a label for destinations found in the IP routing table.

Another aspect of label distribution focuses on how labels are allocated, as listed here:

- Frame-mode MPLS uses independent control mode, where all routers can start propagating labels independently of one another.
- Cell-mode MPLS requires LSRs to already have the next-hop label if the LSRs are to generate and propagate their own local labels. This option is called ordered control mode.

The last aspect of label distribution looks at labels that are received but not used, as listed here:

- Frame-mode MPLS may result in multiple labels being received but only one being used. Unused labels are kept, and this mode is usually referred to as liberal label retention mode.
- Cell-mode MPLS only keeps labels that it previously requested. This mode is called conservative label retention mode.

What Is Label Space?

This topic describes the features of label space.

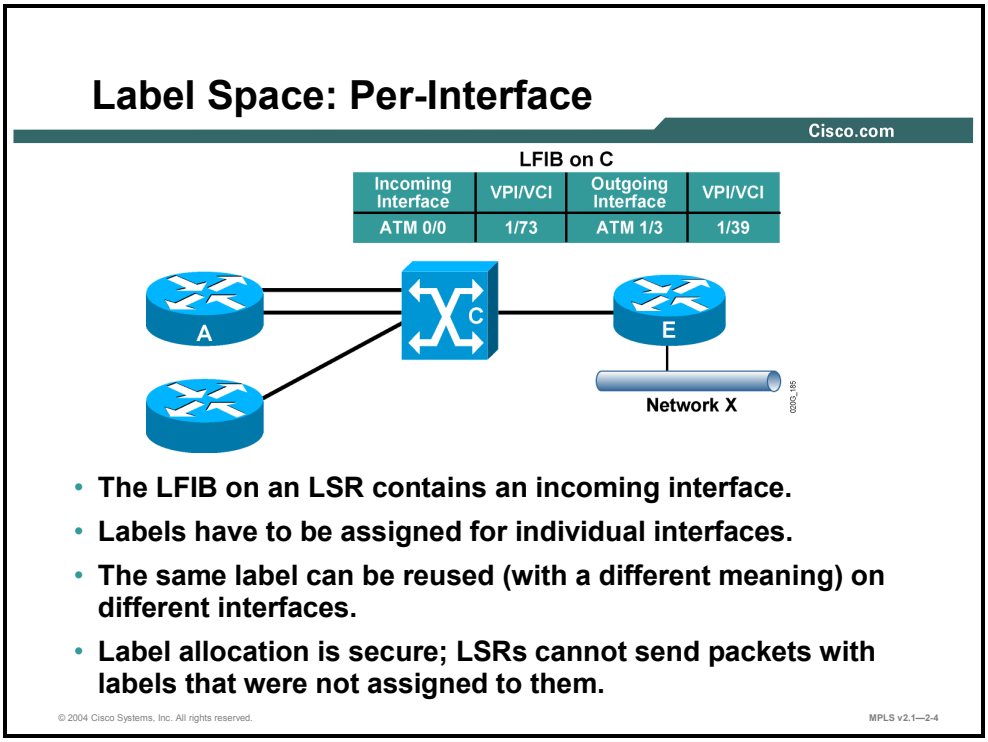

The LFIB table used with cell-mode MPLS maps a local label bound to an input interface to a next-hop label pointing to the outgoing interface. The label assigned to an input interface can be reused on another interface, and it can have a different meaning (assigned to a different destination).

Per-interface label space prevents label spoofing by not allowing cell forwarding for labels (VPI/VCI values) that are not bound to the interface where the cell was received.

Label Space: Per-Platform

Cisco.com

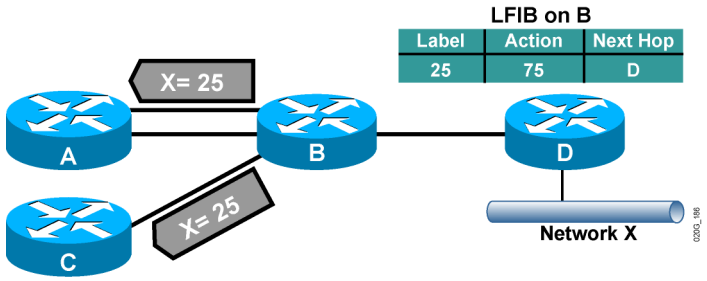

- **The LFIB on an LSR does not contain an incoming interface.**
- **The same label can be used on any interface and is announced to all adjacent LSRs.**
- **The label is announced to adjacent LSRs only once and can be used on any link.**
- **Per-platform label space is less secure than per-interface label space.**

MPLS v2.1—2-5

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 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.

Distributing Labels

This topic describes the two ways in which labels are distributed to neighbors.

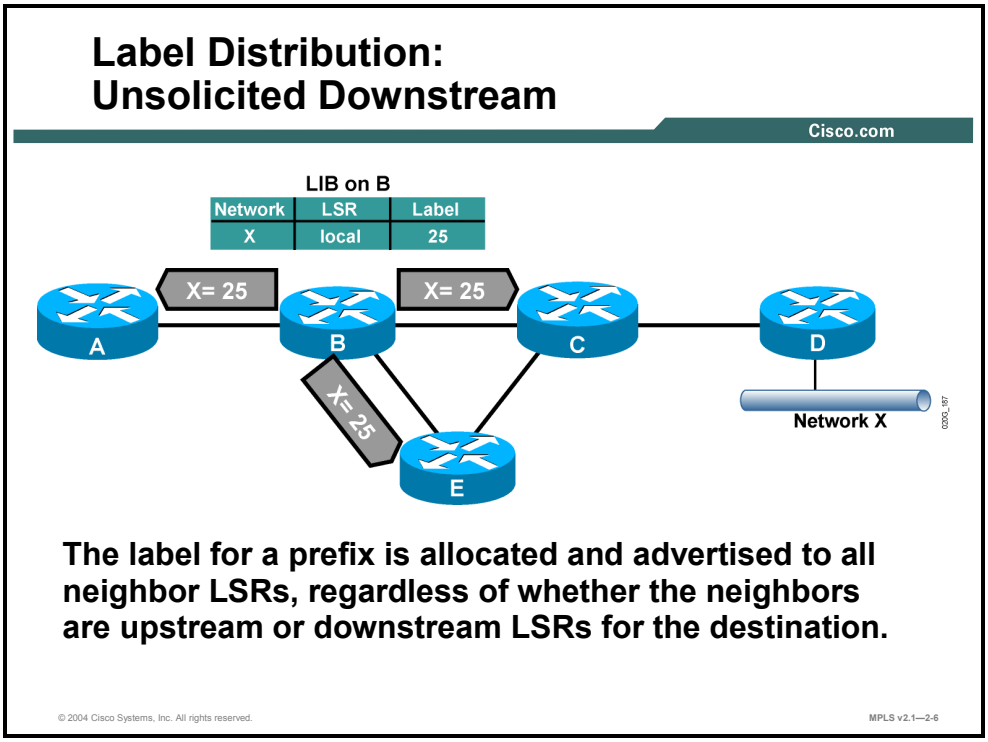

Unsolicited downstream distribution of labels is a method where each router independently assigns a label to each destination IP prefix in its routing table. This mapping is stored in the LIB table, which sends it to all LDP peers. There is no control mechanism to govern the propagation of labels in 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 IGP has put network X into the main routing table.

Each neighbor then decides upon one of the following 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

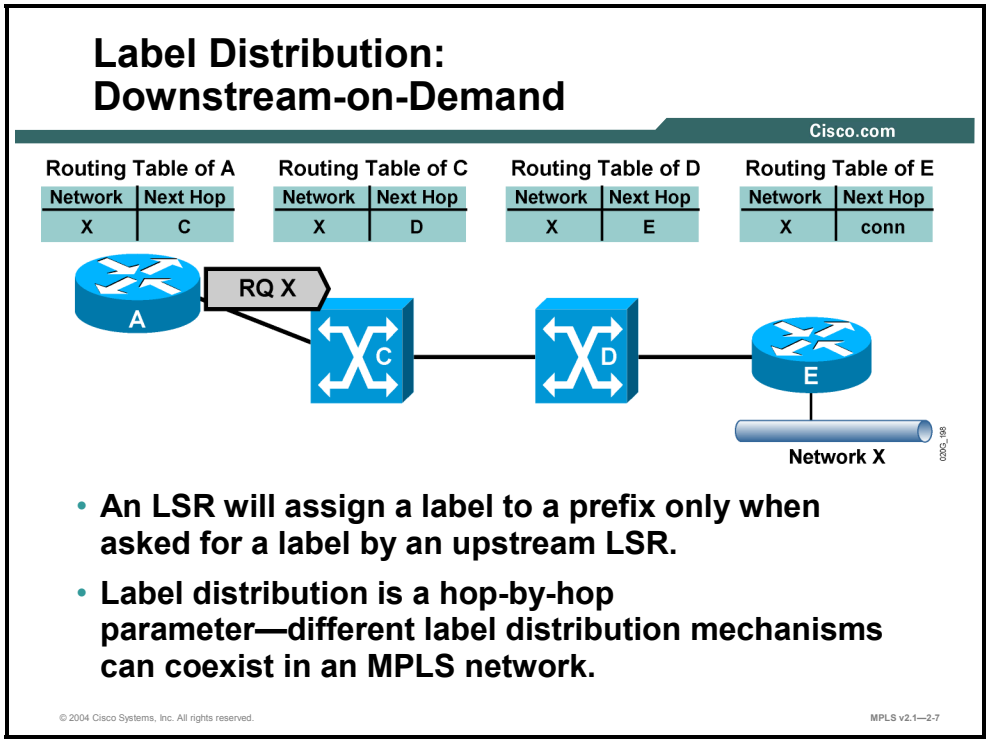

Downstream-on-demand distribution of labels requires each LSR to specifically request a label from its downstream neighbor. The figure shows how router A requests a next-hop label from its downstream LDP peer.

Unsolicited downstream and downstream-on-demand label distribution can be combined because labels are assigned and propagated hop by hop. The usual situation is that frame-mode MPLS uses unsolicited downstream label propagation, and cell-mode MPLS uses downstream-on-demand label propagation.

Allocating Labels

This topic describes the two ways in which labels are allocated to neighbors.

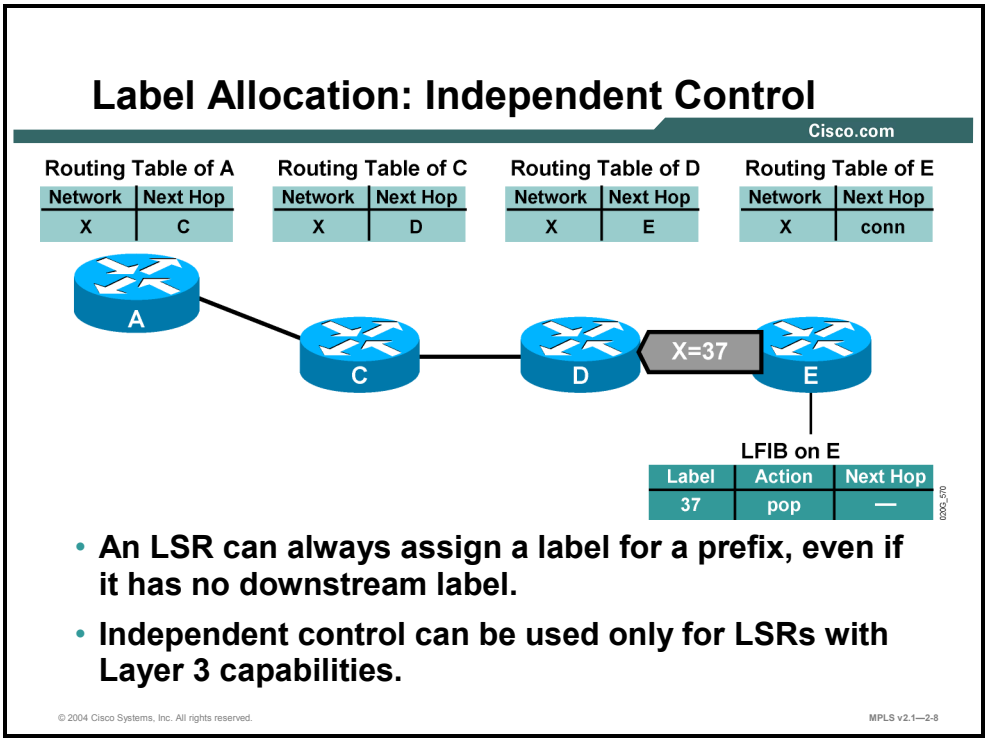

Independent control mode is usually combined with unsolicited downstream propagation of labels, where labels can be created and propagated independently of any other LSR. When independent control mode is used, an LSR might be faced with an incoming labeled packet where there is no corresponding outgoing label in the 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.

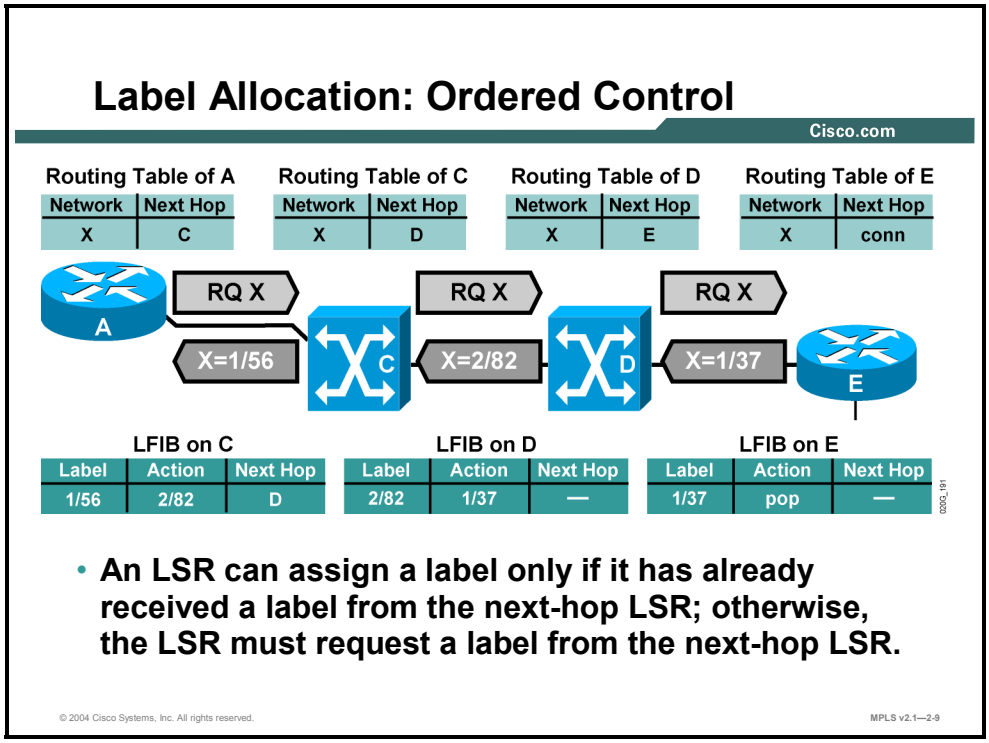

Ordered control mode is usually combined with downstream-on-demand propagation of labels, where a local label can be assigned and propagated only if a next-hop label is available. This requirement results in an ordered sequence of downstream requests until an LSR is found that already has a next-hop label or an LSR is reached that uses independent control mode.

Although ordered control mode could be used with frame-mode MPLS, its use is mandatory on ATM switches, which cannot perform Layer 3 lookups.

Example: Ordered Control

The figure illustrates how both ATM LSRs forward requests until an edge is reached. The edge LSR uses independent control mode and can respond to the request.

Retaining Labels

This topic describes the two ways in which labels are retained.

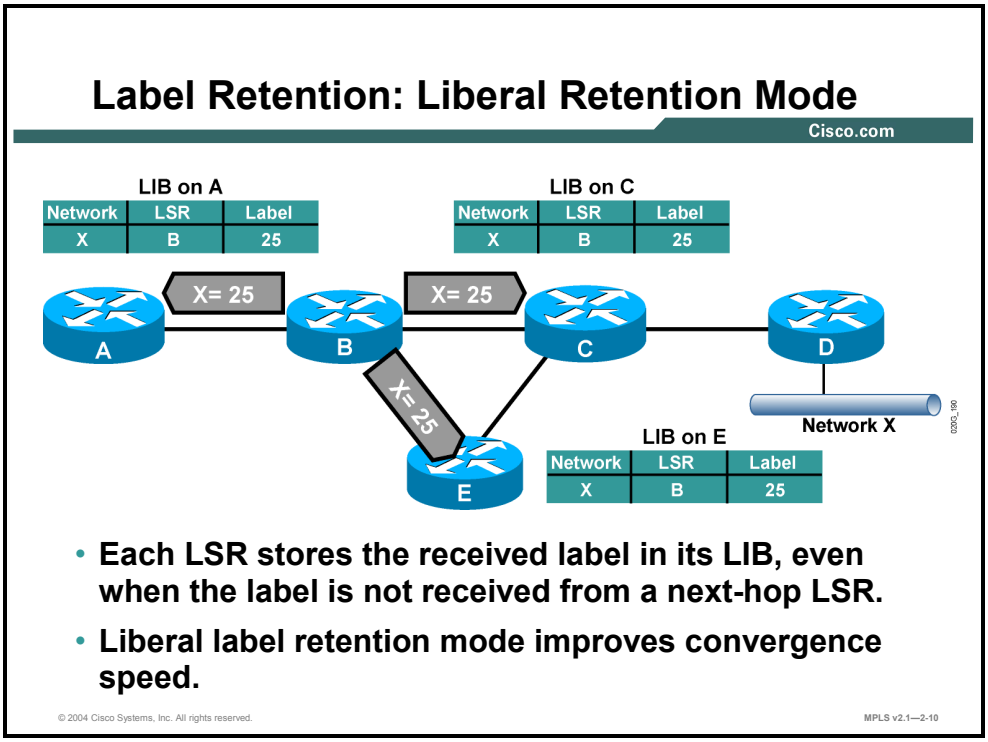

Liberal label retention mode dictates that each LSR keeps all labels received from LDP peers, even if they are not the downstream peers for network X.

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.

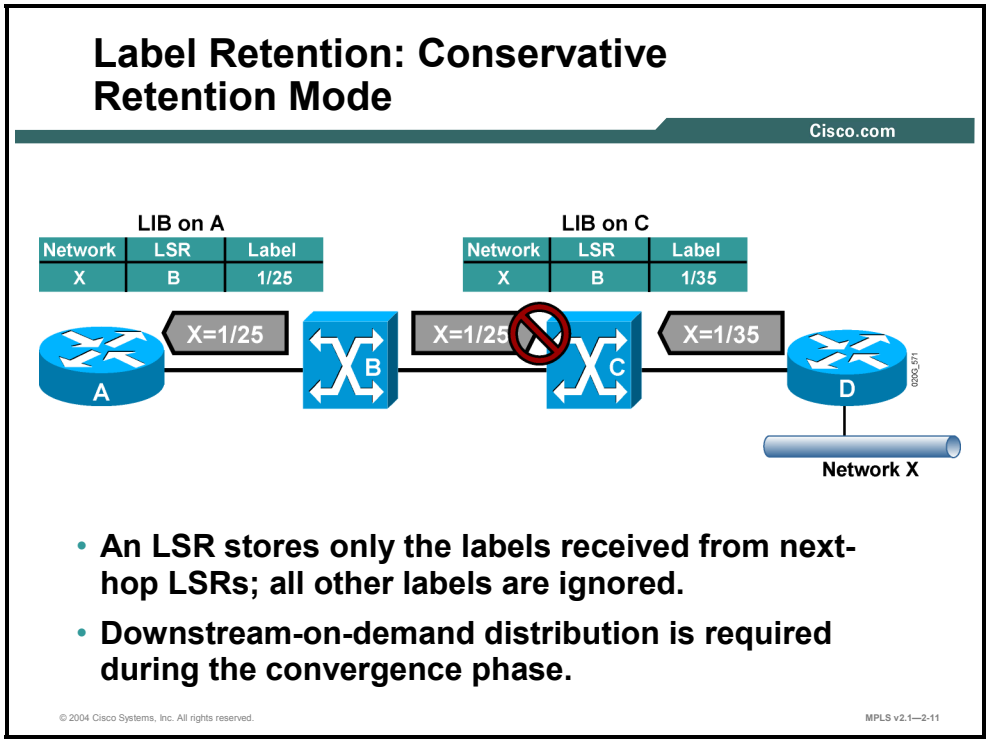

Conservative label retention mode keeps only labels that can immediately be used for normally routed traffic paths. Conservative label retention, downstream-on-demand, and ordered control mode help conserve label resources that affect usage of limited ATM virtual circuits.

Example: Conservative Retention Mode

The figure illustrates how ATM switch C does not consider switch B to be the next hop for network X and, therefore, drops the labels received from router B.

Note	Conservative label retention mode requires downstream-on-demand label allocation after network convergence.

What Are Standard Parameter Sets in MPLS Implementation?

This topic describes the default parameters of Cisco routers when MPLS is implemented.

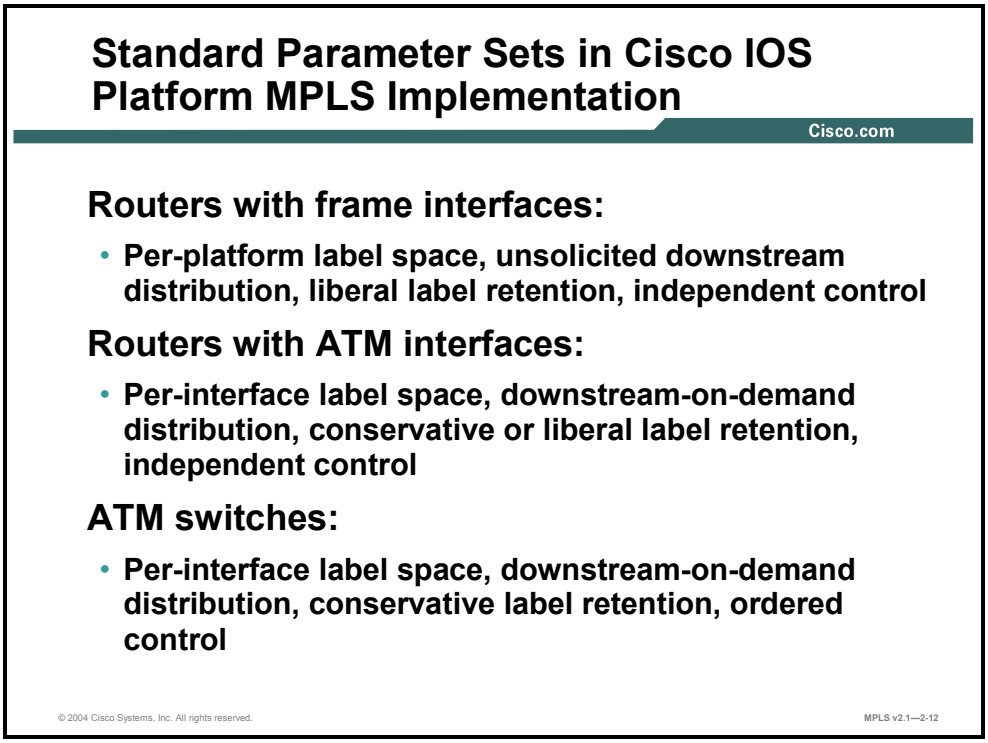

The following default operation applies to routers using frame-mode MPLS (LSRs):

- **Per-platform label space:** Platform-wide incoming labels are used for interfaces. Interfaces can share the same labels.
- **Unsolicited downstream propagation of labels:** Every LSR can propagate a label mapping to its neighbors without a request.
- **Liberal label retention mode:** This mode allows for easy failover if a link fails.
- **Independent control mode:** This mode makes label propagation faster (less time needed for LDP convergence), because LSRs do not have to wait to get the next-hop label from their downstream neighbors.

The following default operation applies to ATM switching using cell-mode MPLS (ATM LSRs):

- **Per-interface label space:** Per-interface label space provides better security and is already available with standard ATM switching functionality.
- **Downstream-on-demand propagation of labels:** LFIB tables on ATM switches are really ATM switching matrices that require full information before switching can start; full information includes next-hop label, which must be requested.
- **Conservative label retention mode:** This mode is implicitly achieved by using the downstream-on-demand propagation of labels; no label is received unless it is requested.
- **Ordered control mode:** This mode is used in combination with downstream-on-demand propagation of labels to ensure that every ATM LSR has all of the information needed to create an entry in the LFIB table (ATM switching matrix), including the next-hop label.

The default operation of routers using cell-mode MPLS (ATM edge LSRs) is similar to that of ATM switches. The exception is that routers use independent control mode because they are the endpoints of the virtual circuits.

Summary

This topic summarizes the key points discussed in this lesson.

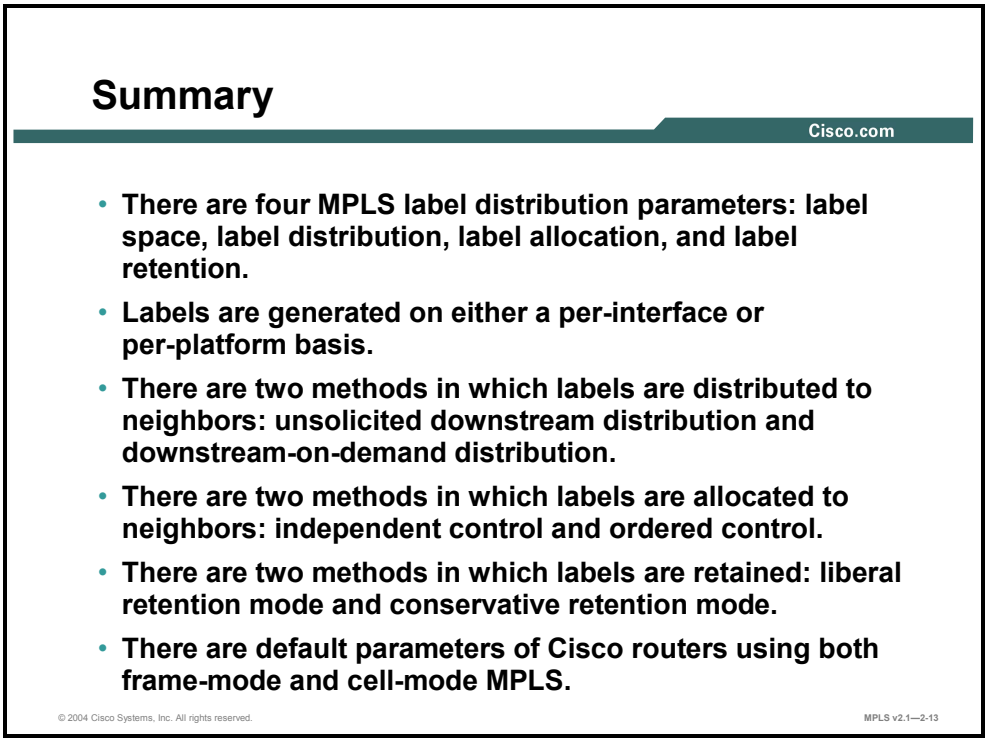

Lesson 5

Discovering LDP Neighbors

Overview

This lesson takes a more detailed look at the 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 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 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 ATM LSRs
- Describe how LDP discovers nonadjacent neighbors

Establishing an LDP Session

This topic describes how LDP sessions are established between neighbors.

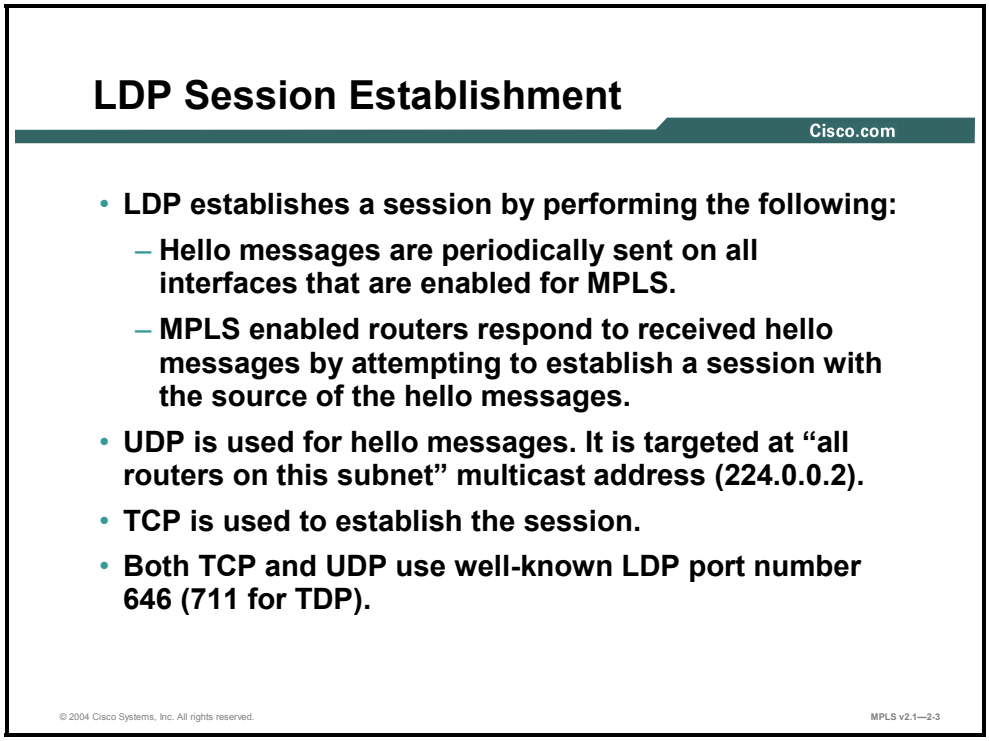

LDP is a standard protocol used to exchange labels between adjacent routers. TDP is a Cisco proprietary protocol that has the same functionality as LDP.

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. The hello messages use UDP packets with a multicast destination address of 224.0.0.2 ("all routers on a subnet") and destination port number of 646 (711 for TDP).

If another router is enabled for LDP (or TDP), it will respond by opening a TCP session with the same destination port number (646 or 711).

What Are LDP Hello Messages?

This topic describes the contents of an LDP hello message.

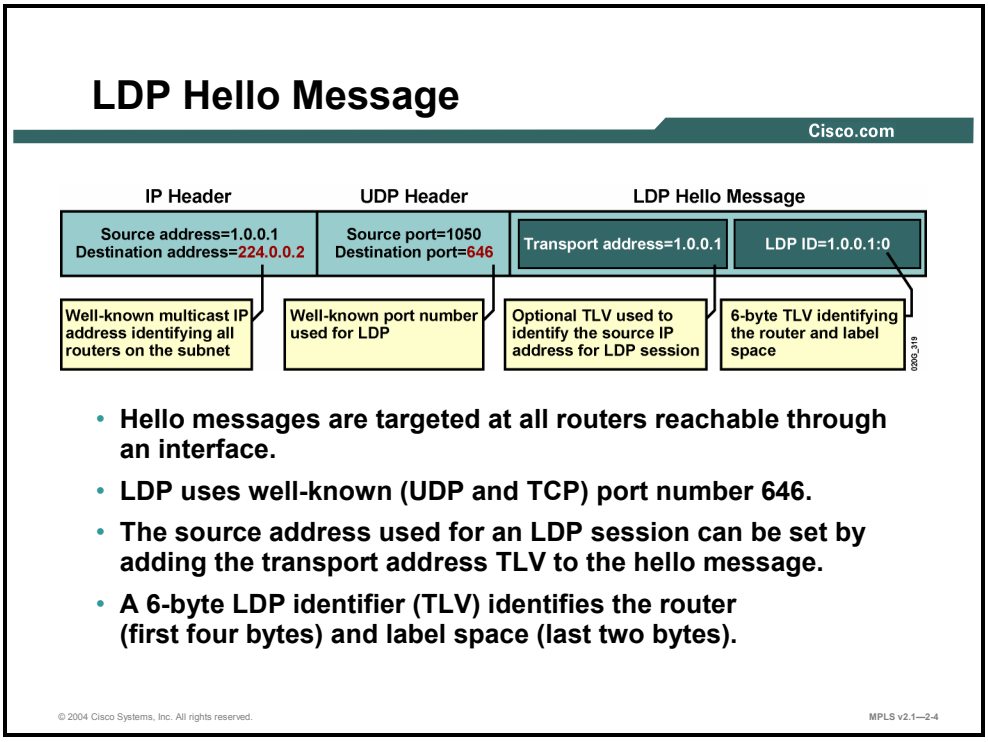

The contents of a hello message are as follows:

- Destination IP address (224.0.0.2), which targets all routes 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 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 is used to uniquely identify the neighbor and the label space; multiple sessions can be established between a pair of LSRs if they use multiple label spaces.

Negotiating Label Space

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

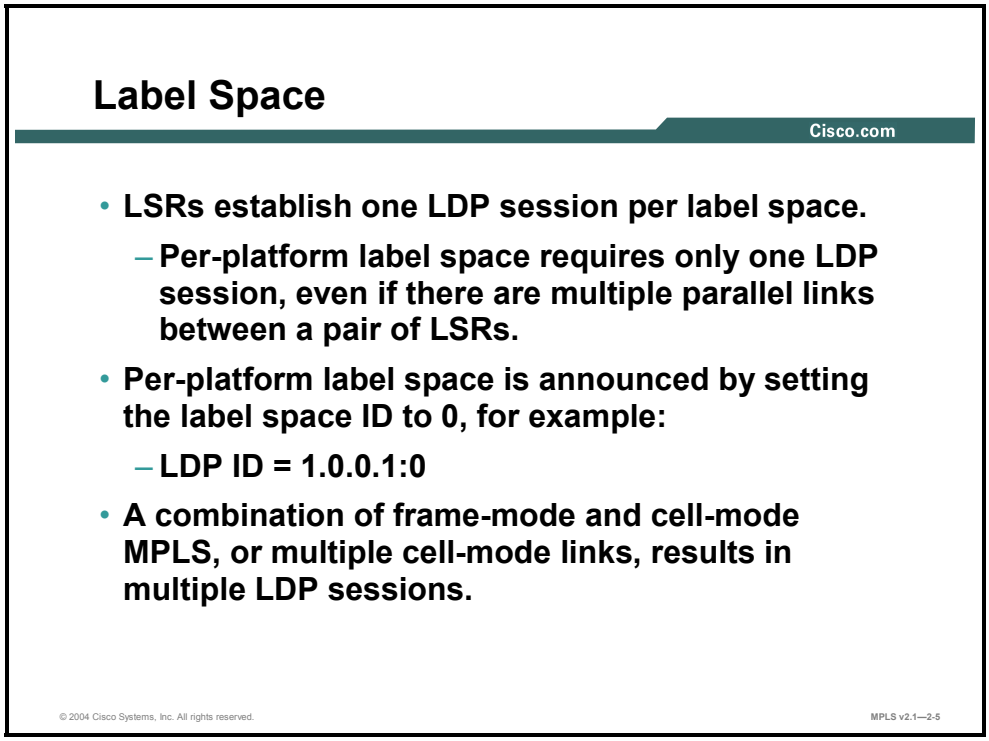

If a pair of routers is connected over two or more parallel links and use frame-mode MPLS, the routers try to establish multiple sessions by using the same LDP identifier. Because the routers are using per-platform label space, this action will result in only one session remaining; the other session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

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

If the two routers use different LDP identifiers (for example, if one link uses frame-mode MPLS and the other uses cell-mode MPLS), they will keep both sessions.

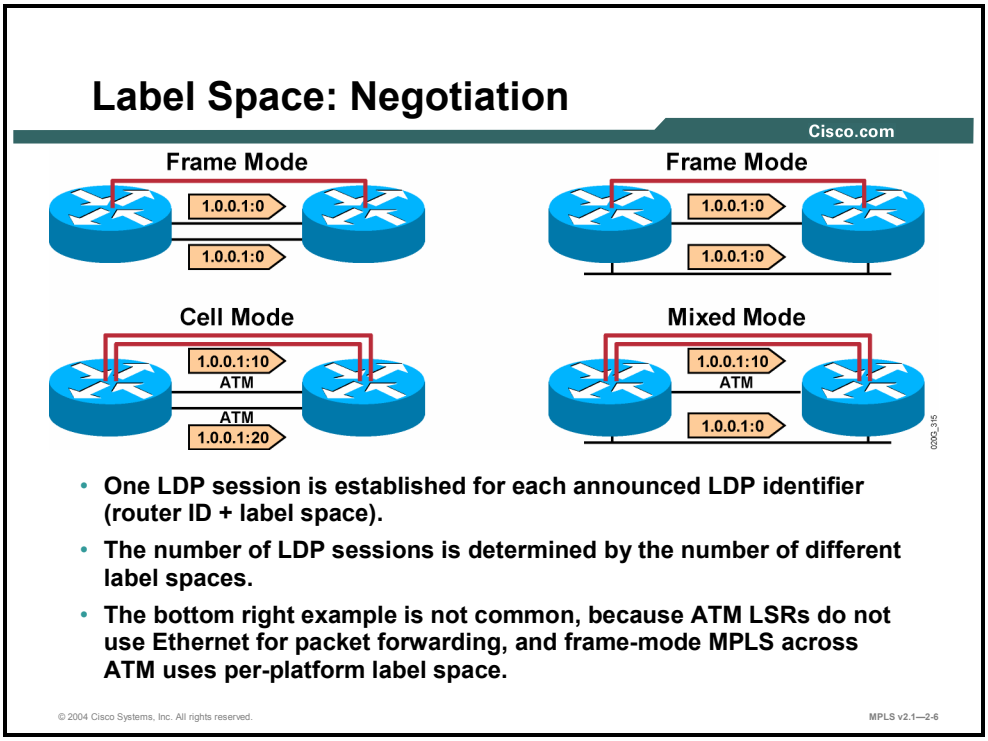

Example: Label Space Negotiation

The figure illustrates four different combinations with two parallel links between a pair of routers. The top routers are frame-mode routers.

A general rule can be extracted from the four examples: An LDP session is established per interface except for all frame-mode interfaces, where only one LDP session between a pair of LSRs is used because frame-mode MPLS uses per-platform label space.

Note	The bottom right example is not common, because ATM LSRs do not use Ethernet for packet forwarding, and frame-mode MPLS across ATM uses per-platform label space.

Discovering LDP Neighbors

This topic describes how LDP neighbors are discovered.

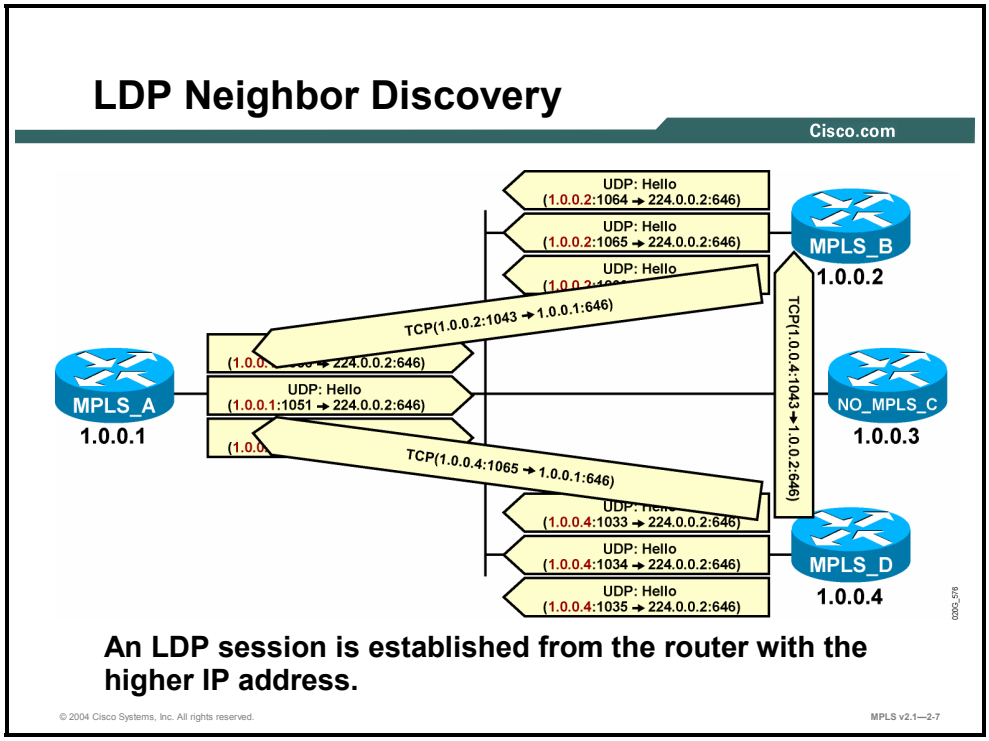

Example: LDP Neighbor Discovery

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 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.

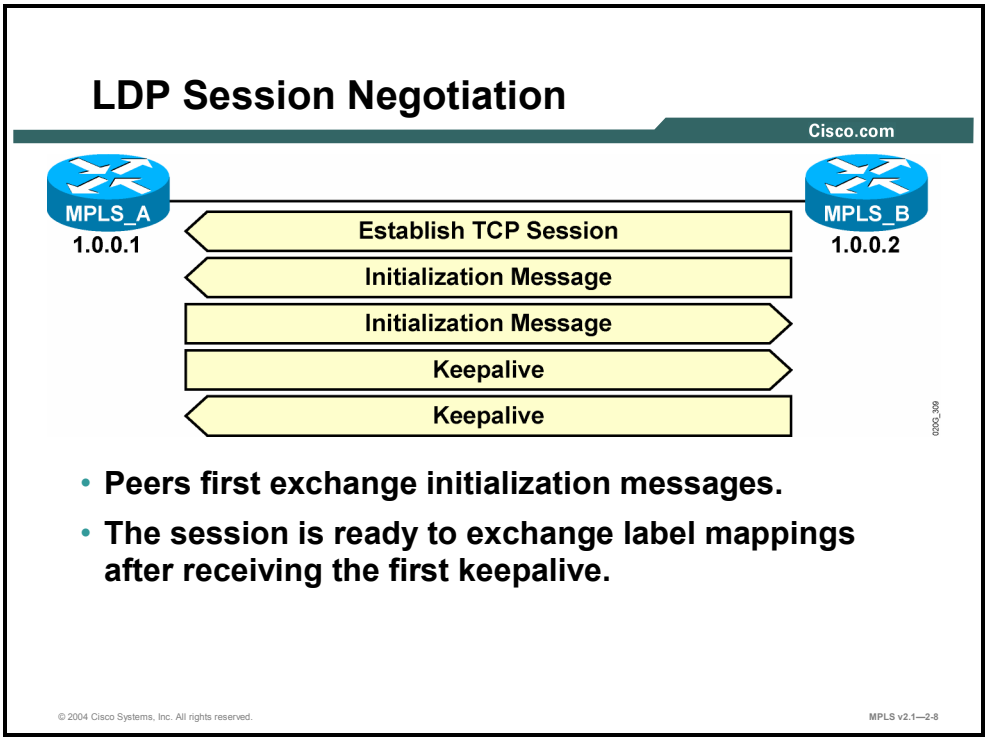

LDP session negotiation is a three-step process, as follows:

Step 1 Establish the TCP session.

Step 2 Exchange initialization messages.

Step 3 Exchange initial keepalive messages.

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

Establishing LDP Sessions Between ATM LSRs

This topic describes how LDP sessions are established between ATM LSRs.

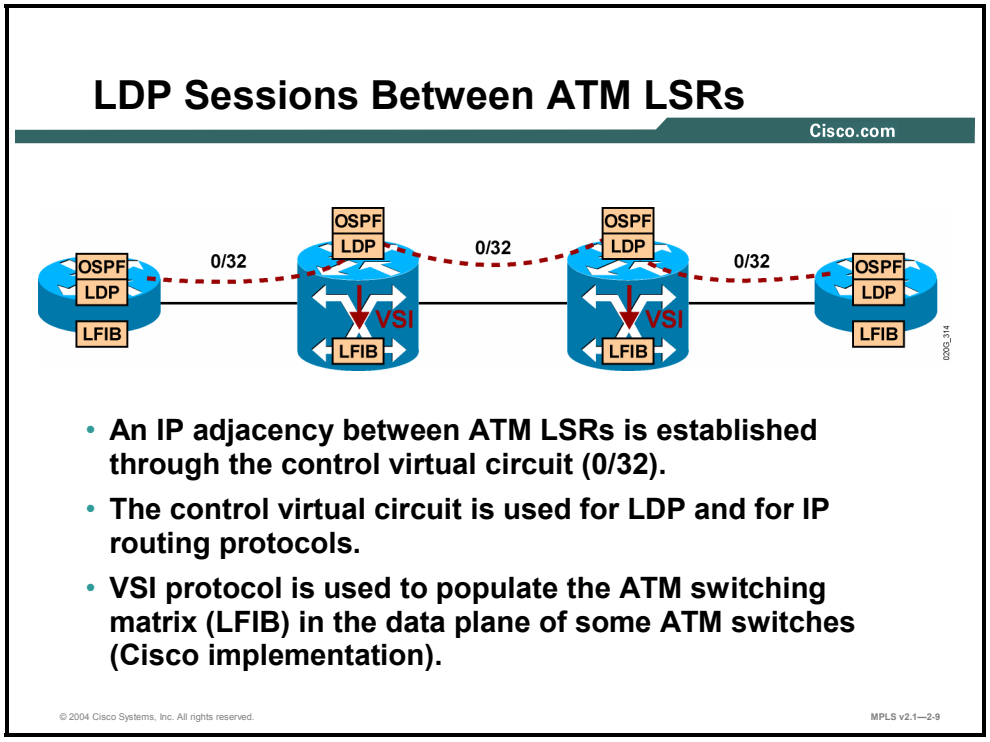

Example: LDP Sessions Between ATM LSRs

The figure illustrates the operation of LDP in ATM networks. ATM LSRs establish the IP adjacency across the MPLS control virtual circuit, which by default has a VPI/VCI value of 0/32.

An IP routing protocol and LDP (or TDP) use this control virtual circuit to exchange IP routing information and labels.

Some Cisco devices use the Virtual Switch Interface (VSI) protocol to create entries in the LFIB table (ATM switching matrix of the data plane) based on the information in the LIB table (control plane). This protocol is used to dynamically create virtual circuits for each IP network.

Discovering Nonadjacent Neighbors

This topic describes how LDP discovers nonadjacent neighbors.

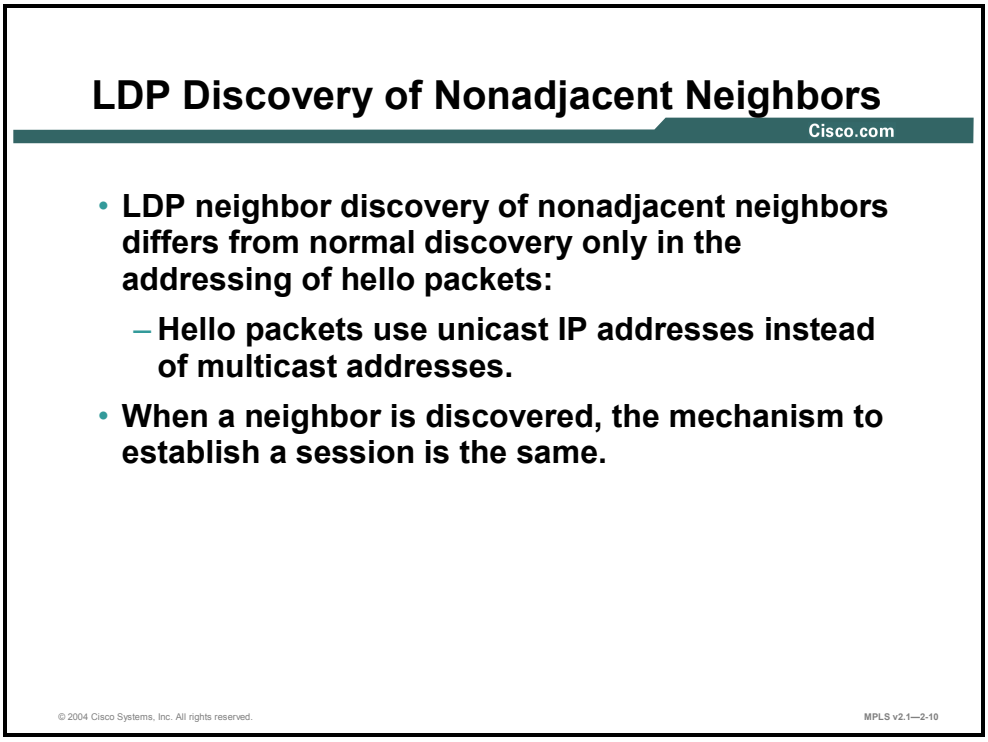

LDP can also be used between nonadjacent routers. However, LDP hello messages use unicast IP addresses instead of multicast. The rest of the session negotiation is the same as for adjacent routers.

Summary

This topic summarizes the key points discussed in this lesson.

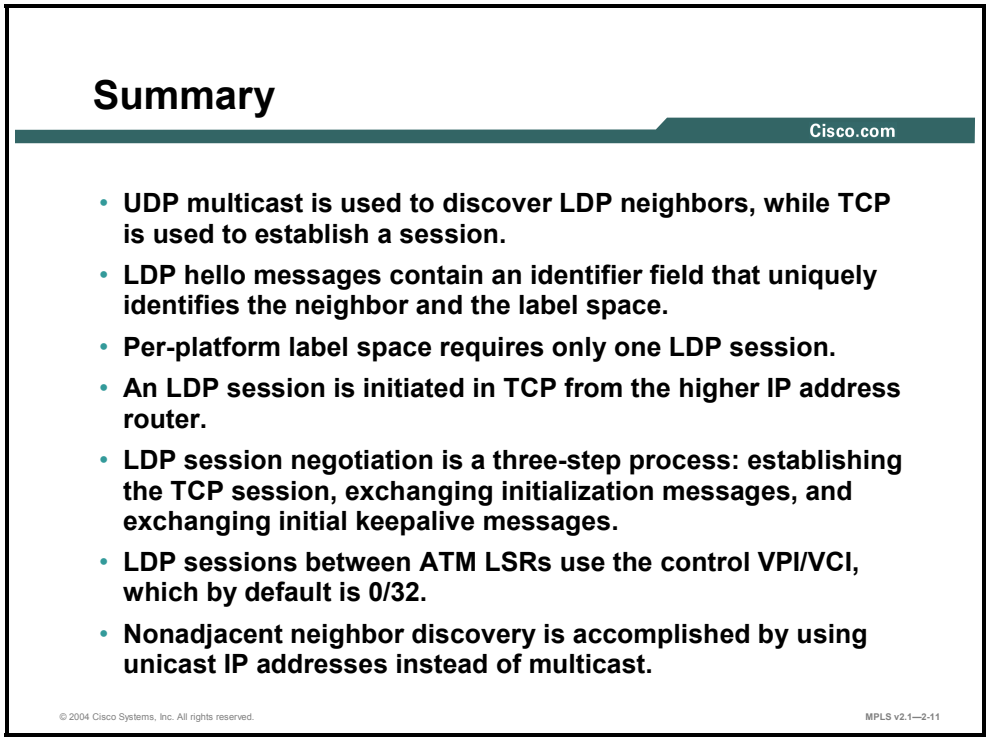

Module Summary

This topic summarizes the key points discussed in this module.

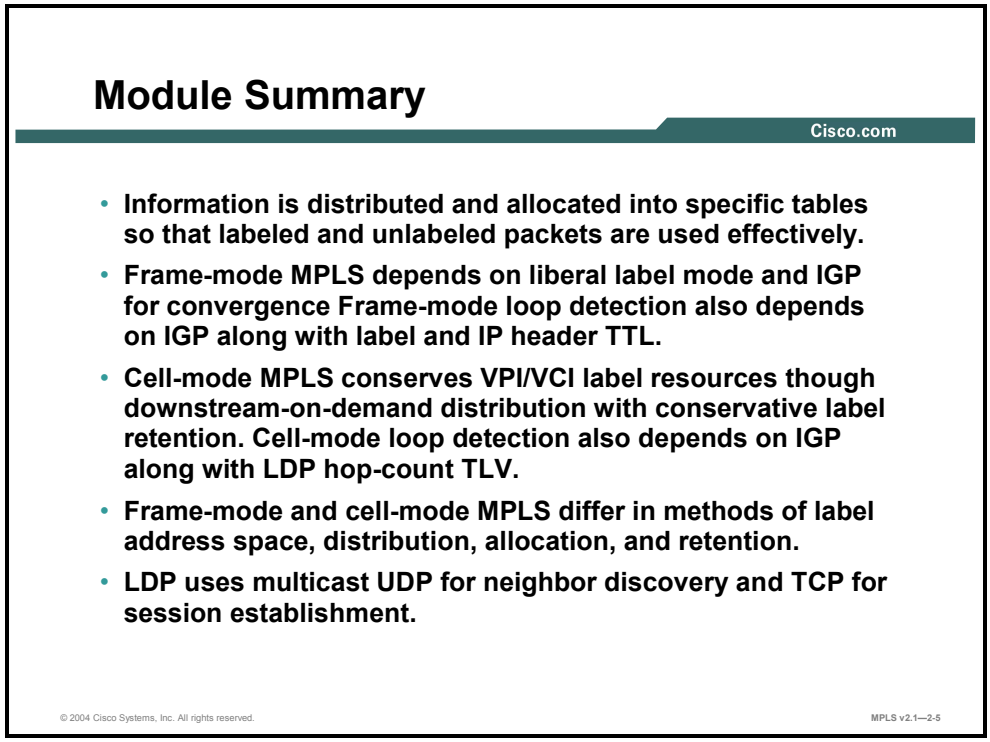

In an MPLS network, labels are assigned and distributed, involving neighbor discovery and session establishment. Label information is populated in LIB, FIB, and LFIB tables.

References

For additional information, refer to these resources:

- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3036, *LDP Specification*

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 of the following statements 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.

Q2) Which of the following descriptions 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

Q3) Which three of the following are contained in the LFIB? (Choose three.) (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

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

Q4) 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

Q5) 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.

Q6) Which two tables contain label information? (Choose two.) (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

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

Q7) Which of the following generates a label update? (Source: Introducing Typical Label Distribution in Frame-Mode MPLS)

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

Q8) 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.

Q9) 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 TTL is not copied back into the TTL of the label.
D) None of the above is correct.

Q10) Which of the following 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.

Q11) 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

Q12) 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

Q13) 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.

Q14) 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

Q15) 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.

Q16) What are two possible solutions to the interleaving of cells in cell-mode MPLS? (Choose two.) (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) Allocate a downstream label for each request.
B) There is no possibility of cells being interleaved if the correct configuration is performed on ATM switches.
C) Buffer the cells of the second packet.
D) There are no issues with the interleaving of cells.

Q17) In cell-mode MPLS networks, where are labels inserted? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) Labels are inserted between the Layer 2 header and Layer 3 header.
B) Labels are inserted in the VPI/VCI field of the ATM header.
C) Labels are not used in cell-mode MPLS networks.
D) Labels are inserted in the Layer 3 header only.

Q18) With regard to VC merge, which statement is NOT true? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) Using VC merge, ATM LSRs can reuse the same downstream label for multiple upstream LSRs.
B) ATM networks are effectively transformed into a frame-mode MPLS network.
C) Jitter and delay across the ATM network decrease.
D) Buffering requirements increase on the ATM LSR.

Q19) Which statement pertains to the IP routing table? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) The IP routing table is NOT built on ATM LSRs.
B) The IP routing table is built on the data plane of each ATM switch.
C) The IP routing table is built on the control plane of each ATM switch.
D) The IP routing table is built on the forwarding plane of each ATM switch.

Q20) Which statement pertains to the IP forwarding table? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) The IP forwarding table is built as in a frame-mode MPLS network.
B) The IP forwarding table is built only after the label requests have been answered (with labels) from upstream LSRs.
C) The IP forwarding table is built only after the label requests have been answered (with labels) from downstream LSRs.
D) There is no need for the IP forwarding table in cell-mode MPLS. Everything is done with the IP routing table.

Q21) Which statement is NOT true? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) Frame-mode MPLS forwards labels based solely on the labels.
B) Cell-mode MPLS forwards labels based on the incoming interface and VPI/VCI field (label).
C) If a router has two parallel links to the same ATM switch, one LDP session will be established, and one label will be requested.
D) Per-interface label allocation prevents label spoofing.

Q22) An ATM switch will respond to a request for a label in which situation? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) The ATM switch will respond when it knows the next-hop label.
B) The ATM switch will always reply to downstream label requests.
C) The ATM switch will always reply to upstream label requests.
D) ATM switches do not use MPLS labels.

Q23) Which of the following describes a task that ATM switches perform? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) upstream-on-demand label allocation
B) downstream-on-demand label allocation
C) unsolicited label allocation
D) Labels are not used in cell-mode MPLS networks.

Q24) Which of the following is used in cell-mode loop detection? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) the TTL field of the IP packet
B) the TTL field in the MPLS label
C) a TLV that counts the number of hops
D) a TLV that counts the number of packets

Q25) Which table holds all labels assigned by an LSR and their mapping to labels that have been received from the neighbors of the LSR? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

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

Q26) What does the term “pop” mean when you are describing penultimate hop popping? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) swap the top label with a new label contained in the LIB
B) swap the top label with a new label contained in the LFIB
C) remove the top label instead of swapping it with the next-hop label
D) remove the bottom label instead of swapping it with the next-hop label

Q27) A solution for cell interleaving that could occur in ATM MPLS networks is which of the following? (Source: Introducing Typical Label Distribution Over LC-ATM Interfaces and VC Merge)

A) PHS
B) VC merge
C) PC merge
D) PSS

Q28) 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

Q29) Cell-mode MPLS uses _____ label space, and frame-mode uses _____ label space. (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

Q30) Which two types of label distribution are used in Cisco MPLS networks? (Choose two.) (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

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

Q31) The modes of label allocation are _____ control and _____ control. (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

Q32) What are the two label retention modes used in Cisco MPLS networks? (Choose two.) (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

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

Q33) Which statement is correct? (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) By default, ATM switches use independent control.
B) By default, ATM switches use per-platform label space.
C) By default, routers with ATM interfaces use per-platform label space.
D) By default, routers with frame interfaces use per-platform label space.

Q34) Which two statements are correct? (Choose two.) (Source: Introducing MPLS Label Allocation, Distribution, and Retention Modes)

A) By default, cell-mode MPLS uses unsolicited downstream label distribution.
B) By default, cell-mode MPLS uses downstream-on-demand label distribution.
C) By default, frame-mode MPLS uses unsolicited downstream label distribution.
D) By default, frame-mode MPLS uses downstream-on-demand label distribution.

Q35) 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

Q36) Per-platform label space requires which of the following? (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.

Q37) 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

Q38) LDP sessions are initiated by using the _____ IP address. (Source: Discovering LDP Neighbors)

Q39) Exchanging initialization messages is what step in the LDP session negotiation process? (Source: Discovering LDP Neighbors)

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

Q40) By default, ATM LSRs establish IP adjacency across which VPI/VCI virtual circuit? (Source: Discovering LDP Neighbors)

A) 0/32
B) 1/32
C) 32/0
D) 32/1

Q41) LDP discovers nonadjacent neighbors by broadcasting _____ IP addresses. (Source: Discovering LDP Neighbors)

Q42) 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.

Q43) In frame-mode MPLS networks, the number of LDP sessions that are required between neighbors is determined by? (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

Module Self-Check Answer Key

Q1) C

Q2) A

Q3) A, B, D

Q4) D

Q5) B

Q6) A, D

Q7) D

Q8) B, D

Q9) C

Q10) C

Q11) D

Q12) A, B, C

Q13) B

Q14) A, C, D

Q15) A

Q16) A, C

Q17) B

Q18) C

Q19) C

Q20) B

Q21) C

Q22) A

Q23) B

Q24) C

Q25) B

Q26) C

Q27) B

Q28) B

Q29) per-interface, per-platform

Q30) A, B

Q31) independent, ordered (or ordered, independent)

Q32) C, D

Q33) D

Q34) B, C

Q35) B

Q36) A

Q37) D

Q38) higher

Q39) B

Q40) A

Q41) unicast

Q42) B, F

Q43) B

Module 3

Frame-Mode and Cell-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 and cell-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 and label-controlled ATM (LC-ATM) 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
- Configure LC-ATM MPLS
- Configure LC-ATM MPLS over ATM Virtual Path
- Monitor LC-ATM MPLS 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 CEF switching, including configuration and monitoring commands.

It is important to understand what part CEF switching plays in an MPLS network. CEF must be running as a prerequisite to running 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
- Explain how to configure IP CEF switching
- Describe how to monitor IP CEF switching

What Are Cisco IOS Platform Switching Mechanisms?

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

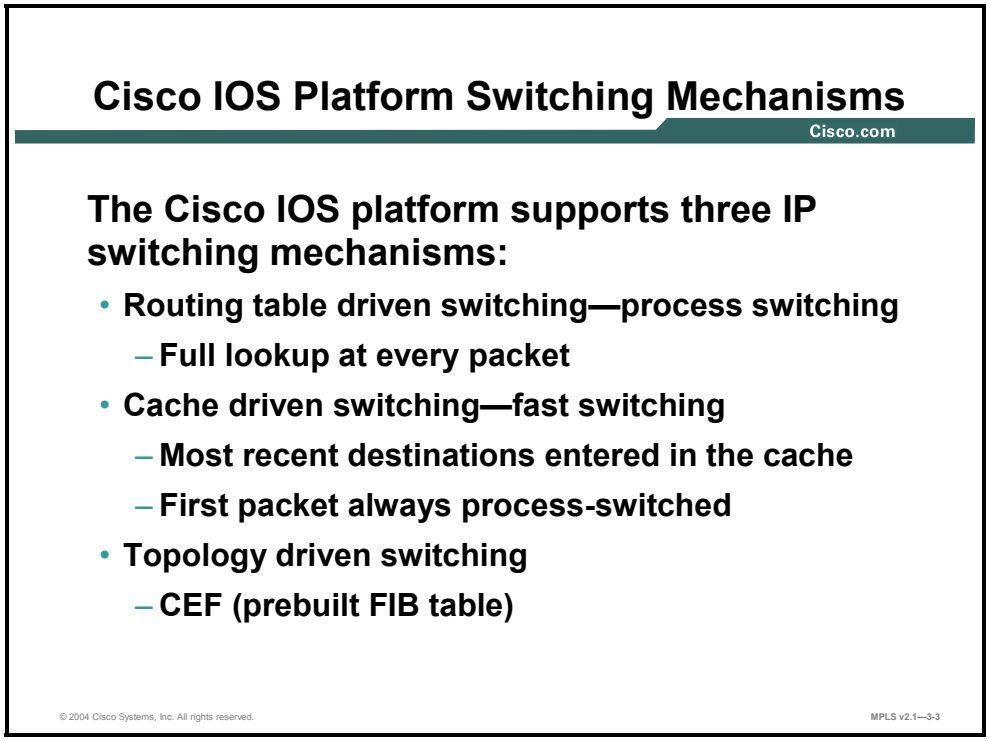

The first and the oldest switching mechanism available in Cisco routers is process switching. Because process switching must find a destination in the routing table (possibly a recursive lookup) and construct a new Layer 2 frame header for every packet, it is very slow and is normally not used.

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.

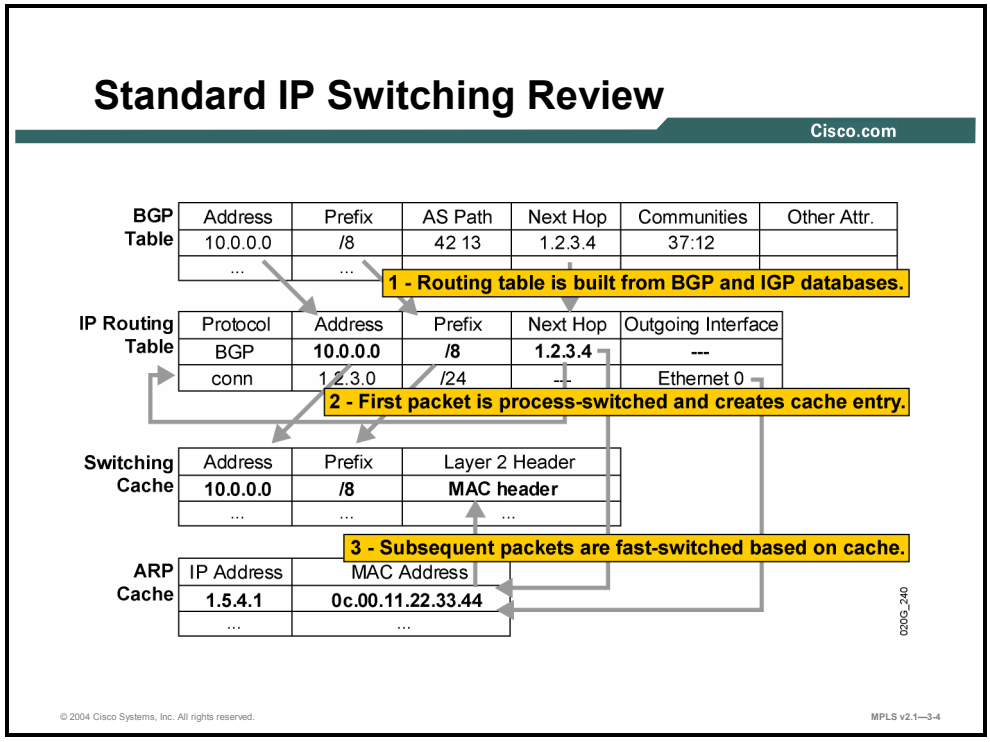

There is a specific sequence of events that occurs when process switching and fast switching are used for destinations learned through Border Gateway Protocol (BGP).

Example: Standard IP Switching

The figure illustrates this process. The following describes 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 CEF Switching Architecture?

This topic describes the architecture of CEF switching.

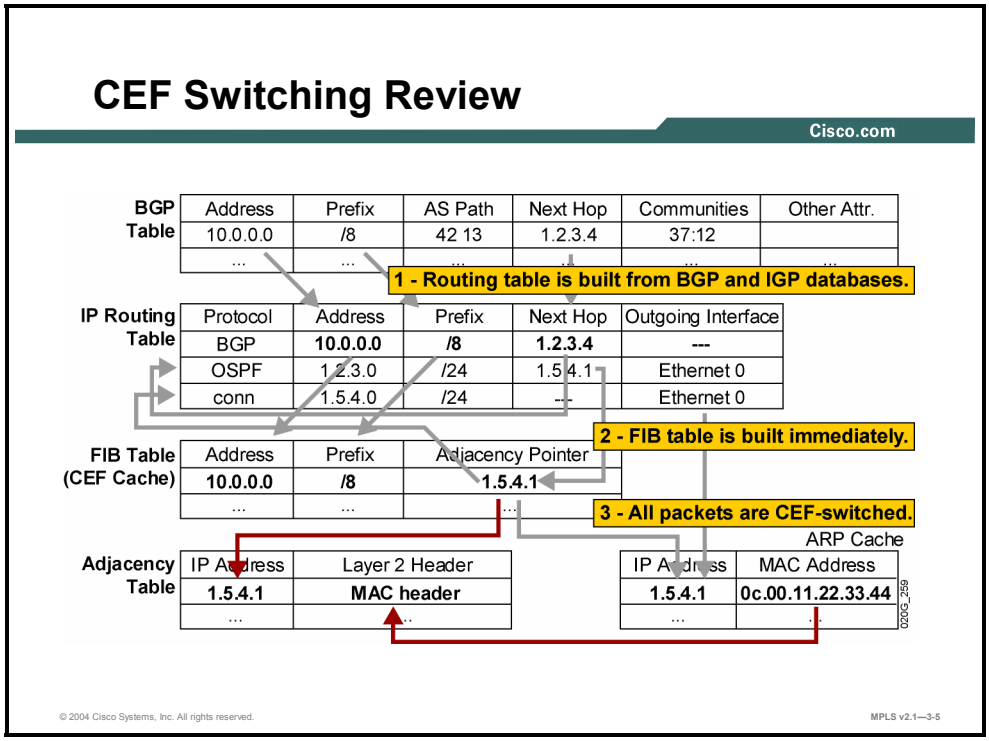

CEF uses a different architecture from process switching or any other cache-based switching mechanism. CEF uses a complete IP switching table, the FIB table, which holds the same information as the IP routing table. The generation of entries in the FIB table is not packet-triggered but change-triggered.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

Because the FIB 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. This 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 (around 100,000+ networks), enabling the CEF may consume additional memory. Enabling the 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.

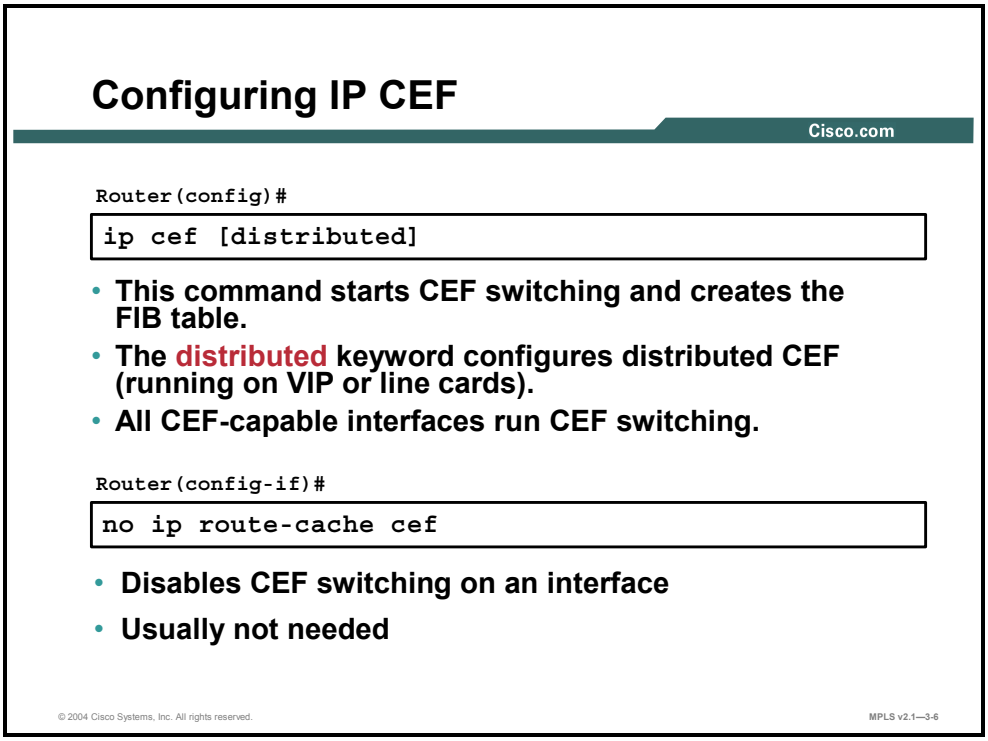

ip cef

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 following form of the two commands:

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

Syntax Description

distributed (optional): Enables the distributed CEF operation. 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 the Cisco 7100 series router.
- CEF is enabled on the Cisco 7200 series router.
- CEF is enabled on the Cisco 7500 series Internet router.
- Distributed CEF is enabled on the Cisco 6500 series router.
- Distributed CEF is enabled on the Cisco 12000 series Internet router.

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 following form 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.

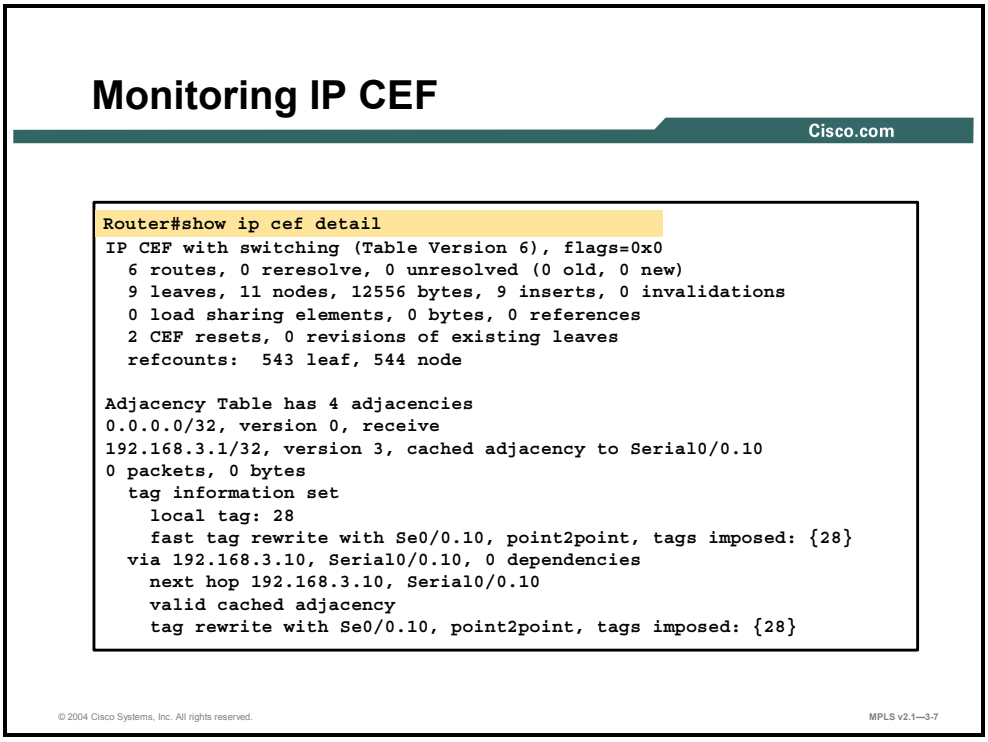

show ip cef

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

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

This 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 discussed in this lesson.

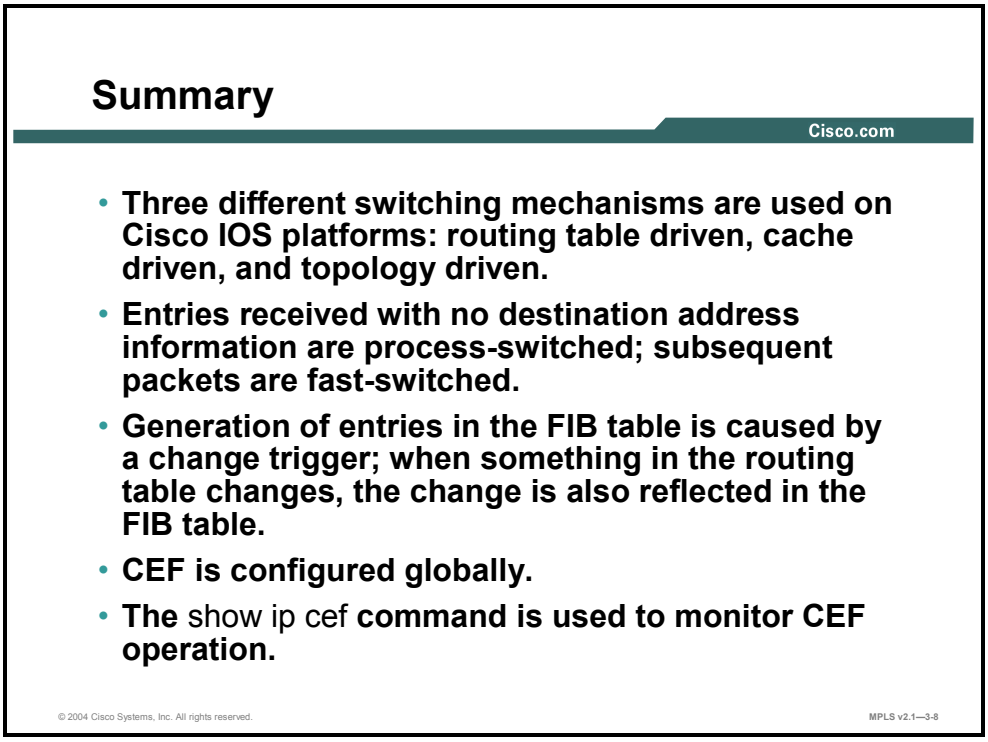

Lesson 2

Configuring Frame-Mode MPLS on Cisco IOS Platforms

Overview

This lesson describes how to configure frame-mode 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 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 for this lesson.

Objectives

This lesson describes 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.

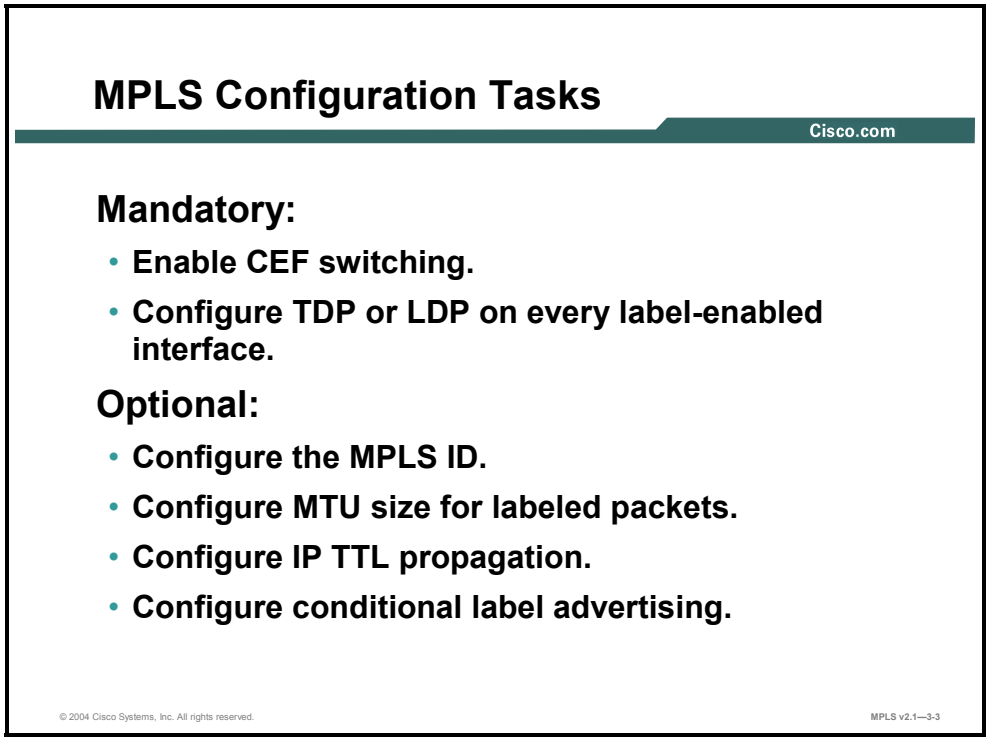

To enable MPLS, you must first enable CEF switching. Depending on the Cisco IOS software version, you may need to establish the range for the label pool.

You must enable Tag Distribution Protocol (TDP) or Label Distribution Protocol (LDP) on the interface by using either tag switching or 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 the conditional label advertising.

Configuring the MPLS ID on a Router

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

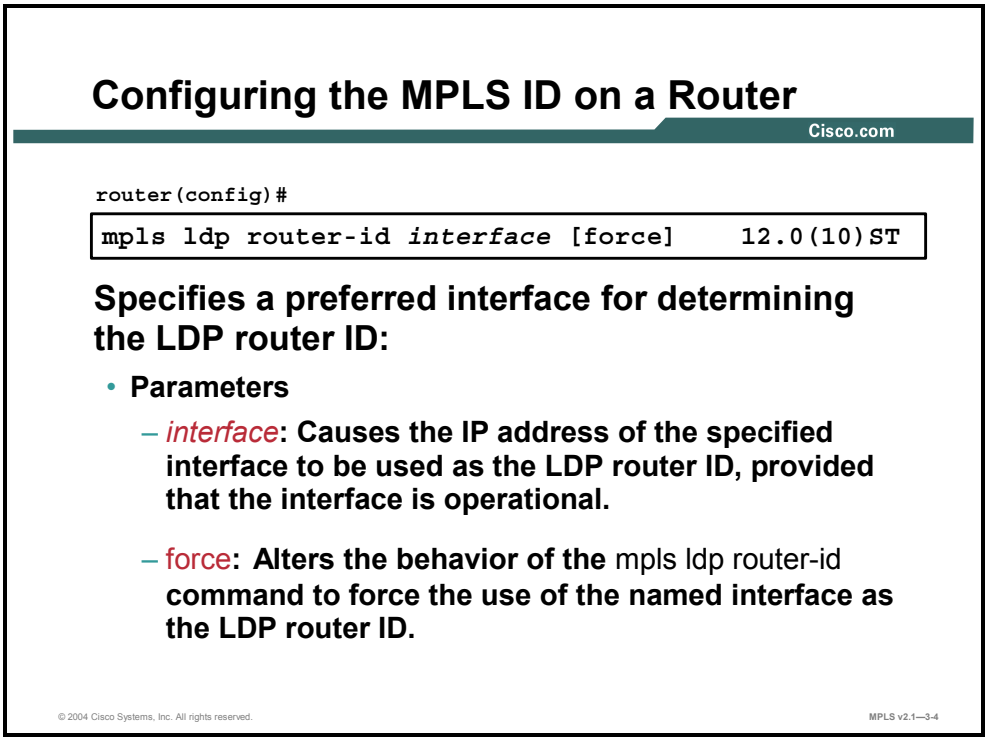

mpls ldp router-id

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. The following 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.

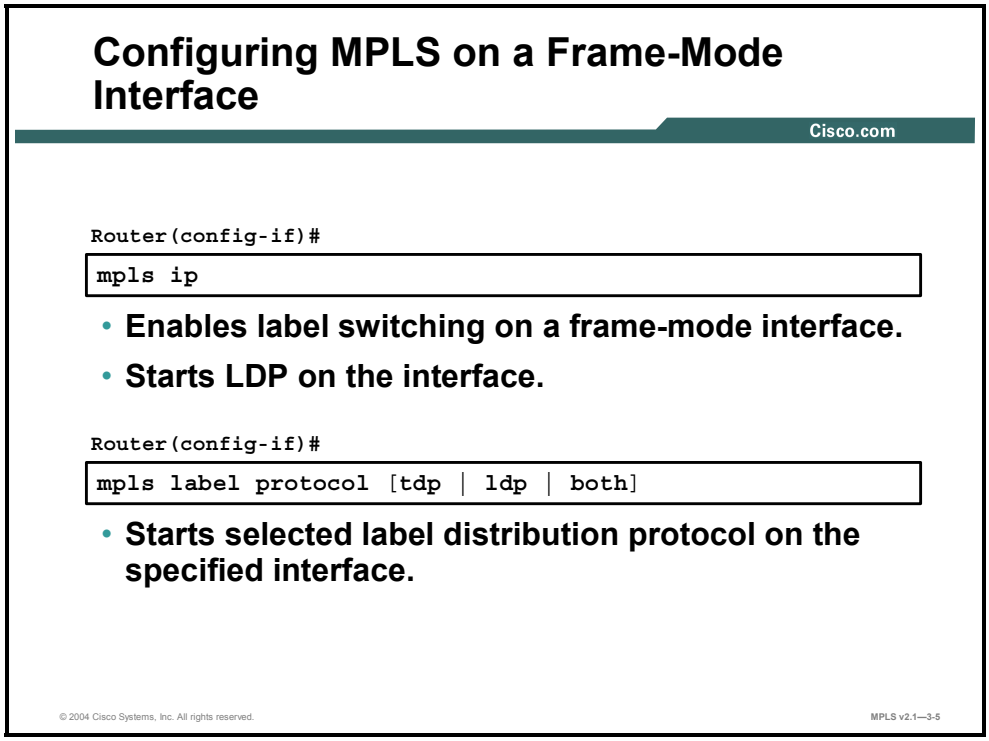

mpls ip

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. The following 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. The following 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 TDP on an interface.
`ldp`	Enables LDP on an interface.
`both`	Enables TDP and LDP on an interface.

Defaults

TDP is the default protocol.

Note For backward compatibility, using the “mpls” syntax will be entered as “tag-switching” syntax in the configuration by the IOS software.

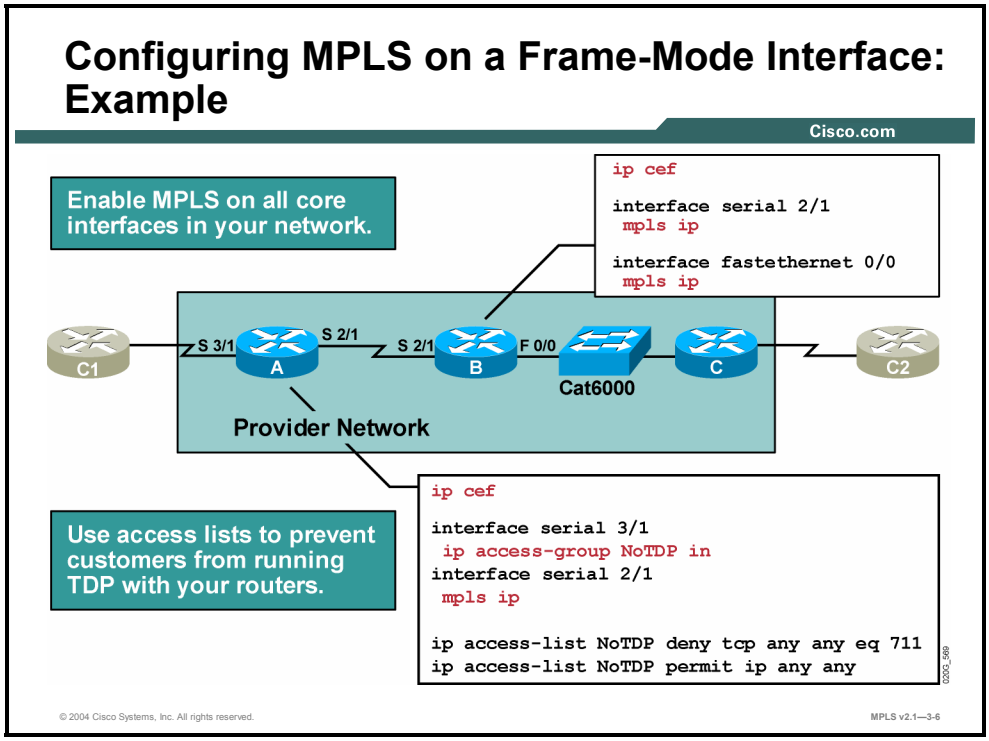

Example: Configuring MPLS on a Frame-Mode Interface

The figure shows the configuration steps needed to enable MPLS on an edge label switch router (LSR). The configuration includes an 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 "NoTDP" access-list on 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).

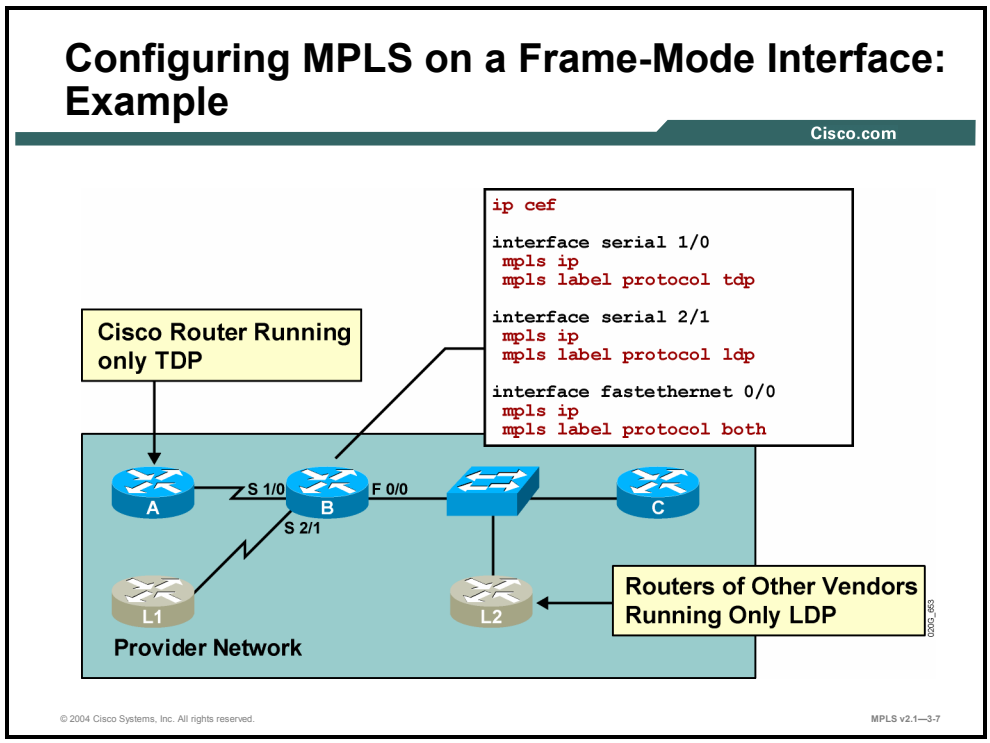

When combining Cisco routers with equipment of other vendors, you may need to use standard LDP (MPLS). TDP (tag switching) can be replaced by LDP on point-to-point interfaces. However, you can also use both protocols on shared media if some devices do not support TDP.

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.

Configuring a Label-Switching MTU

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

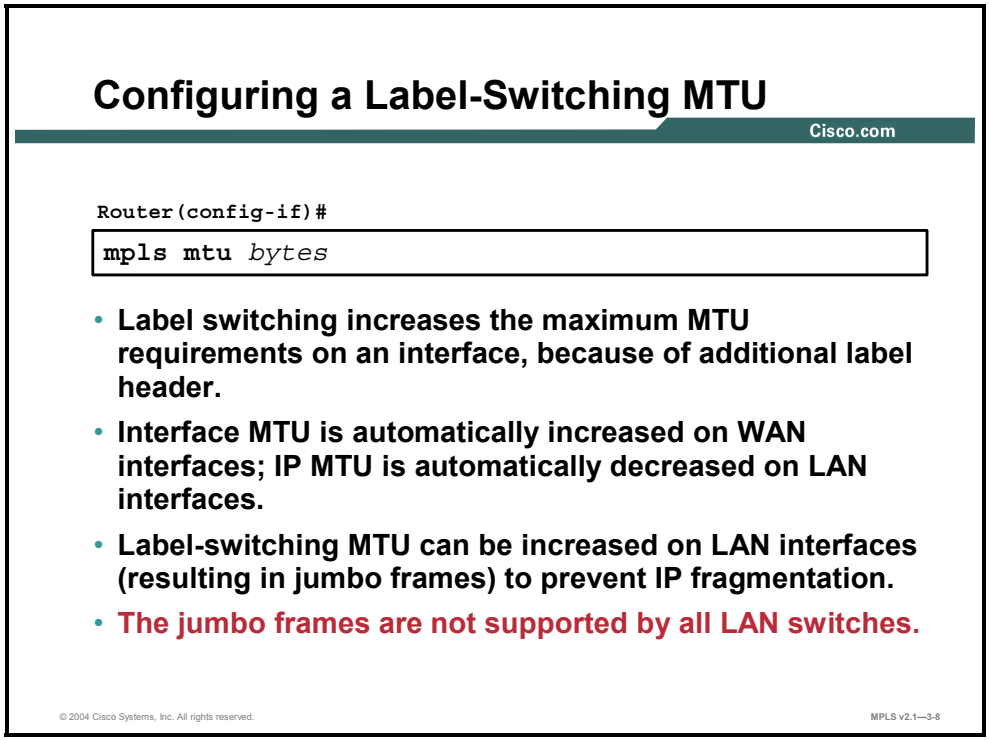

mpls mtu

To set the per-interface MTU for labeled packets, use the **mpls mtu** interface configuration command. The following 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 (B). The maximum depends on the type of interface medium.

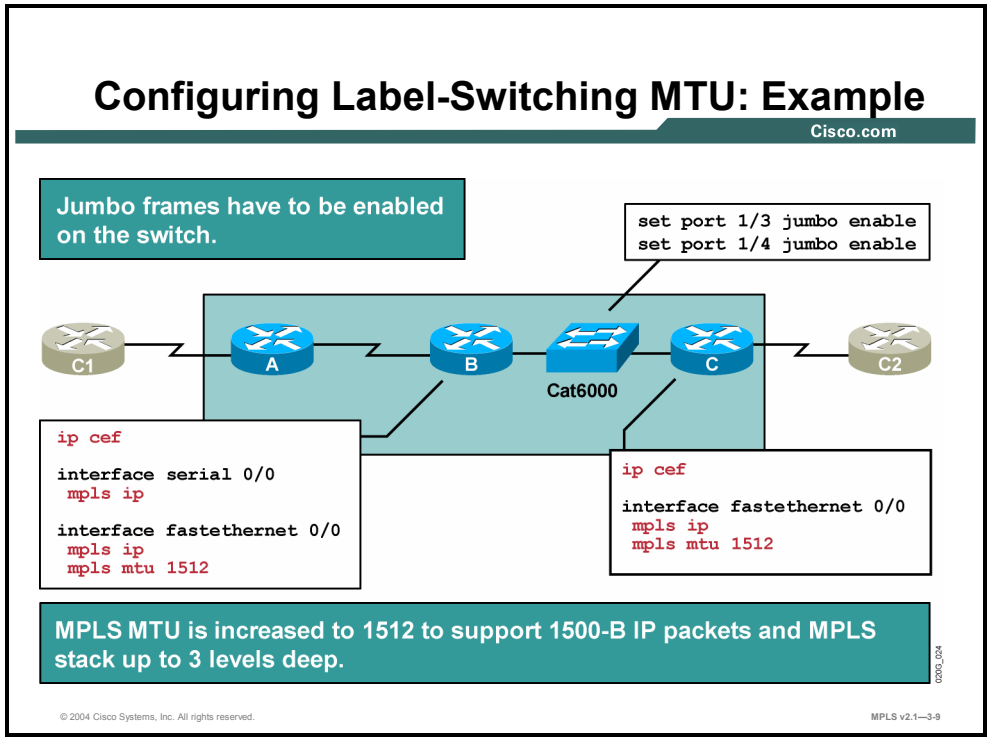

One way of preventing labeled packets from exceeding the maximum size (and being fragmented as a result) is to increase the MTU size of labeled packets for all segments in the label-switched path (LSP) tunnel. The problem will typically occur on LAN switches, where it is more likely that a device does 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 (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.

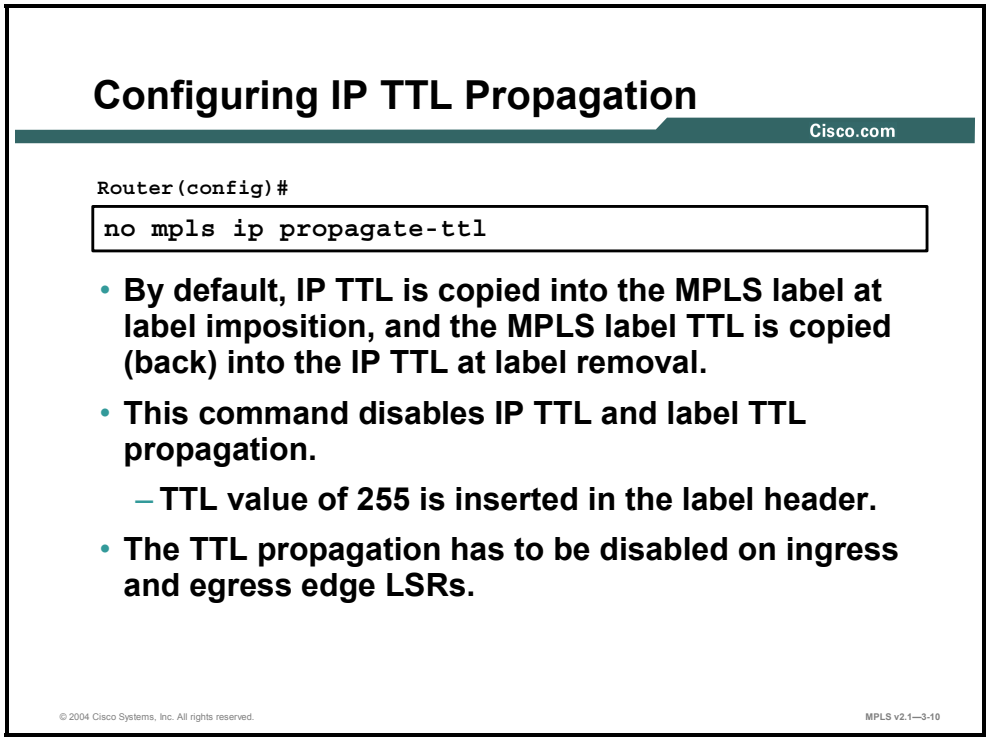

mpls ip propagate-ttl

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. The following 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.

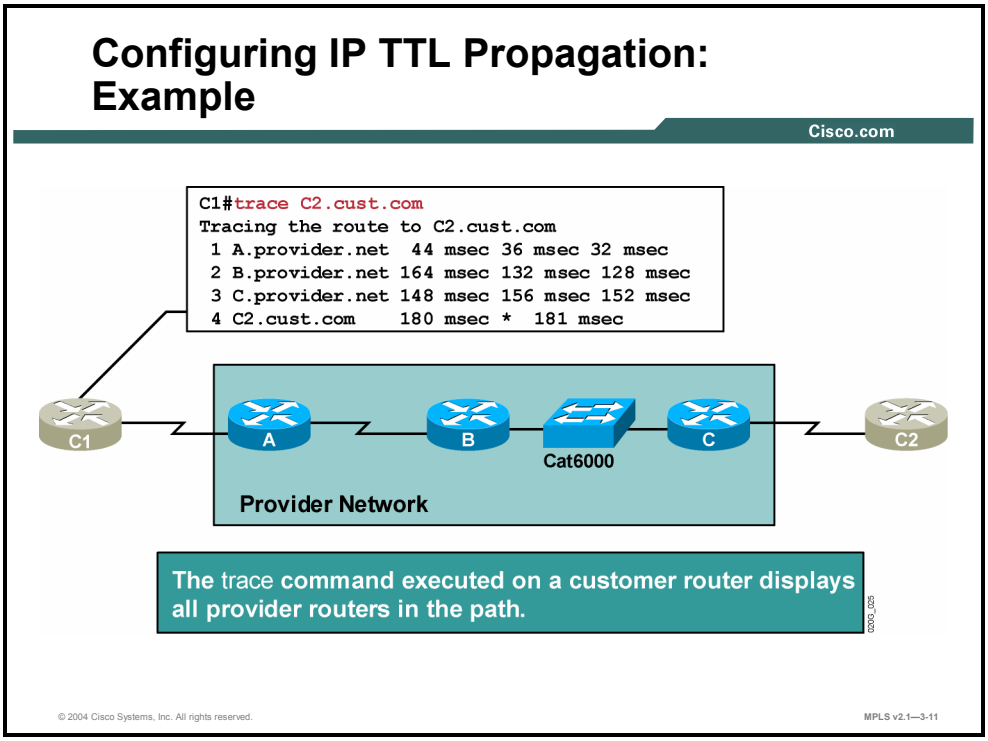

Example: Configuring IP TTL Propagation

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.

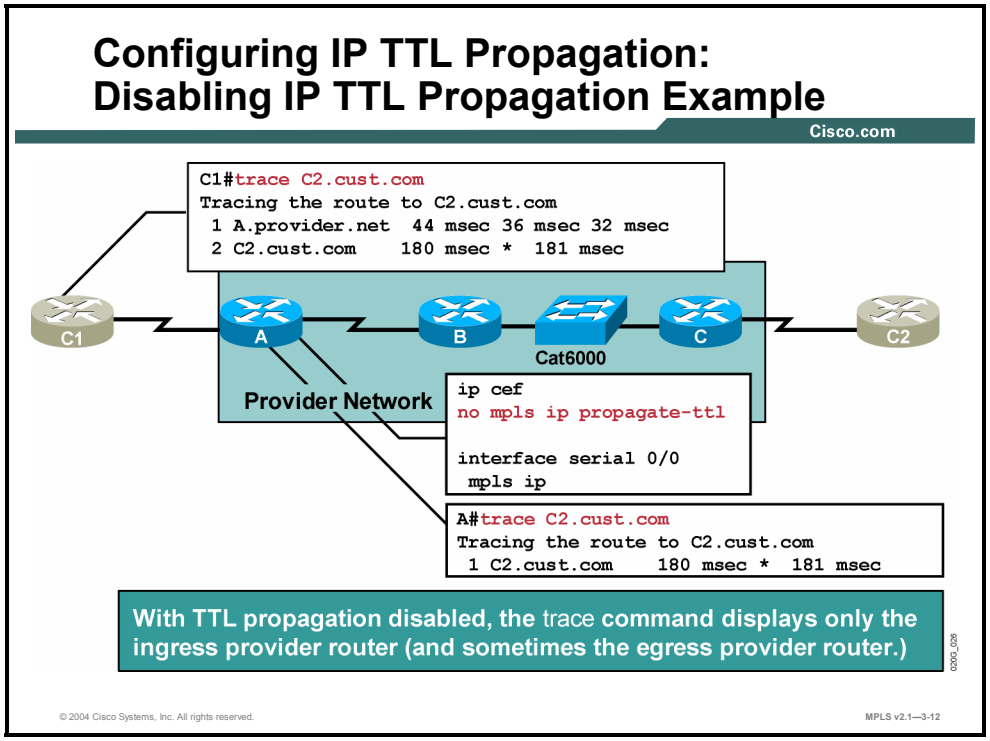

If TTL propagation is disabled, the TTL value is not copied into the label header. Instead, the label TTL field is set to 255. The probable result is that no router in the TTL field in the label header will be decreased to 0 inside the MPLS domain (unless there is a forwarding loop inside the MPLS network).

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 the following 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 BGP, not in the Interior Gateway Protocol (IGP).

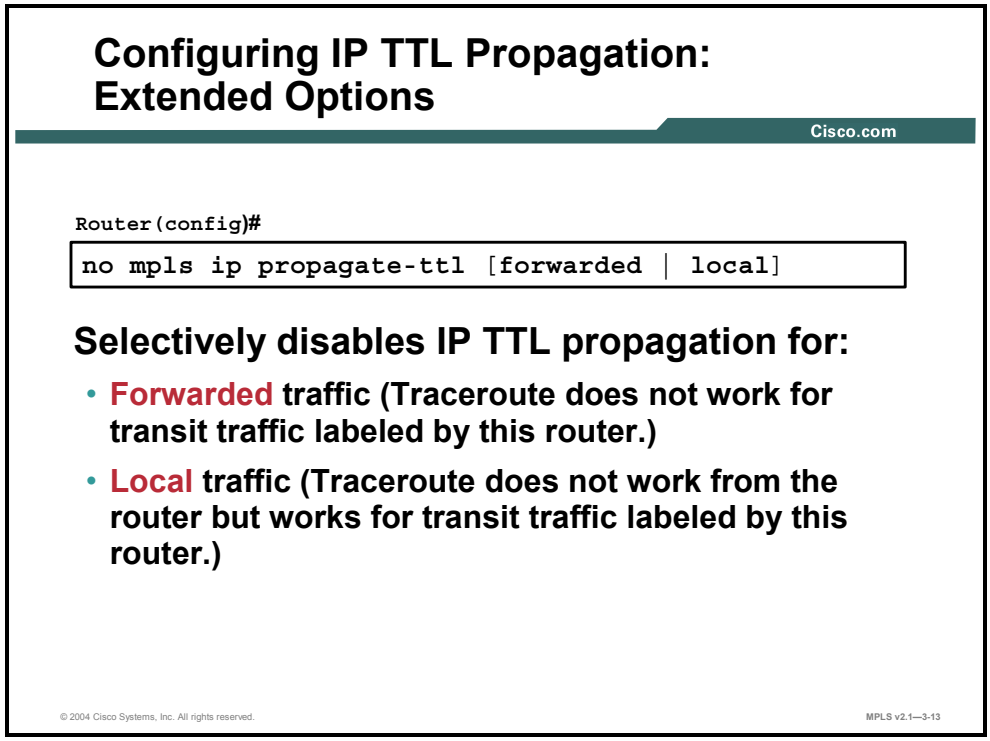

mpls ip propagate-ttl

Use the **mpls ip propagate-ttl** 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.

Command Modes

The **mpls ip propagate-ttl** command is used in global configuration mode.

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."

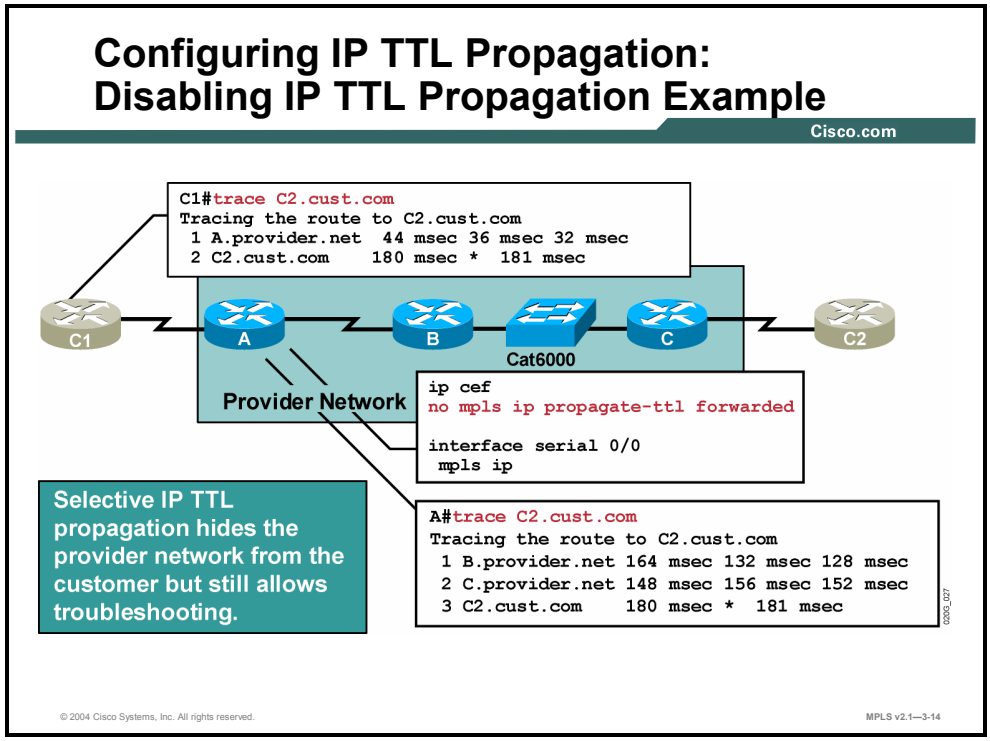

Typically, a service provider likes to hide the backbone network from outside users but allow inside traceroute to work for easier troubleshooting of the network.

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.

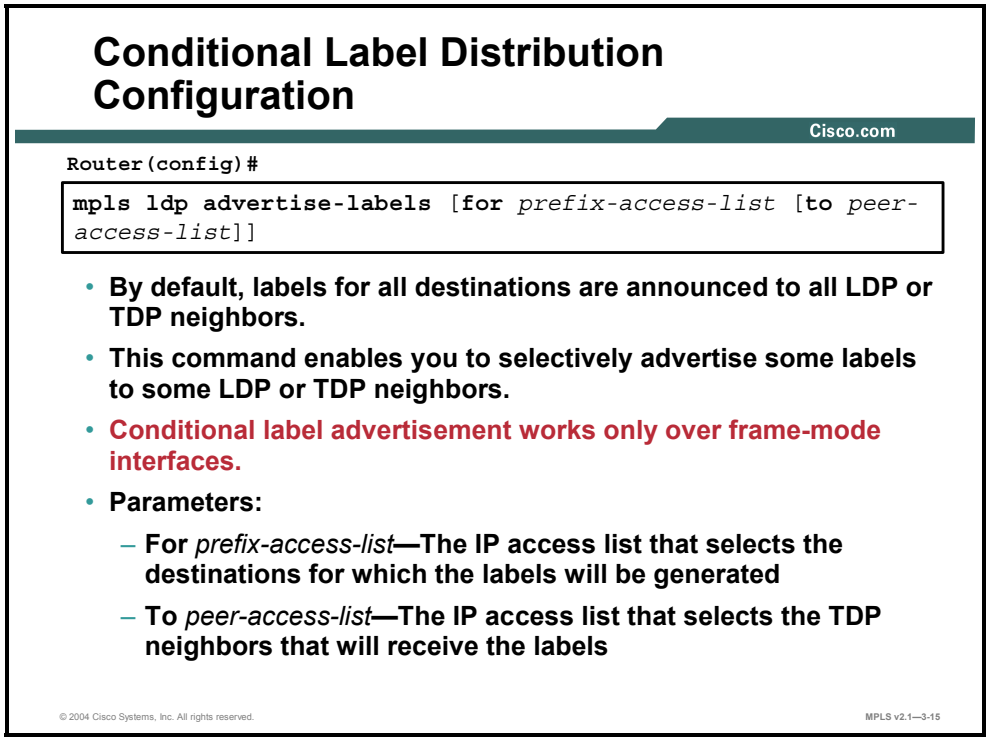

mpls ldp advertise-labels

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) Specifies which destinations should have their labels advertised.
to *peer-access-list*	(Optional) Specifies which LSR neighbors should receive label advertisements. An LSR is identified by its router ID, which consists of the first 4 bytes (B) of its 6-B LDP identifier.

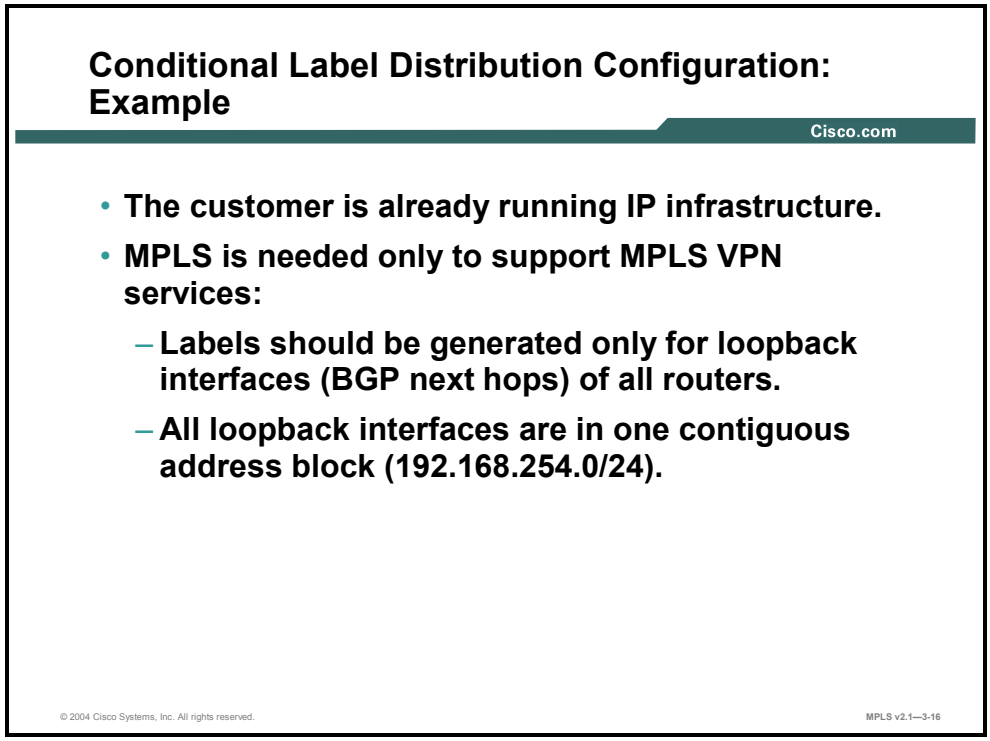

Example: Conditional Label Distribution Configuration

The example here describes where conditional label advertising can be used. The existing network still performs normal IP routing, but the MPLS label-switched path (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.

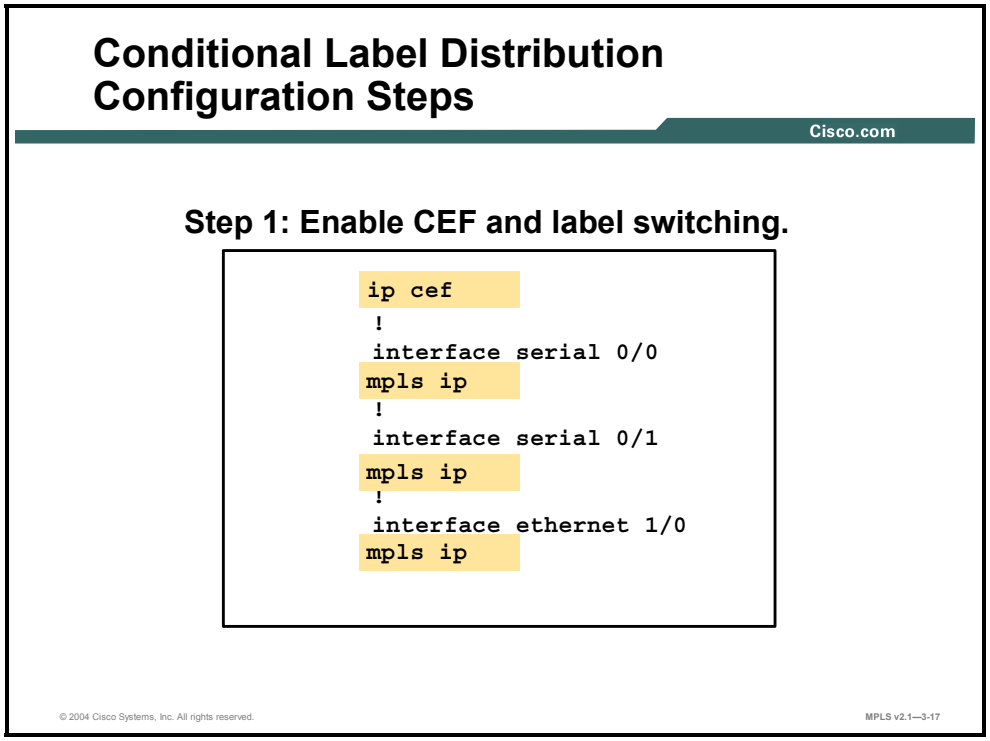

In the first step, CEF switching and MPLS have to be enabled on all core interfaces. The MPLS MTU size may be adjusted on the LAN interfaces.

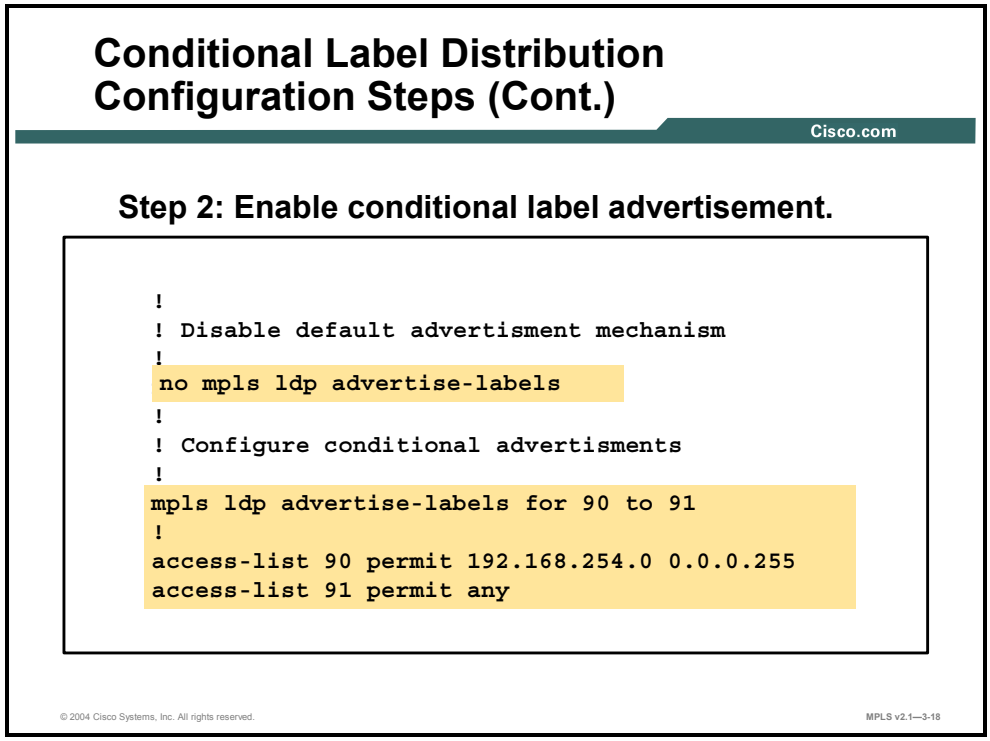

In the second step, disable label propagation and enable conditional label advertising. Within the **mpls ldp advertise-labels** command, specify the neighbors to which the labels are to be sent and the networks for which the labels are to be advertised.

Example: Enabling Conditional Label Advertisement

In the figure, the labels for all networks permitted by access control list (ACL) 90 are sent to all neighbors matched by ACL 91 (in this example, that 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.

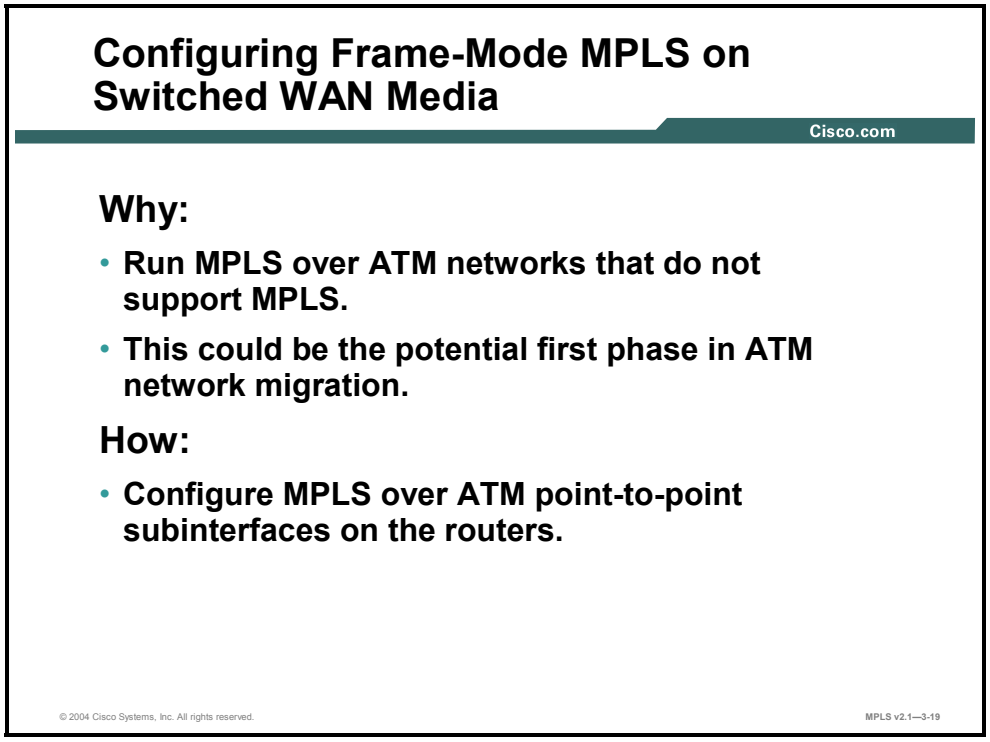

When an underlying ATM infrastructure that does not support cell-mode MPLS is used, MPLS can still be used across point-to-point permanent virtual circuits (PVCs). The MPLS configuration is equal to that on any other Layer 2 media.

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

Cisco.com

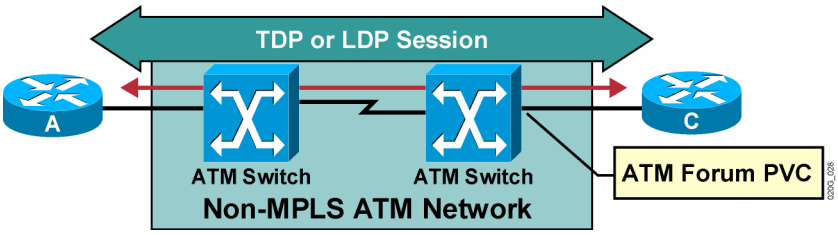

- **Routers view the ATM PVC as a frame-mode MPLS interface.**
- **TDP or LDP is run between the adjacent routers.**
- **Many LSPs can be established over one ATM PVC.**
- **The ATM network is not aware of MPLS between the routers.**

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.)

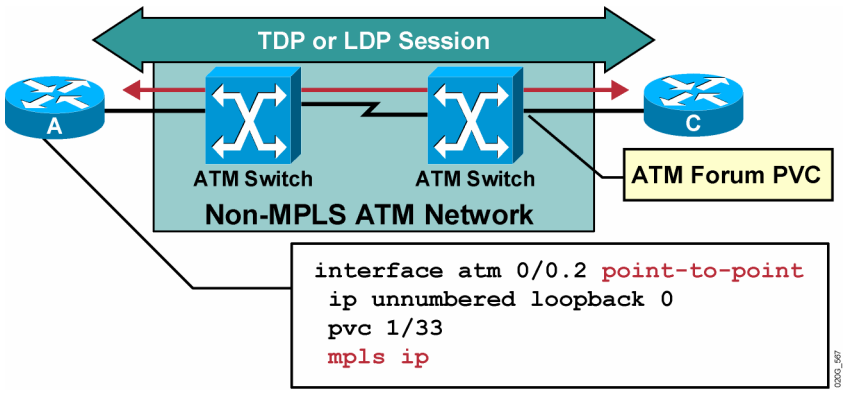

- **Create a point-to-point ATM subinterface.**
- **Configure ATM PVC on the subinterface.**
- **Start label switching and LDP or TDP on the interface.**

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.

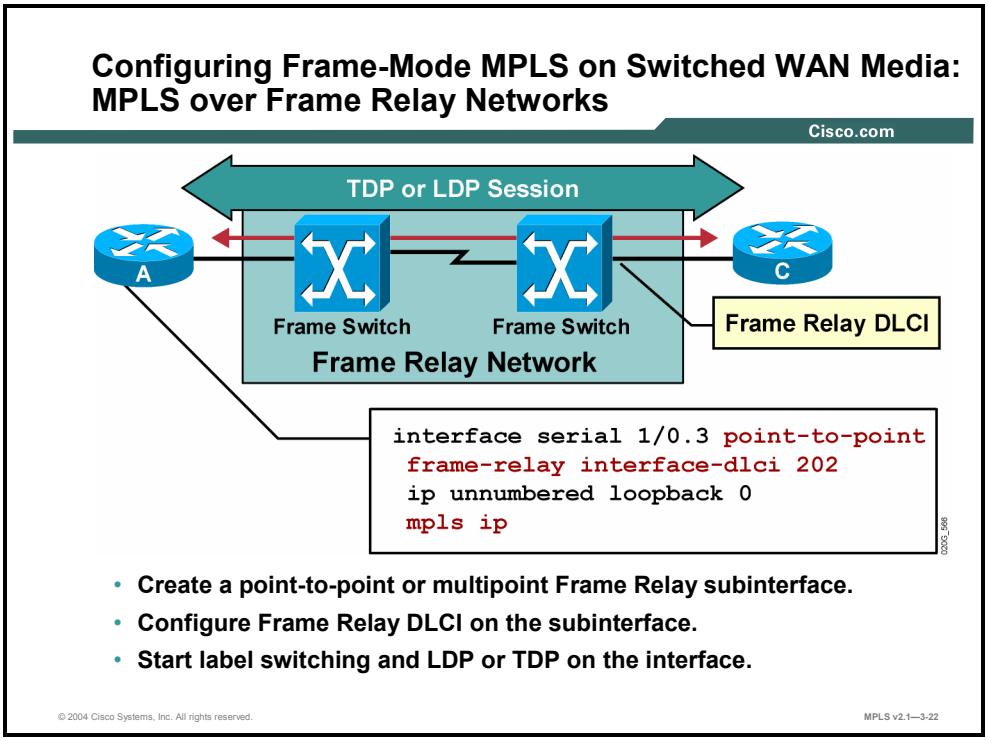

Enabling MPLS on a Frame Relay PVC, also called a data-link connection identifier (DLCI), is no different from doing so on any other point-to-point media.

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 discussed in this lesson.

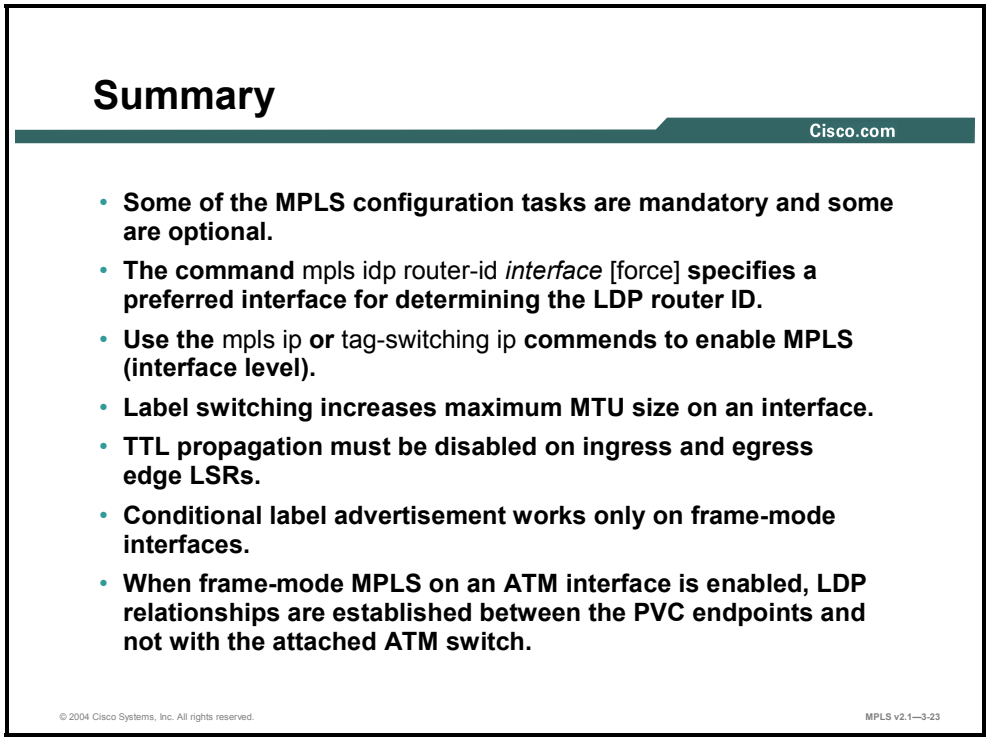

Lesson 3

Monitoring Frame-Mode MPLS on Cisco IOS Platforms

Overview

This lesson covers the procedures for monitoring MPLS on Cisco IOS platforms by listing the syntax and parameter descriptions; looking at interfaces, neighbor nodes, and LIB and label forwarding information base (LFIB) tables; and outlining the usage guidelines for the commands. The lesson also looks at common frame-mode MPLS symptoms and issues.

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:

- Describe how to monitor MPLS
- Describe how to monitor LDP
- Describe how to monitor label switching
- Describe how to debug MPLS and LDP

Monitoring MPLS

This topic describes how to monitor MPLS.

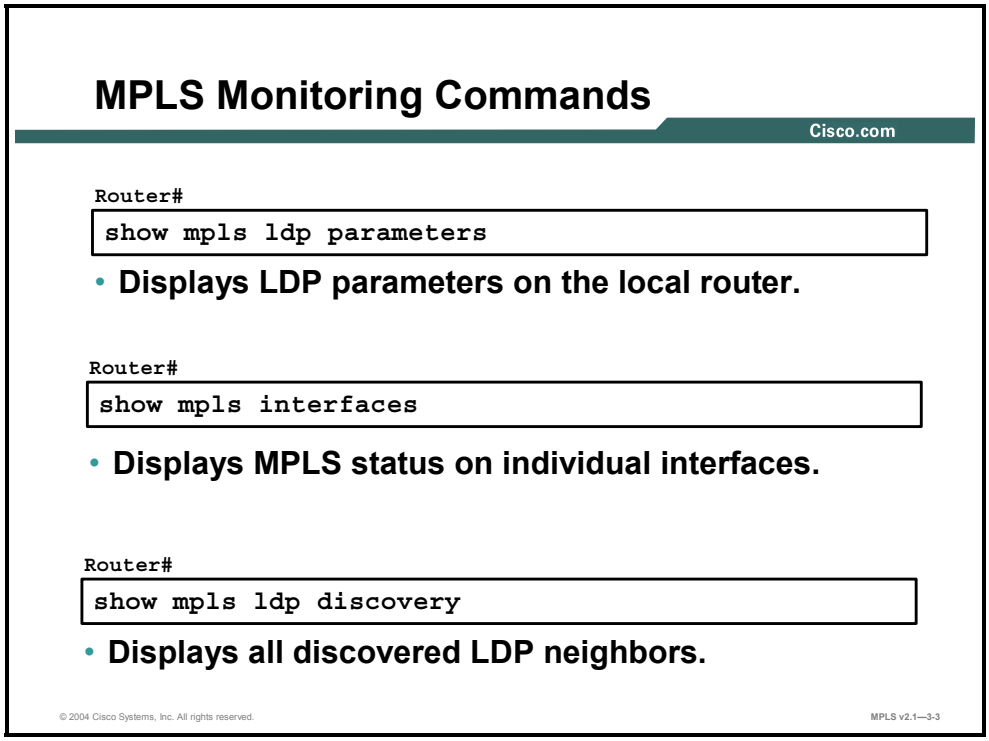

show mpls ldp parameters

To display available LDP parameters, use the following **show mpls ldp parameters** command in privileged EXEC mode: **show mpls ldp parameters**.

show mpls interfaces

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

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

show mpls interfaces Syntax Description

Parameter	Description
interface	(Optional) The interface about which to display MPLS information.
detail	(Optional) Displays information in long form.

show mpls ldp discovery

To display the status of the LDP discovery process (Hello protocol), use the **show mpls ldp discovery** command in privileged EXEC mode. This command displays all MPLS-enabled interfaces and the neighbors that are present on the interfaces.

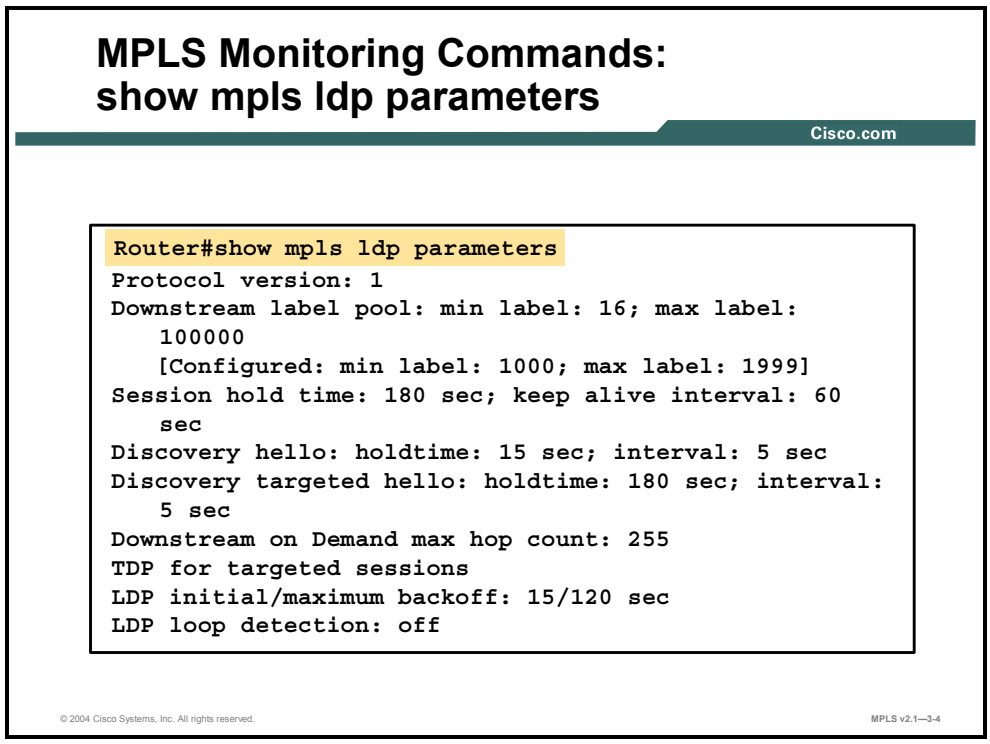

show mpls ldp parameters

To display available LDP parameters, use the following **show mpls ldp parameters** command in privileged EXEC mode: **show mpls ldp parameters**.

Syntax Description

This command has no arguments or keywords.

This table describes the significant fields in the display.

show mpls ldp parameters Field Description

Field	Description
Protocol version	Indicates the version of LDP running on the platform.
Downstream label pool	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	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	Indicates the interval of time between consecutive transmissions of LDP keepalive messages to an LDP peer.
Discovery hello	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	Indicates the amount of time to remember that a neighbor platform wants an LDP session when one of the following 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. 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	Reports the parameters that have been set by the **show mpls atm-ldp bindings** command.
LDP initial/maximum backoff	Reports the parameters that have been set by the **mpls ldp backoff** command.

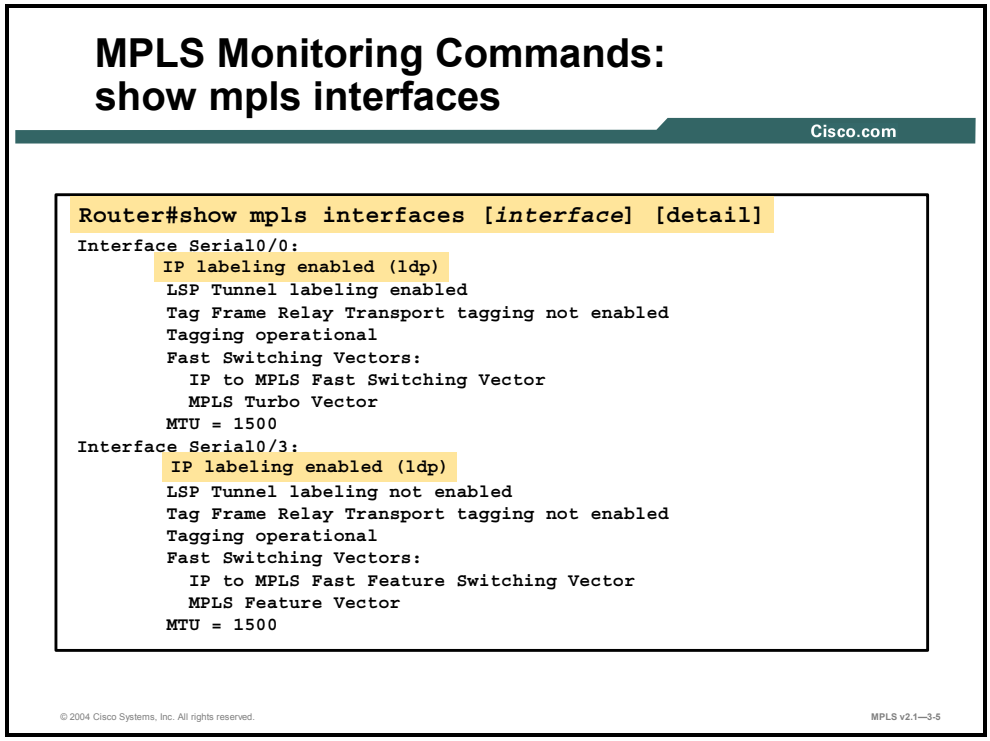

The **show mpls interfaces** command will show only those interfaces on which MPLS has been configured.

show mpls interfaces

To display information about one or more or all interfaces that are configured for label switching, use the following **show mpls interfaces** command in privileged EXEC mode: **show mpls interfaces** [**all**].

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.

This 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 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.

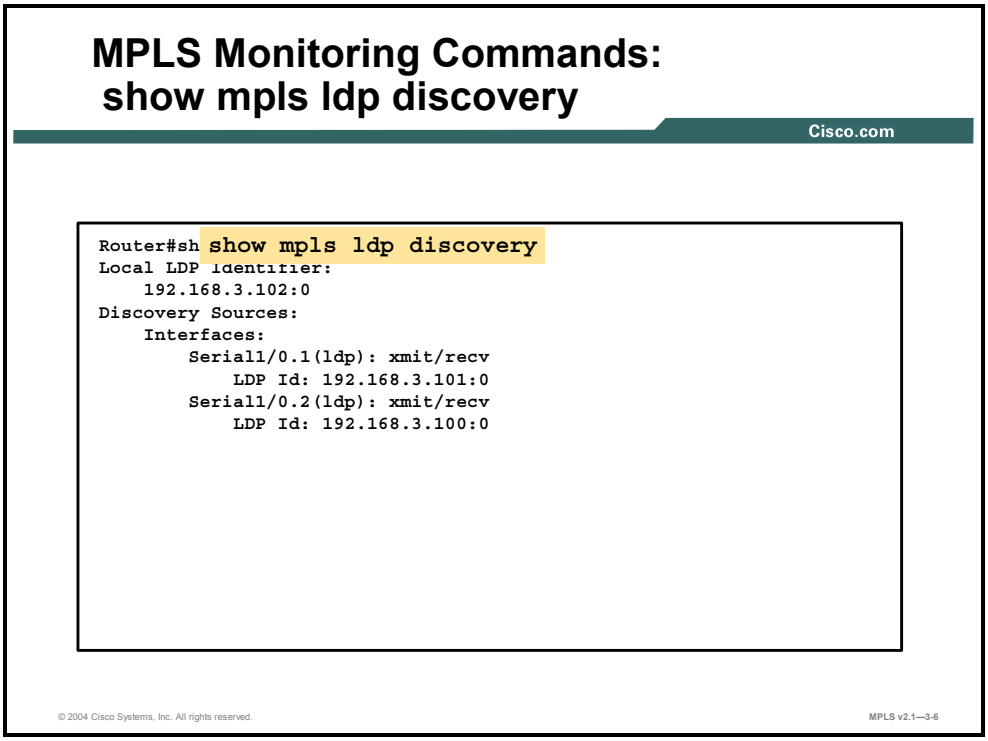

show mpls ldp discovery

To display the status of the LDP discovery process, use the **show mpls ldp discovery** command in privileged EXEC mode. This command generates a list of interfaces over which the LDP discovery process is running. The following shows these commands:

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

show mpls ldp discovery Syntax Description

Parameter	Description
vrf *vpn-name*	(Optional) Displays the neighbor discovery information for the specified 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.

This table describes the significant fields in the display.

show mpls ldp discovery Field Description

Field	Description
`Local LDP Identifier`	The LDP identifier for the local router. An LDP identifier is a 6-B construct displayed in the form “IP address:number.” By convention, the first 4 bytes of the LDP identifier constitute the router ID; integers, starting with 0, constitute the final 2 bytes of the IP address: number construct.
`Interfaces`	Lists the interfaces that are engaging in LDP discovery activity, 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 TDP) identifiers indicate LDP (or TDP) neighbors discovered on the interface.
`Targeted Hellos`	Lists the platforms to which targeted hello messages are being sent, described here: ■ The xmit, recv, and (ldp) or (tdp) fields are as described for the Interfaces field. ■ The active field indicates that this 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.

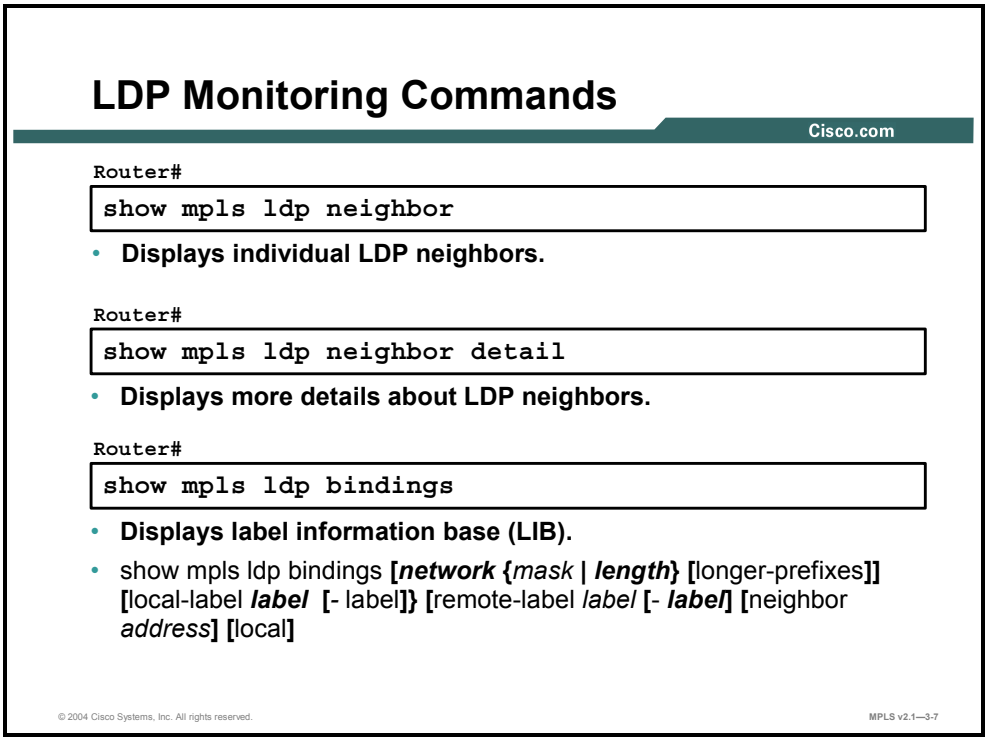

show mpls ldp neighbor

To display the status of LDP sessions, use the following **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) 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 the following **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) Displays the label bindings for the specified VPN routing or forwarding instance (vpn-name).
network	(Optional) Defines the destination network number.
mask	(Optional) Specifies the network mask, written as A.B.C.D.
length	(Optional) Specifies the mask length (1 to 32 characters).
longer-prefixes	(Optional) Selects any prefix that matches *mask* with a length from 1 to 32 characters.
local-label *label-label*	(Optional) Displays entries matching local label values. Use the *label-label* argument to indicate the label range.
remote-label *label-label*	(Optional) Displays entries matching the label values assigned by a neighbor router. Use the *label-label* argument to indicate the label range.
neighbor *address*	(Optional) Displays the label bindings assigned by the selected neighbor.
local	(Optional) Displays the local label bindings.

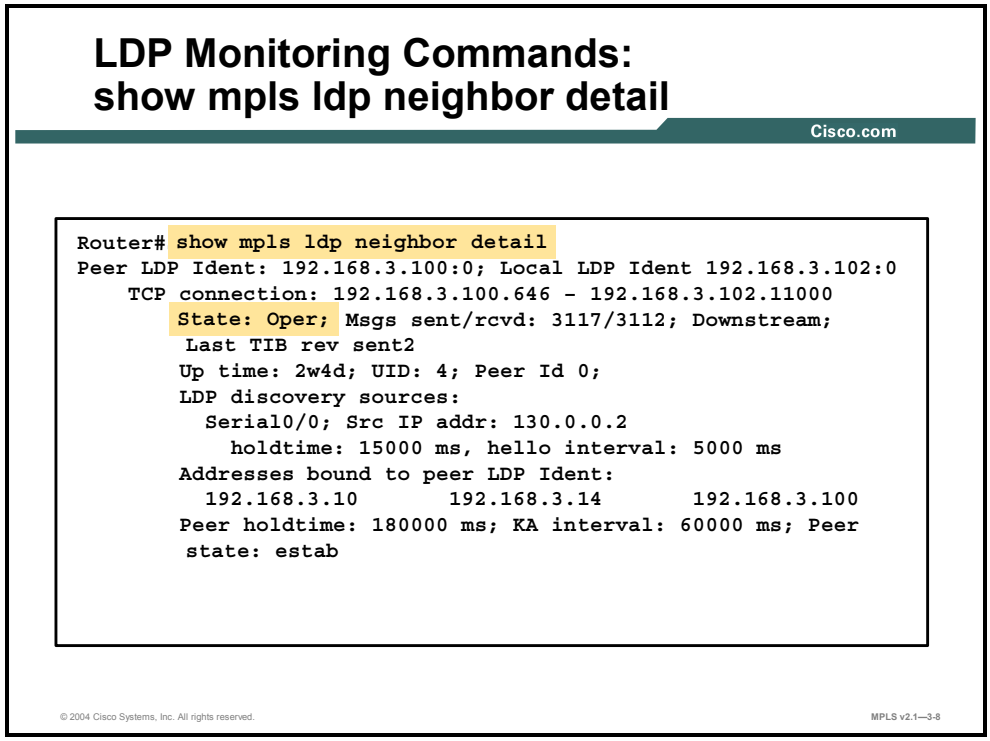

The status of the LDP (TDP) session is indicated by “State: Oper” (operational).

show mpls ldp neighbor

To display the status of LDP sessions, issue the following **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	Displays LDP identifier of the neighbor (peer) for this session.
Local LDP Ident	Displays LDP identifier for the local LSR for this session.
TCP connection	Displays TCP connection used to support the LDP session, shown in the following format: ■ peer IP address.peer port ■ local IP address.local port
State	Displays state of the LDP session. Generally, this is “Oper” (operational), but “transient” is another possible state.
Msgs sent/rcvd	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	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	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	Displays length of time that the LDP session has existed.
LDP discovery sources	Displays source(s) of LDP discovery activity that led to the establishment of this LDP session.
Addresses bound to peer LDP Ident	Displays 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	Displays the time that it takes to remove the relationship if no keepalives are received within this period.
KA interval	Displays the keepalive interval.
Peer state	Shows the status of the neighbor relationship.

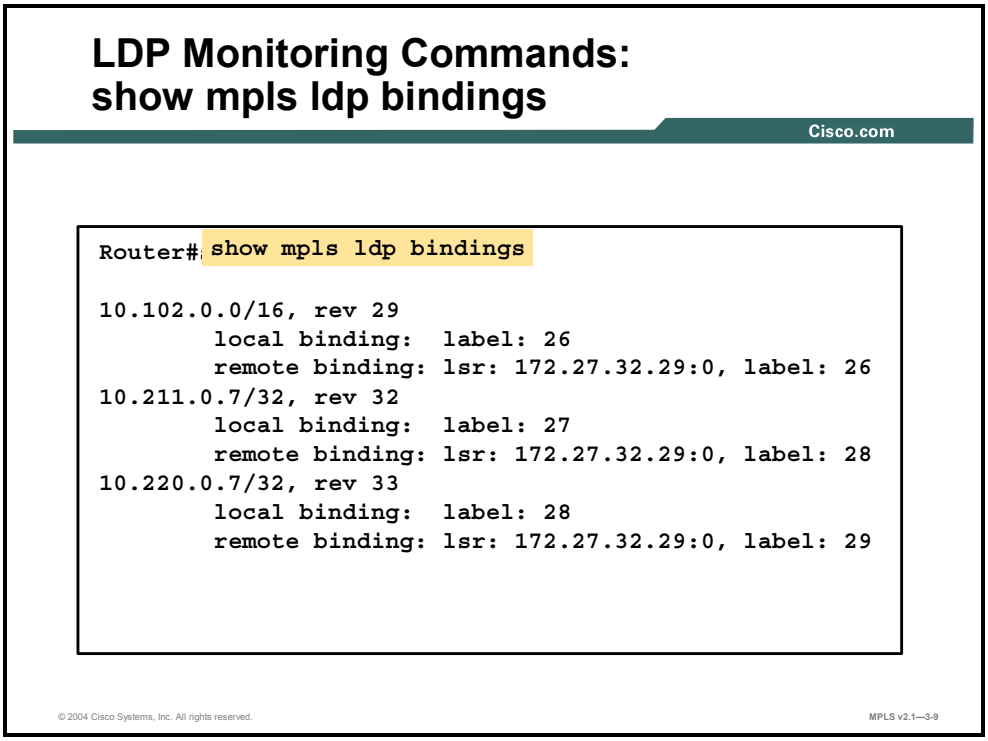

show mpls ldp bindings

To display the contents of the LIB, use the following **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) Displays the label bindings for the specified VPN routing or forwarding instance (vpn-name).
network	(Optional) Defines the destination network number.
mask	(Optional) Specifies the network mask, written as A.B.C.D.
length	(Optional) Specifies the mask length (1 to 32 characters).
longer-prefixes	(Optional) Selects any prefix that matches *mask* with a length from 1 to 32 characters.
local-label *label-label*	(Optional) Displays entries matching local label values. Use the *label-label* argument to indicate the label range.
remote-label *label-label*	(Optional) Displays entries matching the label values assigned by a neighbor router. Use the *label-label* argument to indicate the label range.
neighbor *address*	(Optional) Displays the label bindings assigned by the selected neighbor.
local	(Optional) 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.

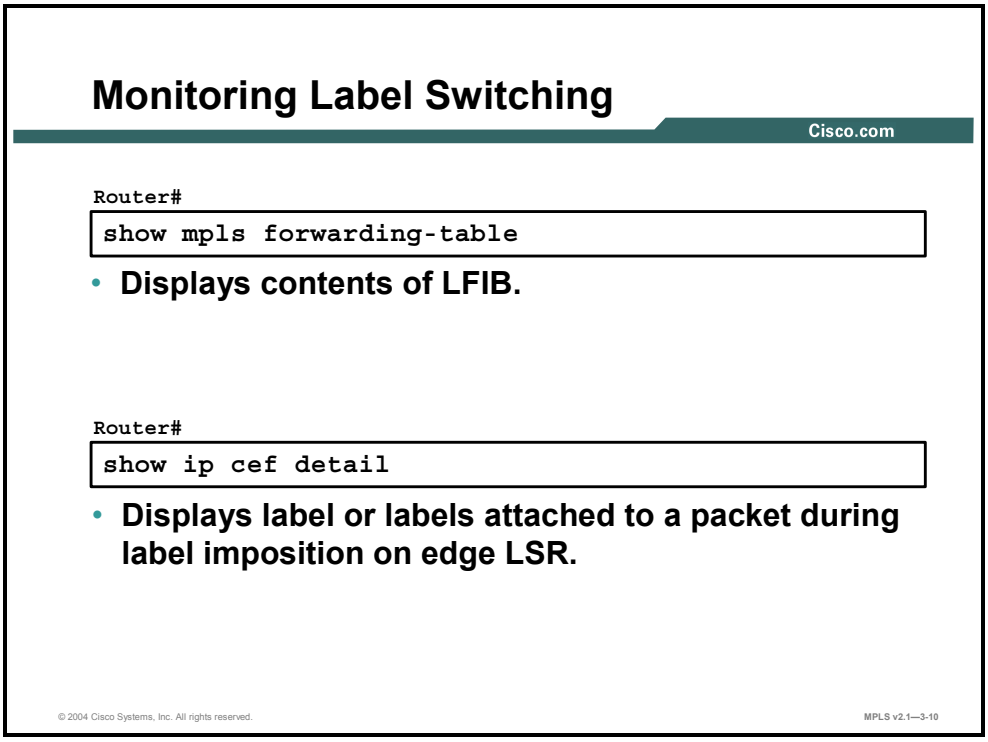

show mpls forwarding-table

To display the contents of the MPLS LFIB, use the following **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 following form of the **show ip cef** in privileged EXEC mode: **show ip cef** [**unresolved** | **summary**].

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

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

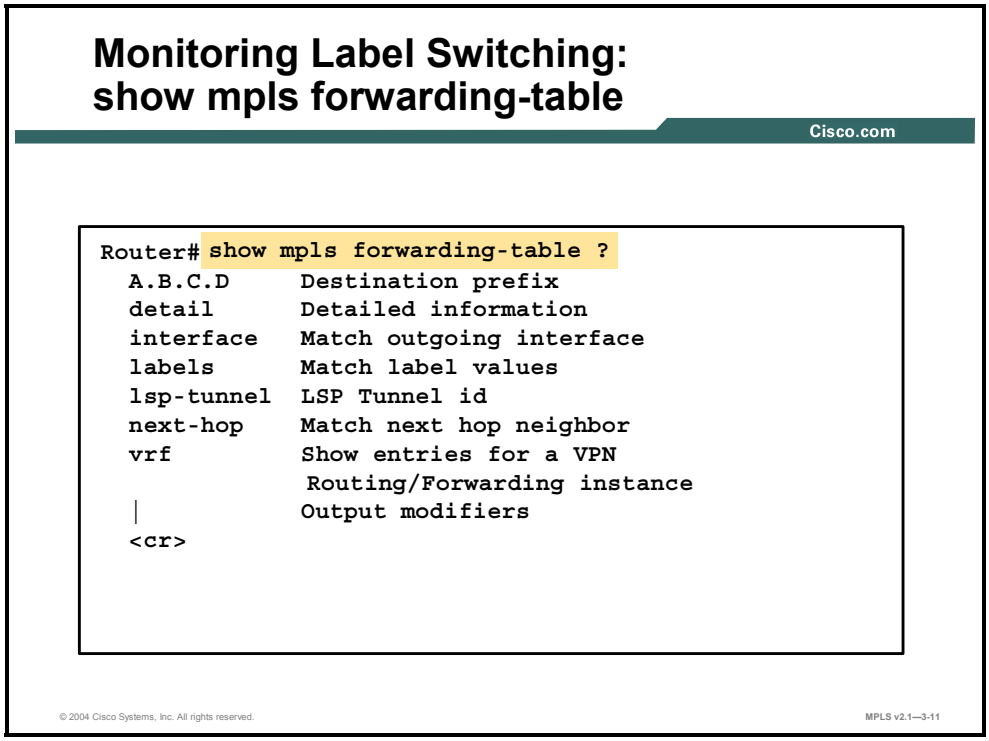

show mpls forwarding-table

To display the contents of the MPLS LFIB, use the following **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) Destination network number.
mask	IP address of destination mask whose entry is to be shown.
length	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, 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     1/3310.15.0.0/16      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 a LSP tunnel.

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

Monitoring Label Switching: show mpls forwarding-table detail

Cisco.com

```
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      point2point
        MAC/Encaps=4/4, MTU=1504, Tag Stack{}
        0F008847
        No output feature configured
    Per-packet load-sharing
71     Pop tag     192.168.3.4/32    0          Se0/0      point2point
        MAC/Encaps=4/4, MTU=1504, Tag Stack{}
        0F008847
        No output feature configured
    Per-packet load-sharing
```


This table describes the significant fields in the display.

show mpls forwarding table Field Description

Field	Description
Local tag	Displays label assigned by this router.
Outgoing tag or VC	Displays label assigned by next hop or 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	Displays address or tunnel to which packets with this label are going.
Bytes tag switched	Displays number of bytes switched with this incoming label.
Outgoing interface	Displays interface through which packets with this label are sent.
Next Hop	Displays IP address of neighbor that assigned the outgoing label.
MAC/Encaps	Displays length in bytes of Layer 2 header, and length in bytes of packet encapsulation, including Layer 2 header and label header.
MTU	Displays MTU of labeled packet.
Tag Stack	Displays all the outgoing labels. If the outgoing interface is transmission convergence-ATM (TC-ATM), the virtual circuit descriptor (VCD) is also shown.
00020900 00002000	Displays the actual encapsulation in hexadecimal form. There is a space shown between Layer 2 and the label header.

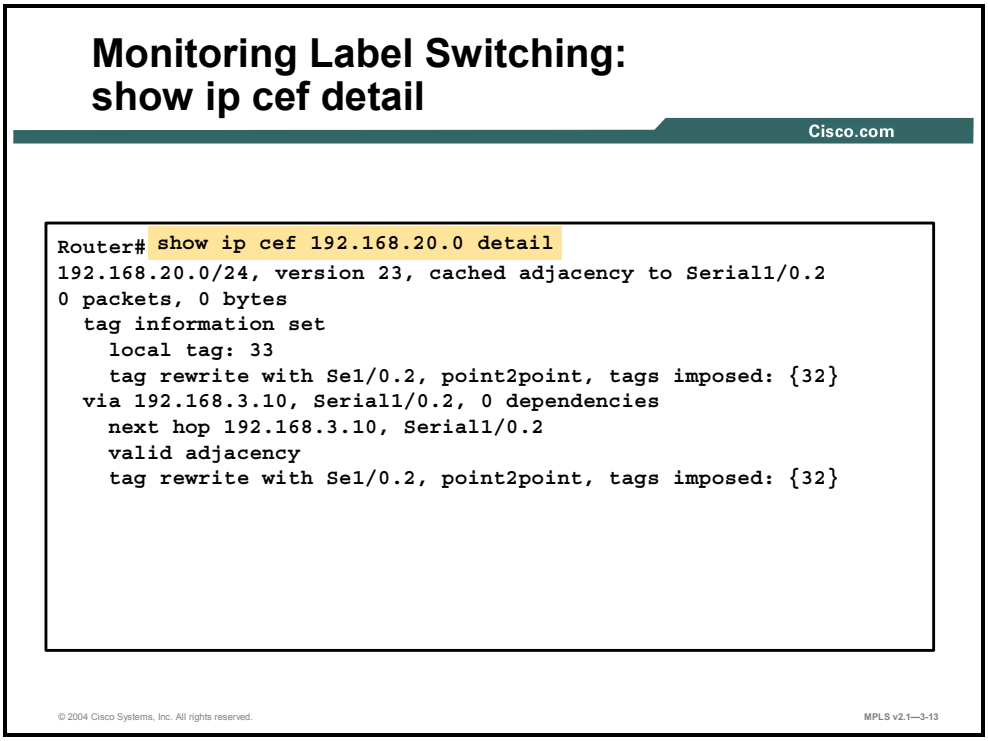

show ip cef detail

To display detailed FIB entry information for all FIB entries, use the following **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 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.

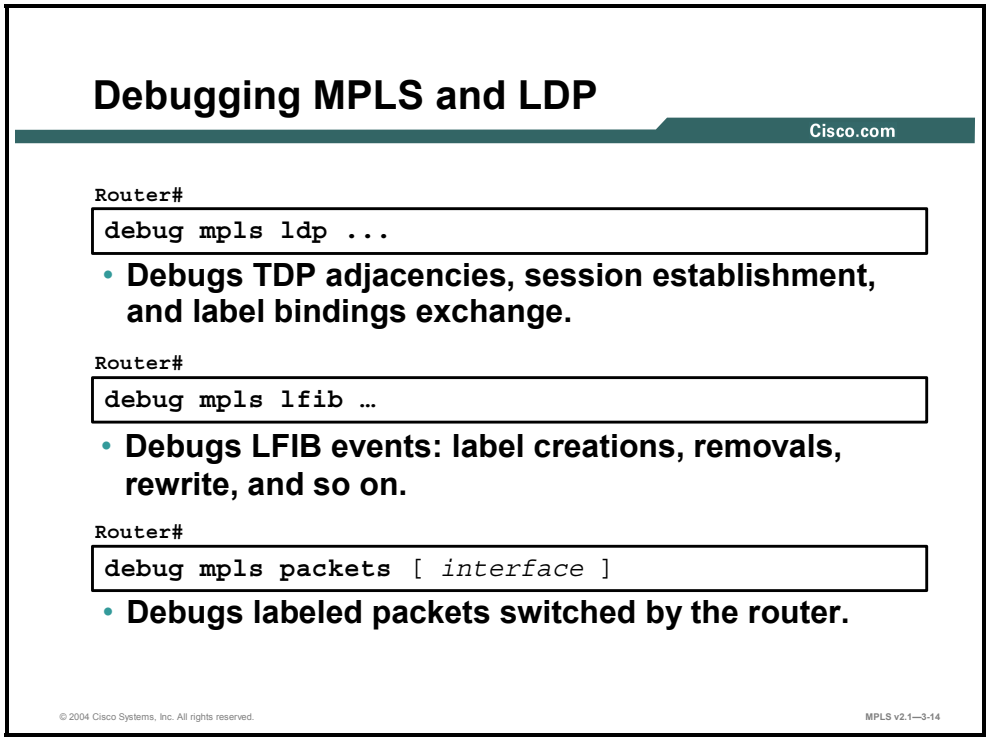

A large number of **debug** commands are associated with MPLS on Cisco IOS platforms. The **debug mpls ldp** set of commands debug various aspects of LDP protocol, from label distribution to exchange of the application-layer data between adjacent LDP-speaking routers.

The **debug mpls lfib** set of 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).

Use this command with care, because it generates output for every packet processed. Furthermore, enabling this 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. The following 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 packet received.
Xmit	Displays packet transmitted.
CoS	Displays class of service (CoS) field from the packet label header.
TTL	Displays TTL field from the packet label header.
(no tag)	Displays 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 discussed in this lesson.

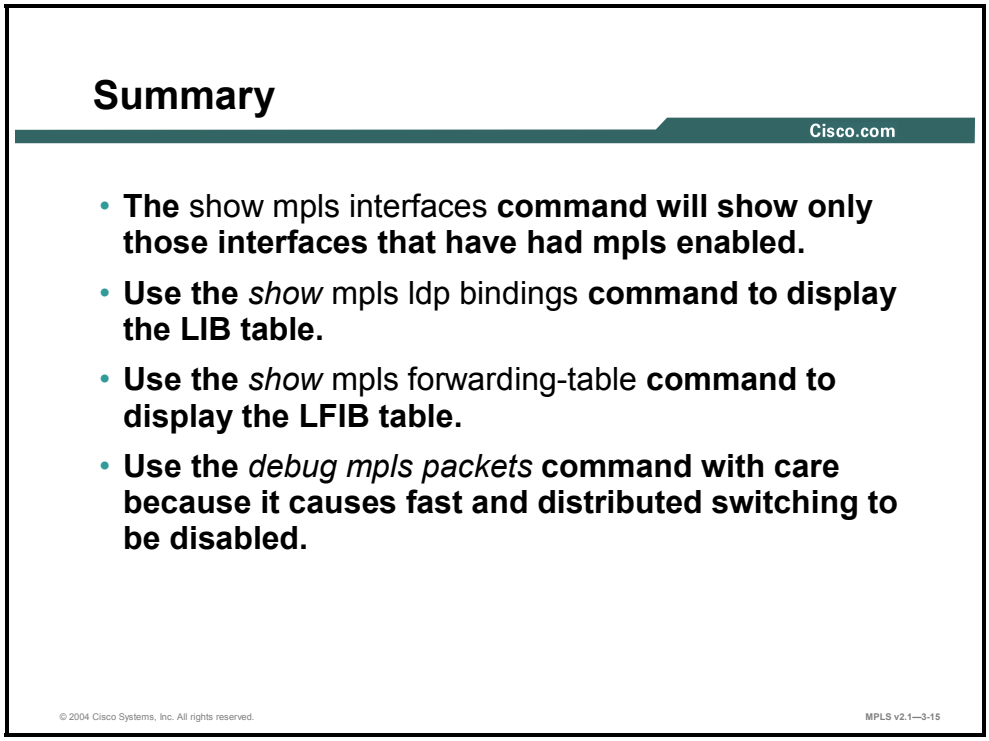

Lesson 4

Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms

Overview

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

It is very important to know what commands that 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
- Describe how to solve LDP session startup issues
- Describe how to solve label allocation issues that can arise in MPLS networks
- Describe how to solve label distribution issues that can arise in MPLS networks
- Describe how to solve packet labeling issues that can arise in MPLS networks
- Describe how to solve intermittent MPLS failures
- Describe how to 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.

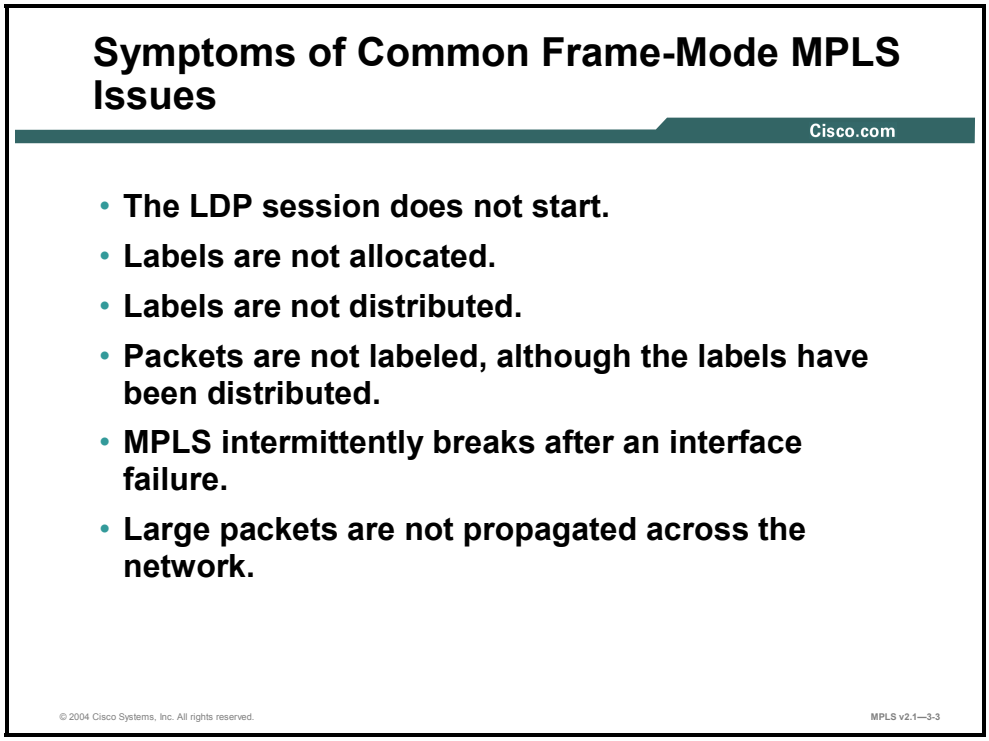

The following describes the common issues that can be encountered while you are troubleshooting a frame-mode MPLS network:

- The 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.

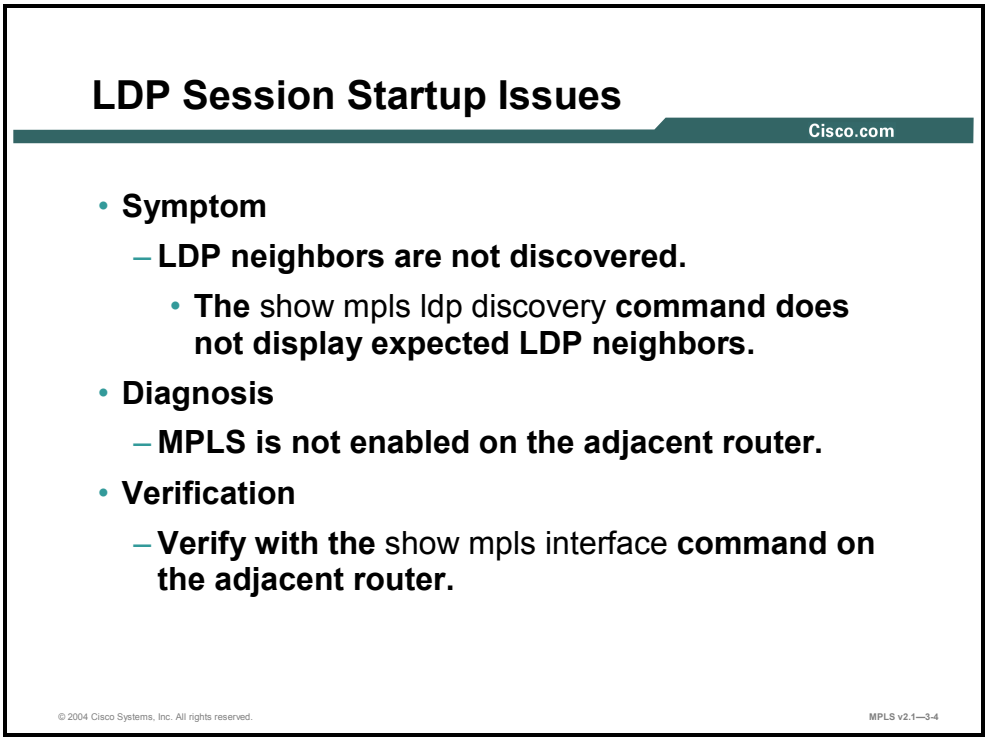

Diagnosis: If MPLS is enabled on an interface, but no neighbors are discovered, it is likely that MPLS is not enabled on the neighbor.

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.

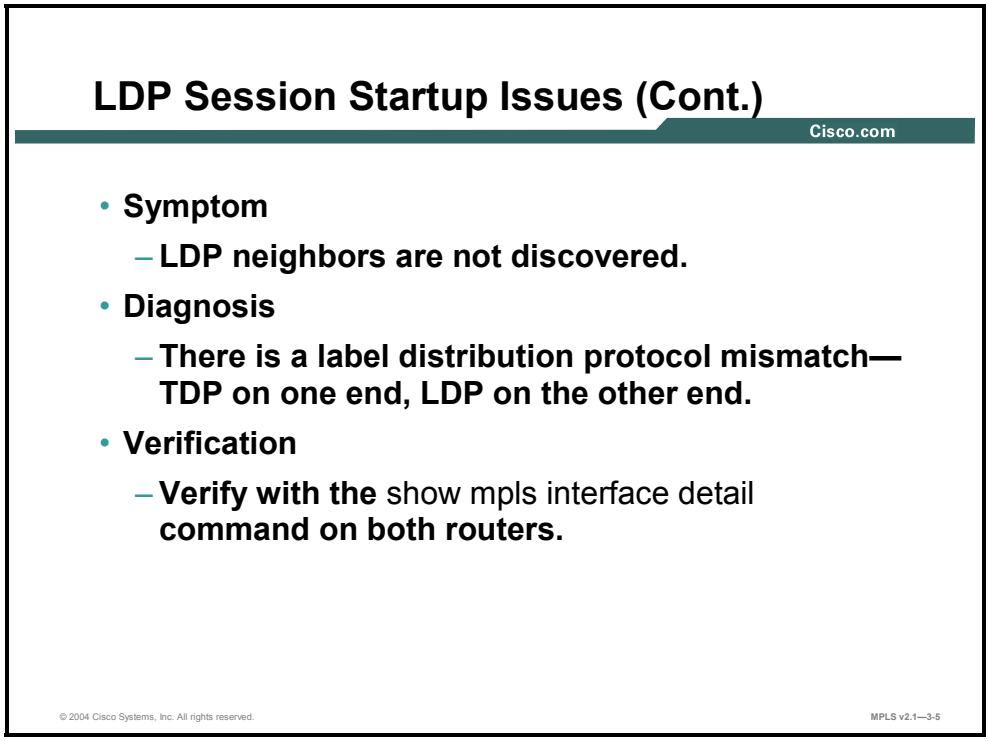

Diagnosis: Another possibility is that the neighbor has a different label distribution protocol enabled on the interface.

Solution: Use one of the following 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.

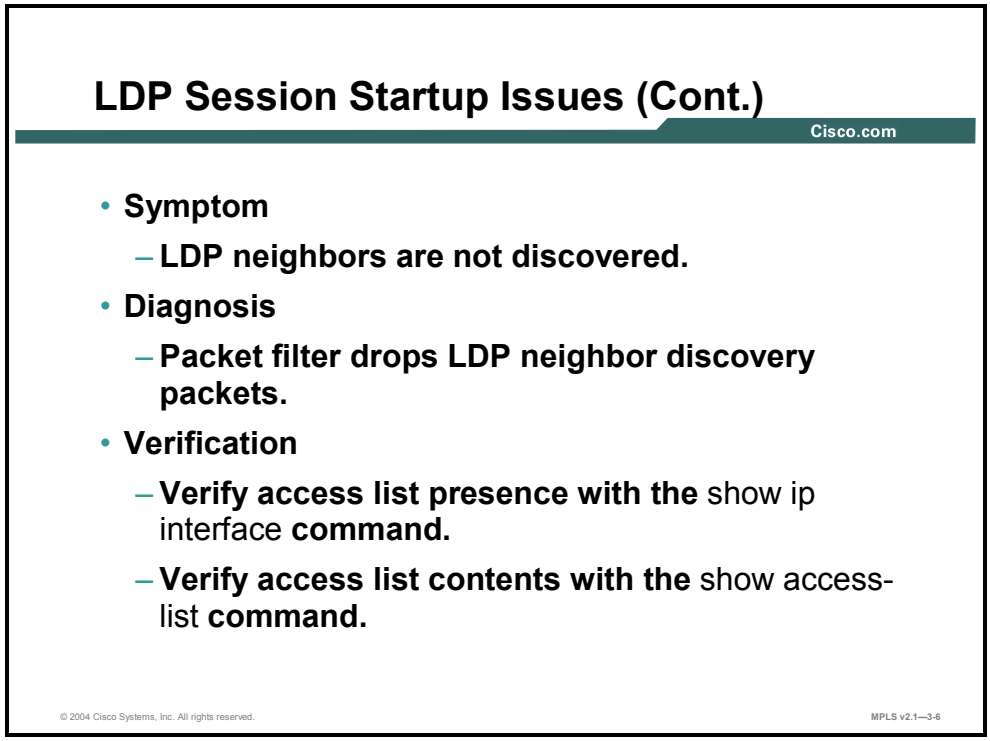

Diagnosis: MPLS configurations match on both ends, but the session still does not get established. Check whether there are any input access lists that deny discovery messages.

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.)

Cisco.com

- **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 IGP used in the network.

Solving Label Allocation Issues

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

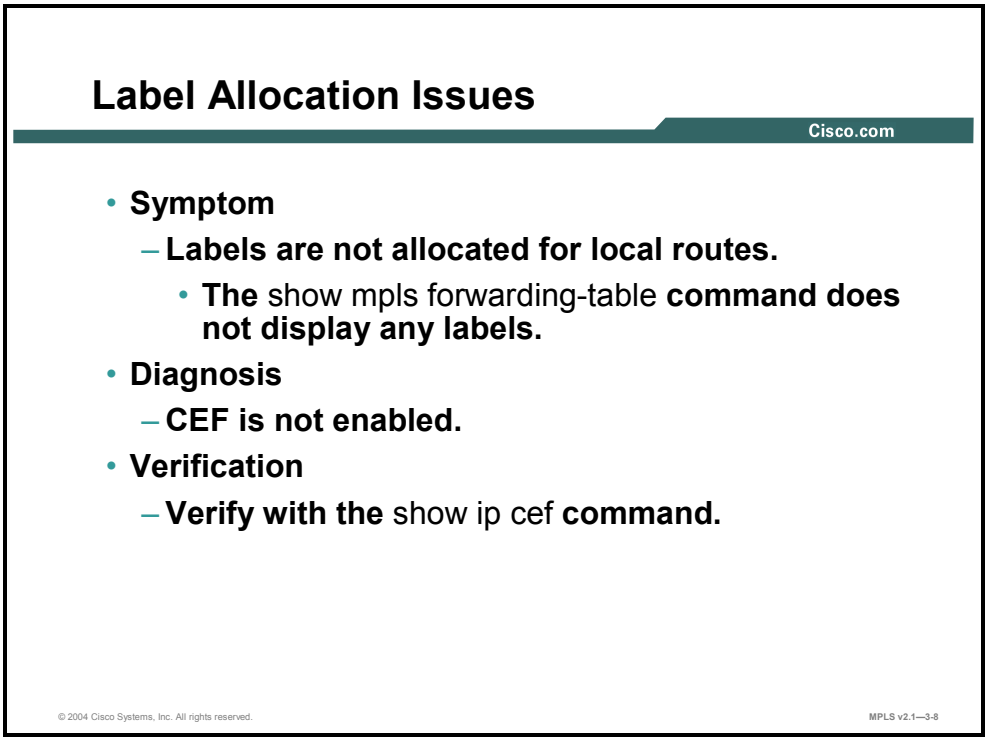

Diagnosis: Labels are not allocated for any or some of the local routes. Use the **show ip cef** command to verify whether CEF switching is enabled on all MPLS-enabled interfaces.

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.

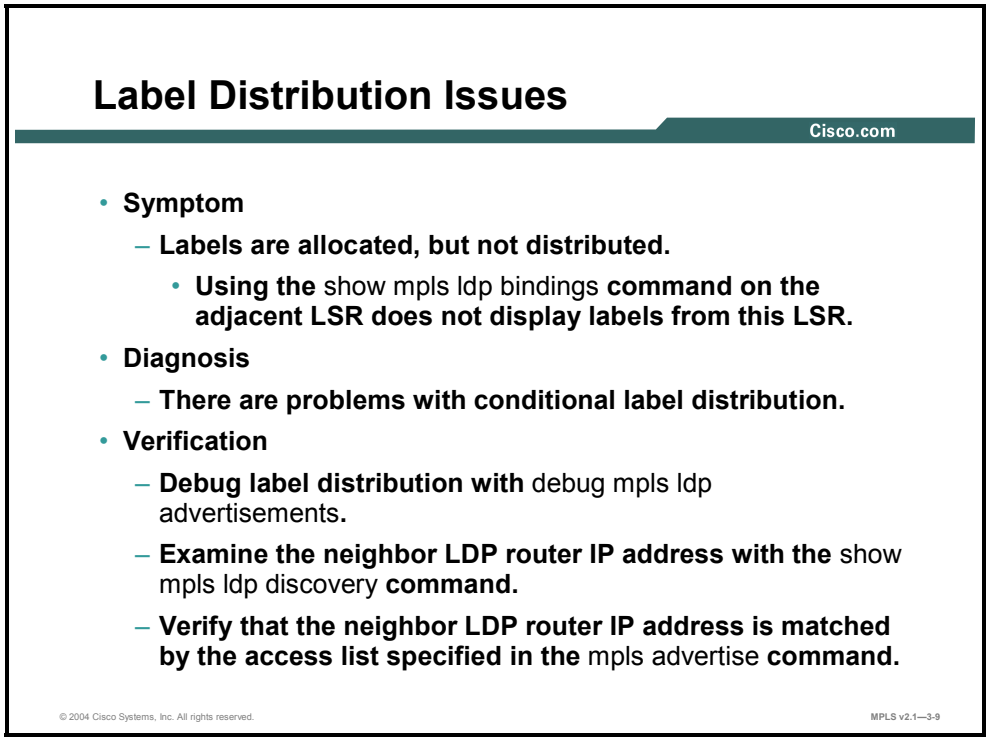

Symptom: Labels are generated for local routes but are not received on neighboring routers.

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.

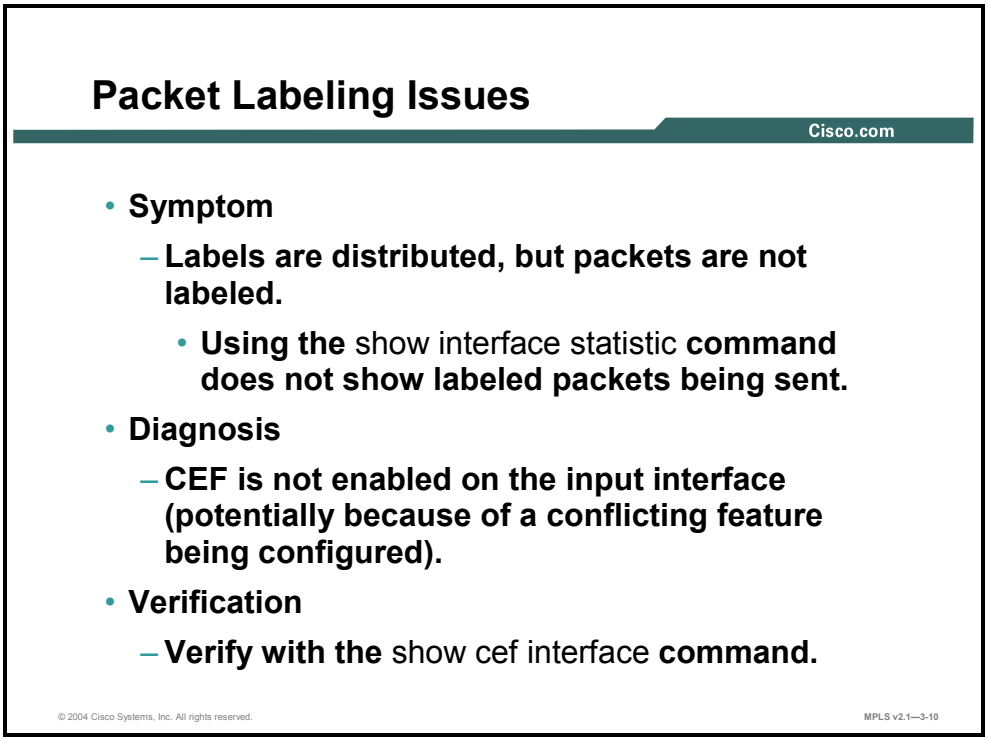

Symptom: Labels exist, but packets are not labeled.

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.

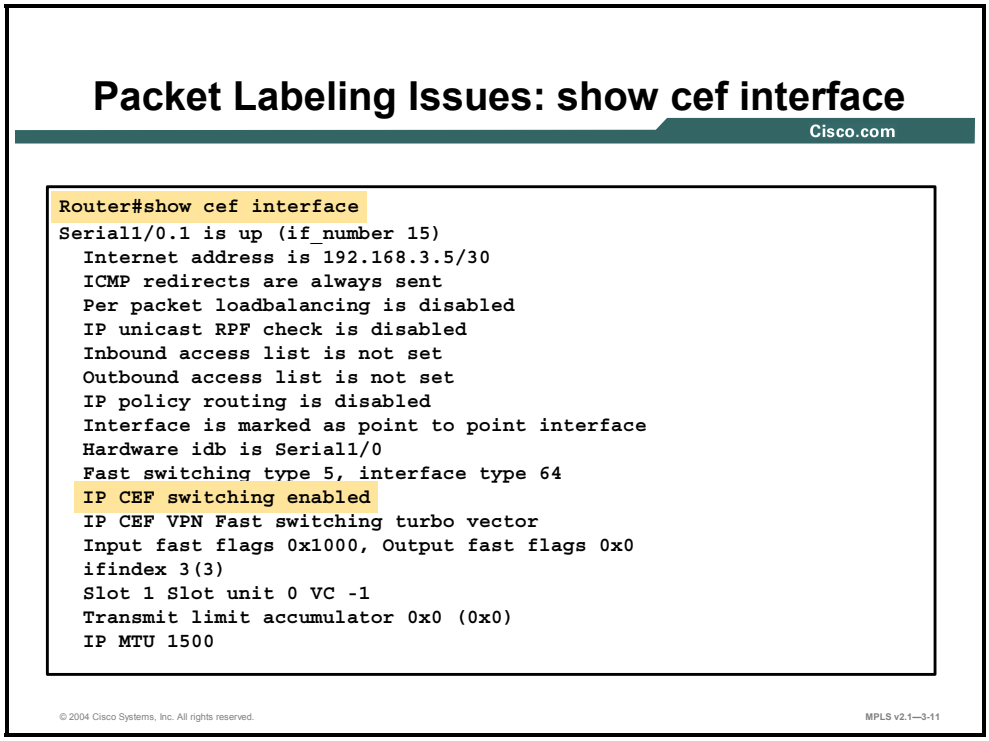

show cef interface

The **show cef interface** command is used to display CEF interface information. The following command is executed in privileged EXEC mode: **show cef interface** *type number* [**detail**].

This table describes the parameters for the **show cef interface** command.

show cef interface Syntax Description

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

This command is available on routers that have route processor (RP) cards and line cards.

The **detail** keyword displays more CEF information for the specified interface.

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

This table describes the significant fields in the display.

show cef interface Field Description

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 MTU size set on the interface.

Solving Intermittent MPLS Failures

This topic describes how to solve intermittent MPLS failures.

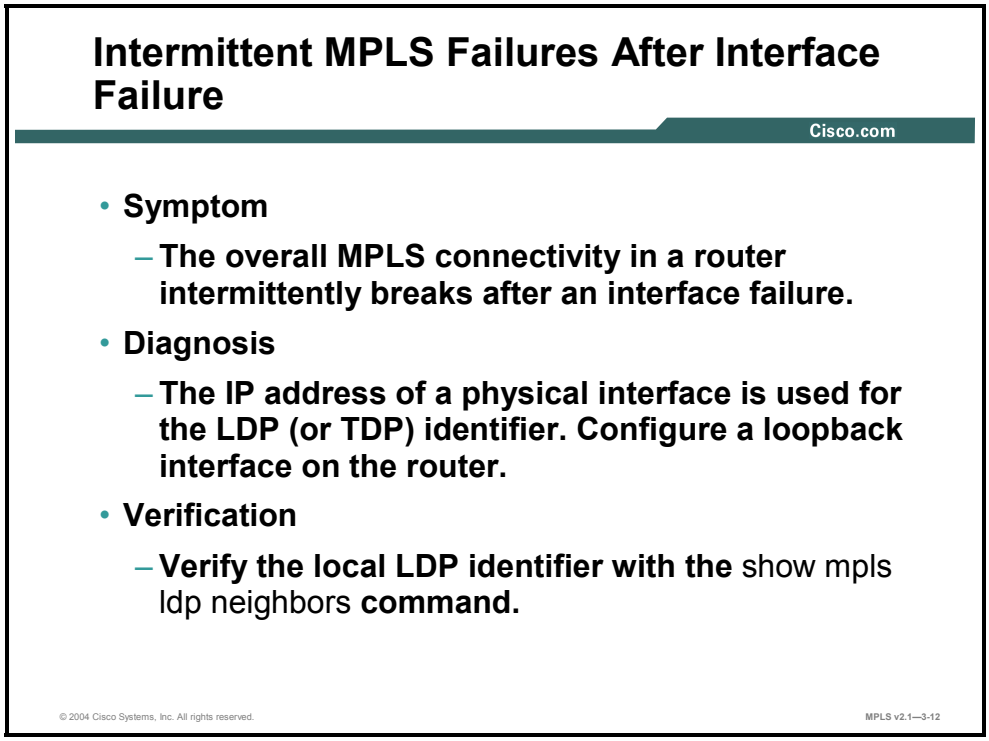

Symptom: MPLS connectivity is established, labels are exchanged, and packets are labeled and forwarded. However, an interface failure can sporadically stop an MPLS operation on unrelated interfaces in the same router.

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

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 identifier 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 identifiers. 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 identifier 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.

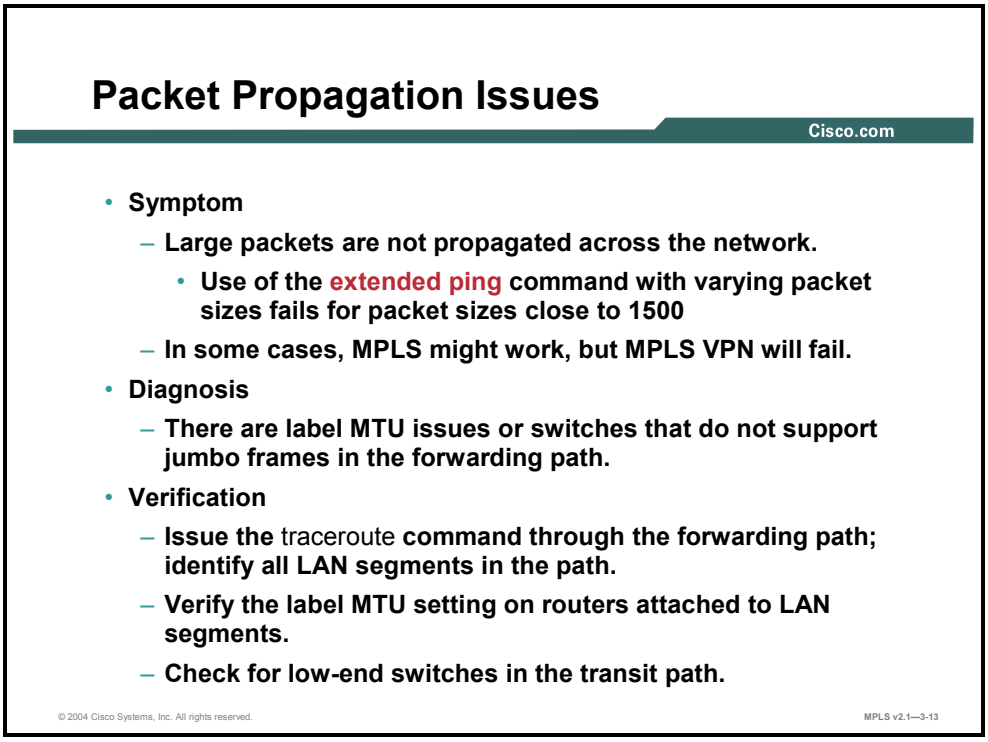

Symptom: Packets are labeled and sent, but they are not received on the neighboring router. A LAN switch between the adjacent MPLS-enabled routers may drop the packets if it does not support jumbo frames.

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 discussed in this lesson.

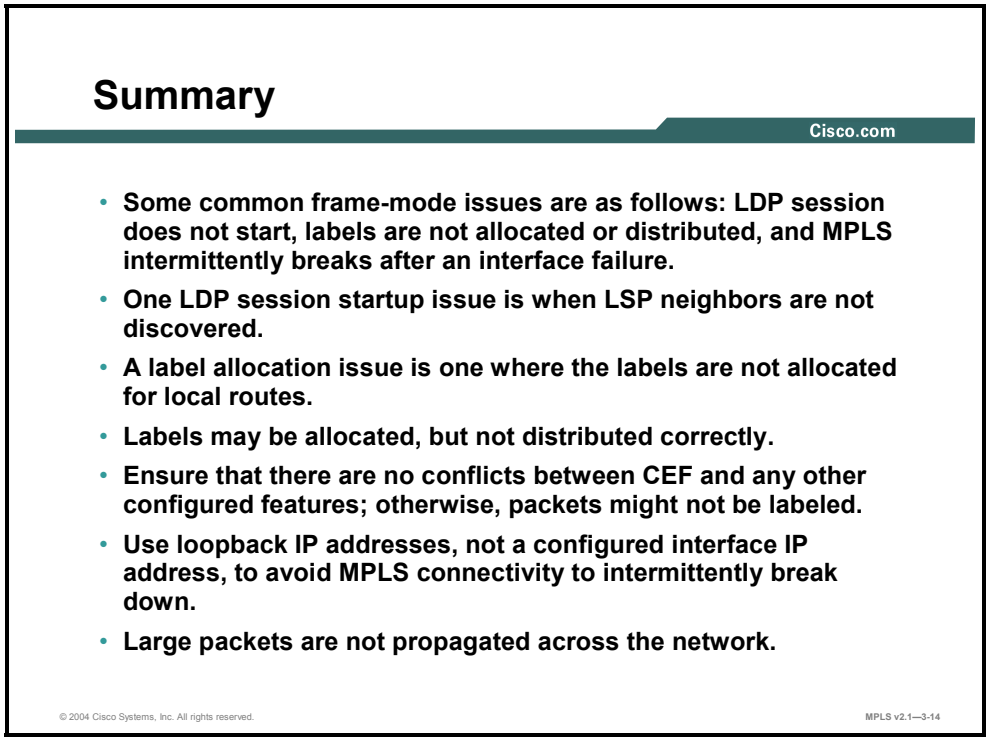

Lesson 5

Configuring LC-ATM MPLS

Overview

This lesson explains how to configure MPLS on router LC-ATM interfaces and Cisco IOS software-based ATM switches. The lesson presents configuration tasks, syntax definitions, and configuration examples.

It is important to understand the differences between frame-based MPLS configuration and cell-based MPLS configuration. This lesson will explain some issues regarding the two technologies and, in particular, how they relate to cell-based MPLS.

Objectives

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

- List the configuration tasks for MPLS on LC-ATM interfacesDescribe how to configure an LC-ATM interface on a routerDescribe how to configure an LC-ATM interface on a Catalyst ATM switchDescribe the guidelines for configuring MPLS between a router and a switch
- Describe some additional LC-ATM parameters that can be configured
- Describe how to disable VC merge

What Are the Configuration Tasks for MPLS on LC-ATM Interfaces?

This topic lists the configuration tasks for configuring MPLS on LC-ATM interfaces.

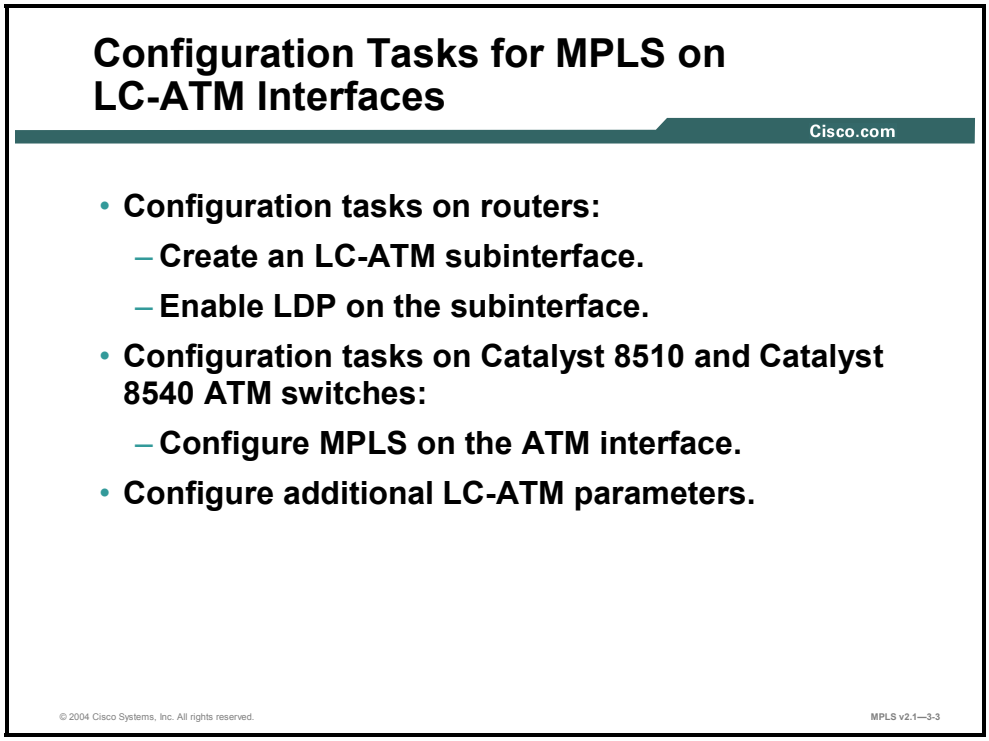

Configuration of cell-mode MPLS differs from configuration of frame-mode MPLS. An additional command specifies the type of subinterface that is to be used.

Instead of enabling a point-to-point or multipoint connection, you set the interface to MPLS mode. (This approach enables cell-mode MPLS instead of the default frame-mode MPLS.)

When the ATM subinterface type is specified, use the MPLS configuration commands to enable MPLS on the interface. MPLS type (cell-mode versus frame-mode) is determined from the type of subinterface.

Note	On ATM switches, there is no need for an additional command because these switches run only cell-mode MPLS.

Configuring an LC-ATM Interface on a Router

This topic describes how to configure an LC-ATM interface on a router.

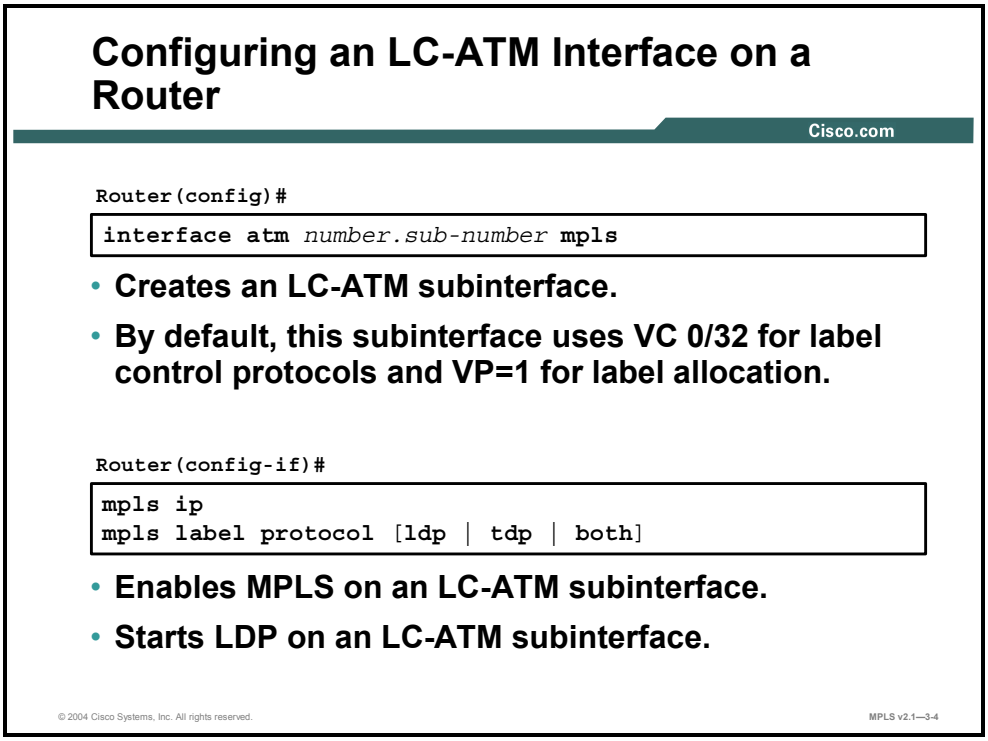

On Cisco IOS platform routers, subinterfaces are typically used. Use the **mpls** keyword to specify the type of subinterface when you are entering interface configuration mode. This command specifies that cell-mode MPLS should be used instead of frame-mode MPLS (which is the default).

Use the **mpls ip** command in configuration mode to enable MPLS.

After the **mpls ip** command is issued, the router creates the control virtual circuit with VPI/VCI=0/32 to establish an IP adjacency with the directly connected ATM switch. This virtual circuit is used for LDP and the routing protocol used in the network.

Optionally, the label distribution protocol can be changed. By default, Cisco routers use TDP. There should be no need to enable both LDP and TDP, because there is only one device on the other side of the link.

To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the **mpls ip** command in interface configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

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

mpls label protocol [tdp | ldp | both]

To specify the label distribution protocol to be used on a given interface, use the **mpls label protocol** command in interface configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

- **mpls label protocol** [**ldp** | **tdp** | **both**]
- **no mpls label protocol** [**ldp** | **tdp** | **both**]

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

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

Parameter	Description
ldp	Specifies use of LDP on the interface.
tdp	Specifies use of TDP on the interface.
both	Specifies use of both label distribution protocols on the interface.

Defaults

TDP is the default protocol.

Configuring an LC-ATM Interface on a Catalyst ATM Switch

This topic explains how to configure an LC-ATM interface on an ATM switch.

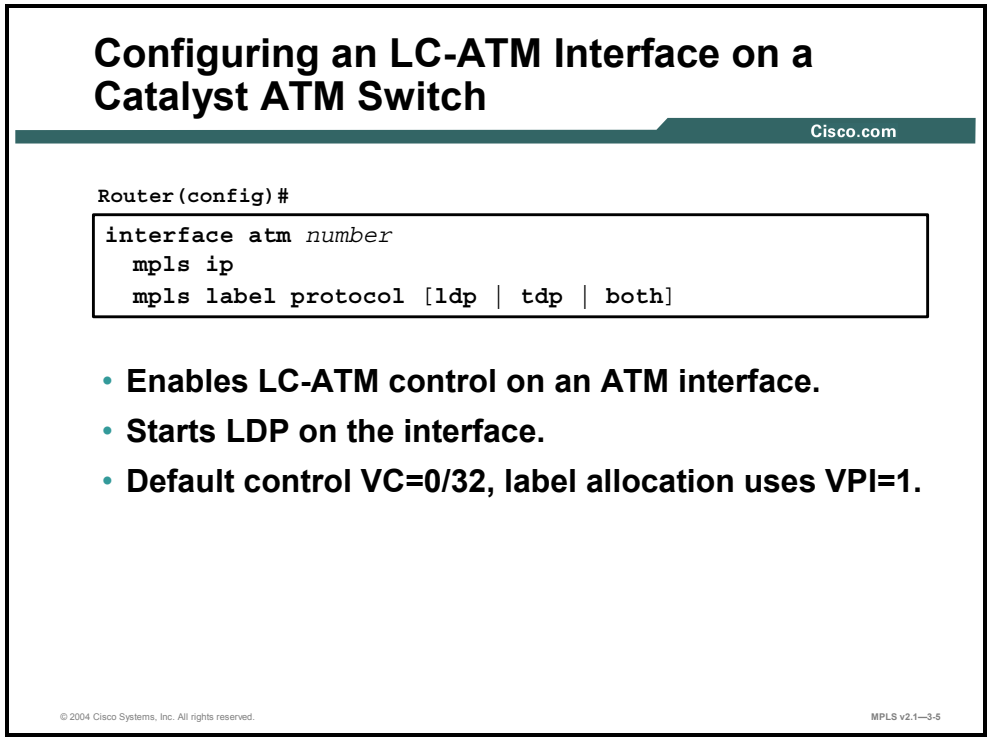

Use these commands to enable MPLS on an interface of a Catalyst ATM switch. Cell-mode MPLS is implied. Enabling both distribution protocols can be useful in a mixed environment when the supported protocol for every device connected to the switch does not need to be determined.

When the LDP or TDP adjacency is established (over virtual circuit 0/32), the devices start negotiating label-switched controlled virtual circuits (LVCs). By default, all LVCs use a VPI value of 1.

mpls ip

To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the **mpls ip** command in interface configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

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

mpls label protocol [tdp | ldp | both]

To specify the label distribution protocol to be used on a given interface, use the **mpls label protocol** command in interface configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

- **mpls label protocol** [**ldp** | **tdp** | **both**]
- **no mpls label protocol** [**ldp** | **tdp** | **both**]

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

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

Parameter	Description
ldp	Specifies use of LDP on the interface.
tdp	Specifies use of TDP on the interface.
both	Specifies use of both label distribution protocols on the interface.

Configuring MPLS Between a Router and a Switch

This topic describes the guidelines for configuring MPLS between a router and a switch.

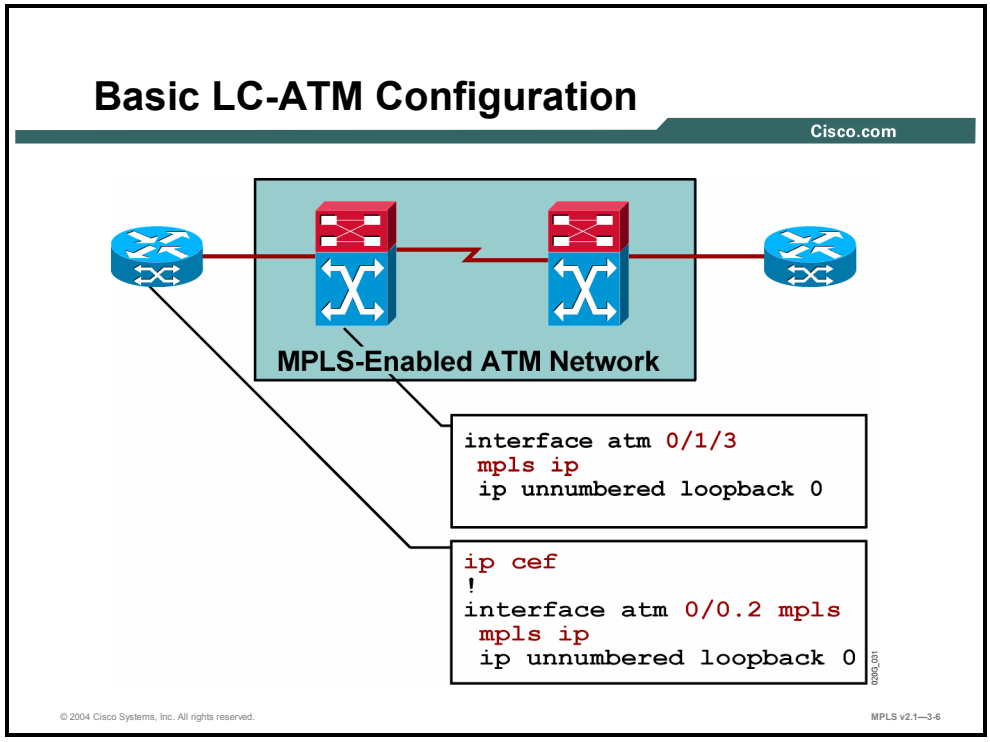

To enable cell-mode MPLS between a router and a switch, ensure that the router uses the MPLS type for the subinterface.

For successful establishment of a label distribution session, both devices need to use the same protocol: LDP (or TDP).

Both devices should use the same parameters for the control virtual circuit (VPI/VCI=0/32). There should be an intersection between the proposed ranges of VPI and VCI values.

By default, all Cisco devices use a VPI value of 1 for dynamically established LVCs.

Additionally, Cisco routers require CEF switching to enable MPLS.

Configuring Additional LC-ATM Parameters

This topic describes some additional LC-ATM parameters that can be configured.

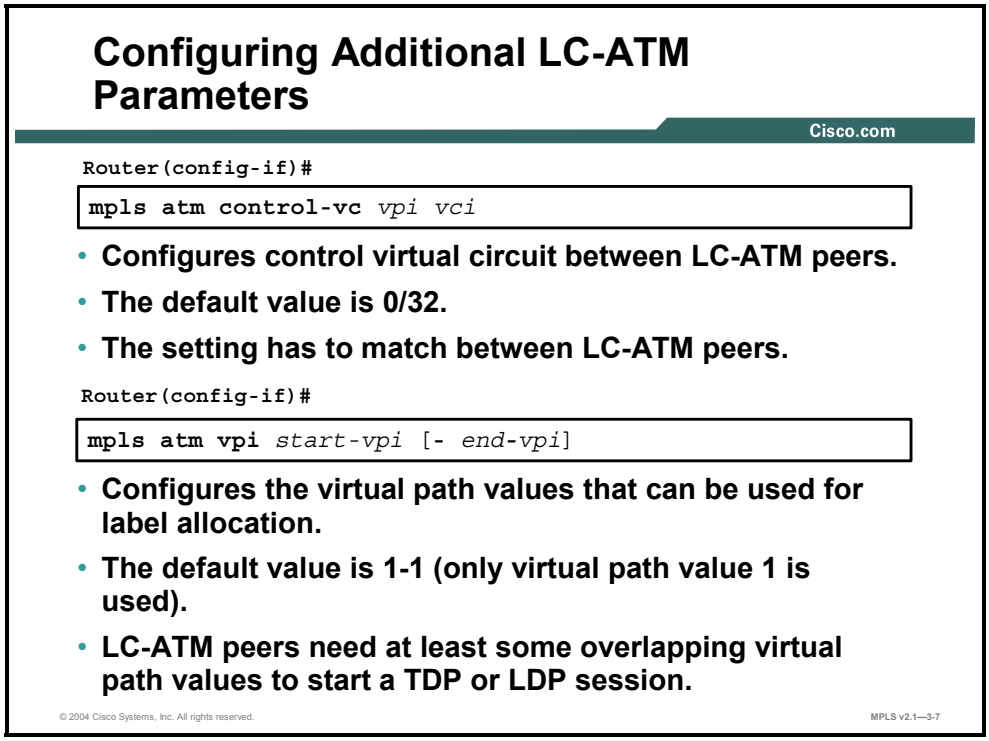

Use the **mpls atm control-vc** command to change the default VPI and VCI numbers used for the control virtual circuit. Use the **mpls atm vpi** command to change the default VPI values for the LVCs.

mpls atm control-vc

To configure VPI and VCI to be used for the initial link to the label-switching peer device, use the **mpls atm control-vc** command in interface configuration mode. The initial link is used to establish the LDP session and to carry non-IP traffic. To clear the interface configuration, use the **no** form of this command. The following illustrates these commands:

- **mpls atm control-vc** *vpi vci*
- **no mpls atm control-vc** *vpi vci*

This table describes the syntax for the **mpls atm control-vc** command.

mpls atm control-vc Syntax Description

Parameter	Description
vpi	Displays VPI.
vci	Displays VCI.

Defaults

If the subinterface has not changed to a virtual path tunnel, the default is 0/32. If the subinterface corresponds to the virtual path tunnel VPI *x*, the default is *x*/32.

mpls atm vpi

To configure the range of values to be used in the VPI field for LVCs, use the **mpls atm vpi** command in interface configuration mode. To clear the interface configuration, use the **no** form of this command. The following illustrates these commands:

- **mpls atm vpi** *vpi* [- *vpi*]
- **no mpls atm vpi** *vpi* [- *vpi*]

This table describes the syntax for the **mpls atm vpi** command.

mpls atm vpi Syntax Description

Parameter	Description
vpi	Displays VPI (low end of range).
- *vpi*	(Optional) Displays VPI (high end of range).

Defaults

The default is 1-1.

Example: Configuring Additional LC-ATM Parameters

The example shows how to change the default VPI range from 1-1 to 5-6. The control virtual circuit can also use the VPI value used for LVCs.

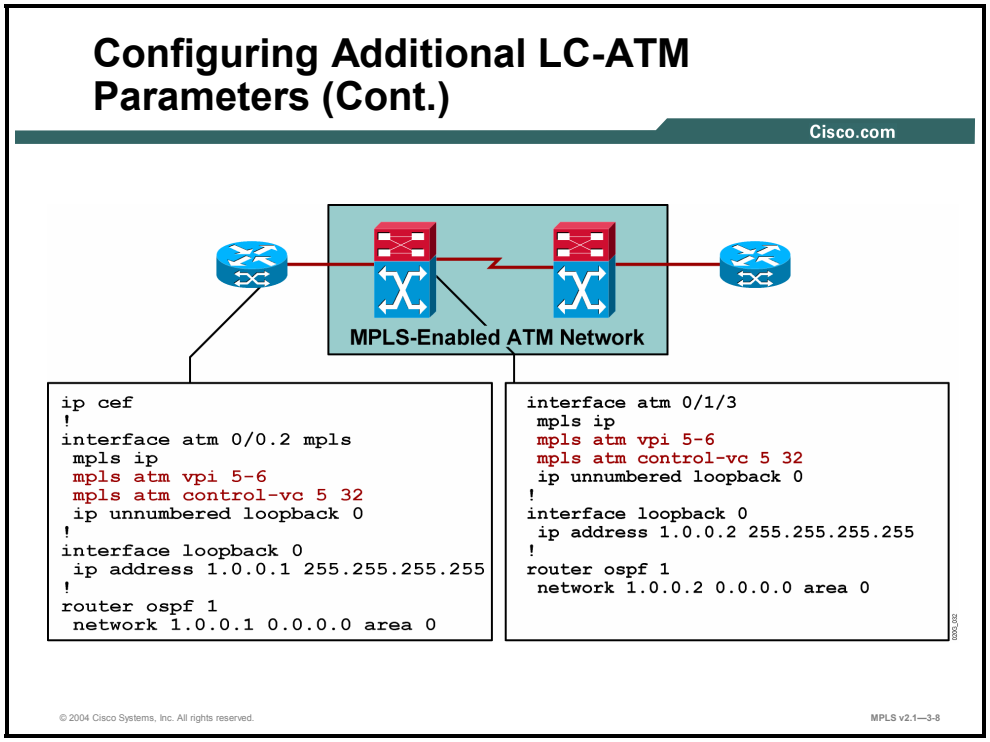

In this example, the control virtual circuit is using VPI=5 and VCI=32. Note that the values must match on each neighbor.

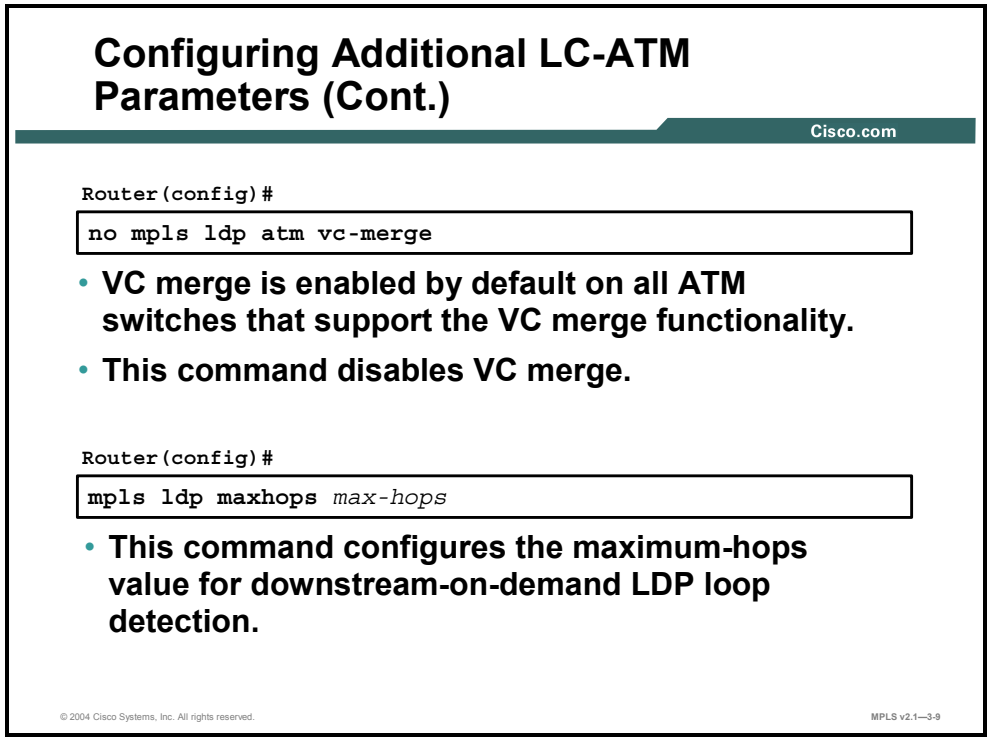

mpls ldp atm vc-merge

To control whether the VC merge (multipoint-to-point) capability is supported for unicast LVCs, use the **mpls ldp atm vc-merge** command in global configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

- **mpls ldp atm vc-merge**
- **no mpls ldp atm vc-merge**

Usage Guidelines

A large ATM network using cell-mode MPLS may experience the problem of having too many LVCs. MPLS itself is very similar to ATM, but it normally merges multiple sources into one destination (label). This is an unusual situation for ATM and can cause mixing of cells belonging to different packets. The end device that needs to reassemble the cells into a packet is not able to differentiate between cells, because the cells use the same VPI/VCI value pair. The following describes the two solutions:

- Create a distinct label for every source-destination pair (may require a large number of LVCs).
- Merge multiple sources to use the same destination label, by buffering the incoming cells in the ATM switch and forwarding them when the complete frame has been assembled. This option is called VC merge.

VC merge is *enabled by default* on all devices that support it, and must be explicitly disabled if it is not desired.

Note The ATM switch that does the VC merge function buffers the entire ATM adaptation layer 5 (AAL5) frame as the individual cells are received and then forwards them contiguously, without mixing cells. The end device, therefore, has no problem reassembling each individual frame correctly. The drawback of using VC merge is the increased store-and-forward delay incurred by the ATM switch.

mpls ldp maxhops

To limit the number of hops permitted in an LSP established by the downstream-on-demand method of label distribution, use the **mpls ldp maxhops** command in global configuration mode. To disable this feature, use the **no** form of this command. The following illustrates these commands:

- **mpls ldp maxhops** *number*
- **no mpls ldp maxhops**

This table describes the syntax for the **mpls ldp maxhops** command.

mpls ldp maxhops Syntax Description

Parameter	Description
number	Displays number from 1 to 255, inclusive, that defines the maximum hop count. The default is 254.

Usage Guidelines

When an ATM LSR initiates a request for a label binding, it sets the hop-count value in the label request message to 1. Subsequent ATM LSRs along the path to the edge of the ATM label-switching region increment the hop count before forwarding the label request message to the next hop.

When an ATM LSR receives a label request message, it does not send a label-mapping message in response, and it does not propagate the request to the destination next hop if the hop in the request equals or exceeds the maximum-hops value. Instead, the ATM LSR returns an error message that specifies that the maximum allowable hop count has been reached. This threshold is used to prevent forwarding loops in the setting up of LSPs across an ATM region.

Disabling VC Merge

This topic describes how to disable VC merge.

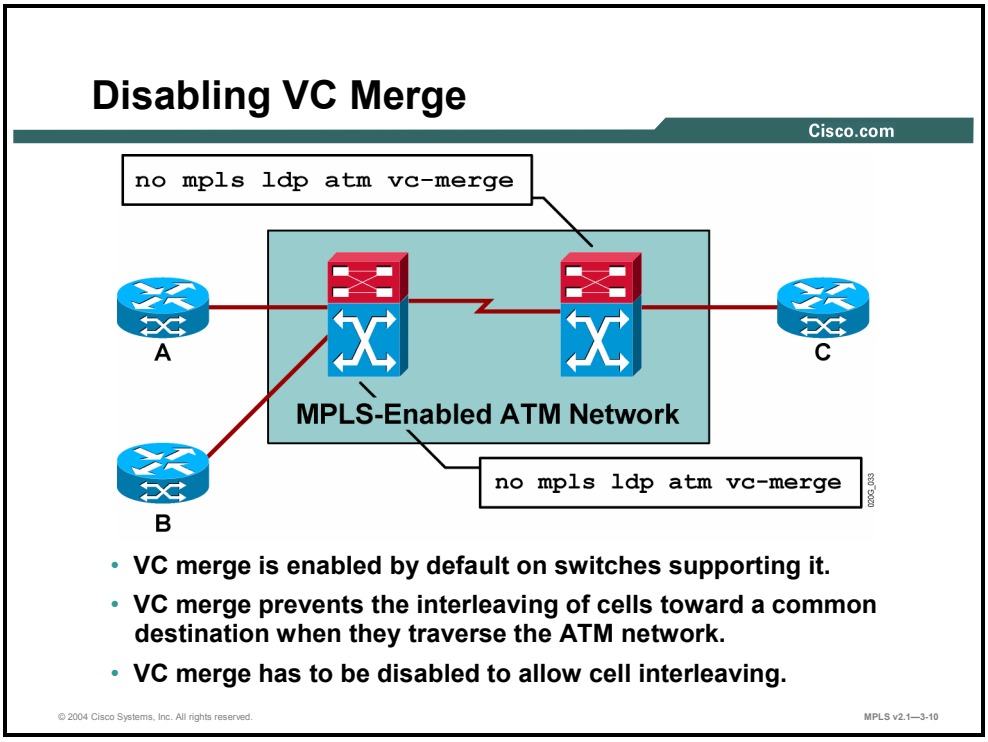

The VC merge feature is enabled by default on all switches that support it. If the feature is not required (that is, because of a small network, different line speeds, or buffering not desired), it can be disabled.

Disabling VC merge results in the ability to interleave cells, but an LVC must be created for every source-destination pair.

Summary

This topic summarizes the key points discussed in this lesson.

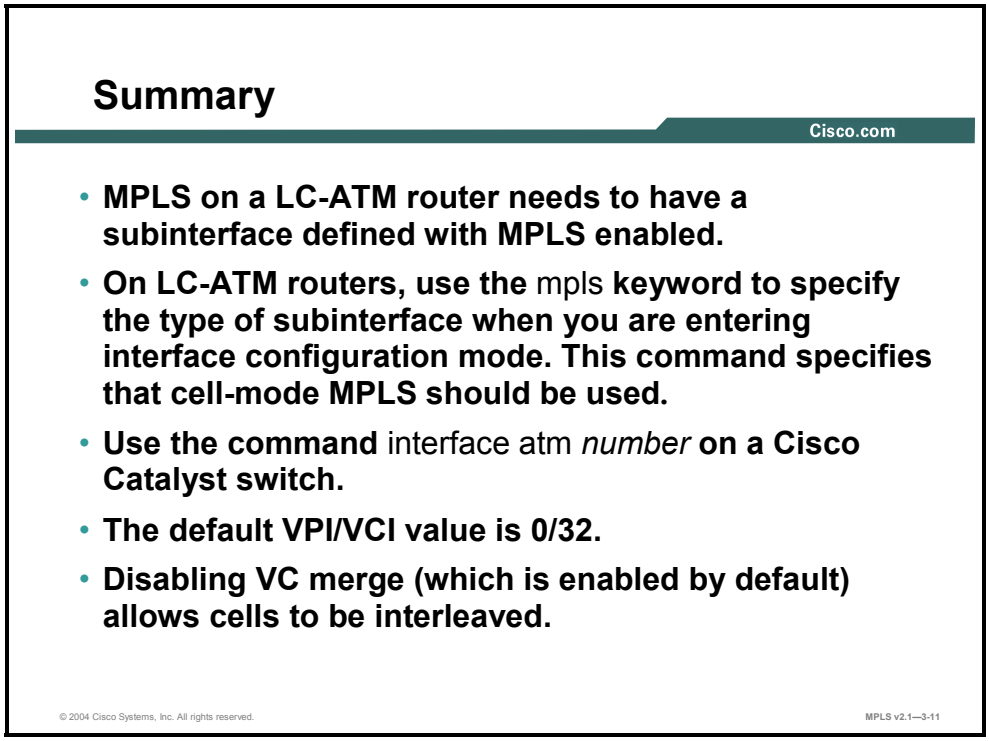

Lesson 6

Configuring LC-ATM MPLS over ATM Virtual Path

Overview

This lesson explains what ATM Virtual Path (ATM VP) is and why it might be used. Also, the configuration of ATM VP for both routers and switches is covered in this lesson.

This lesson explains what to do when an MPLS network must travel across an ATM network that does not support MPLS. This situation is somewhat typical when you are migrating from a standard ATM network to an IP+ATM network, or when the need arises to connect sites across a public ATM network.

Objectives

Upon completing this lesson, you will be able to describe how to configure LC-ATM MPLS over ATM VP. This ability includes being able to meet these objectives:

- Describe the function of ATM VP
- Describe how ATM VP can be used
- Describe how to configure MPLS over ATM VP for switches
- Describe how to configure MPLS over ATM VP for routers

What Is ATM Virtual Path?

This topic describes the function of ATM VP.

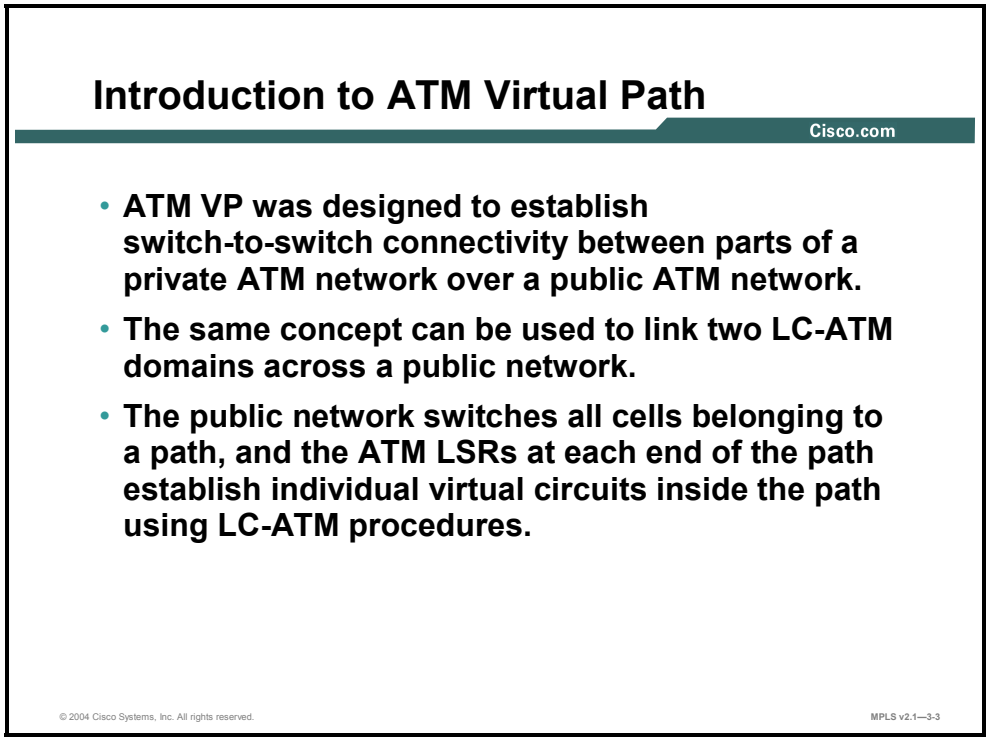

A virtual path is a collection of virtual circuits with a common Virtual Path Identifier (VPI).

ATM switches forward cells based on the VPI only (the VCI is ignored). This approach is useful if one or more switches in the network do not support MPLS.

A static virtual path can be established between switches that support MPLS. Switches can establish a control virtual circuit across the virtual path and negotiate LVCs with the virtual path VPI used to set the label range.

This solution is typically used when a public ATM network interconnects remote sites that use ATM switches.

ATM Virtual Path Usages

This topic describes how ATM VP can be used.

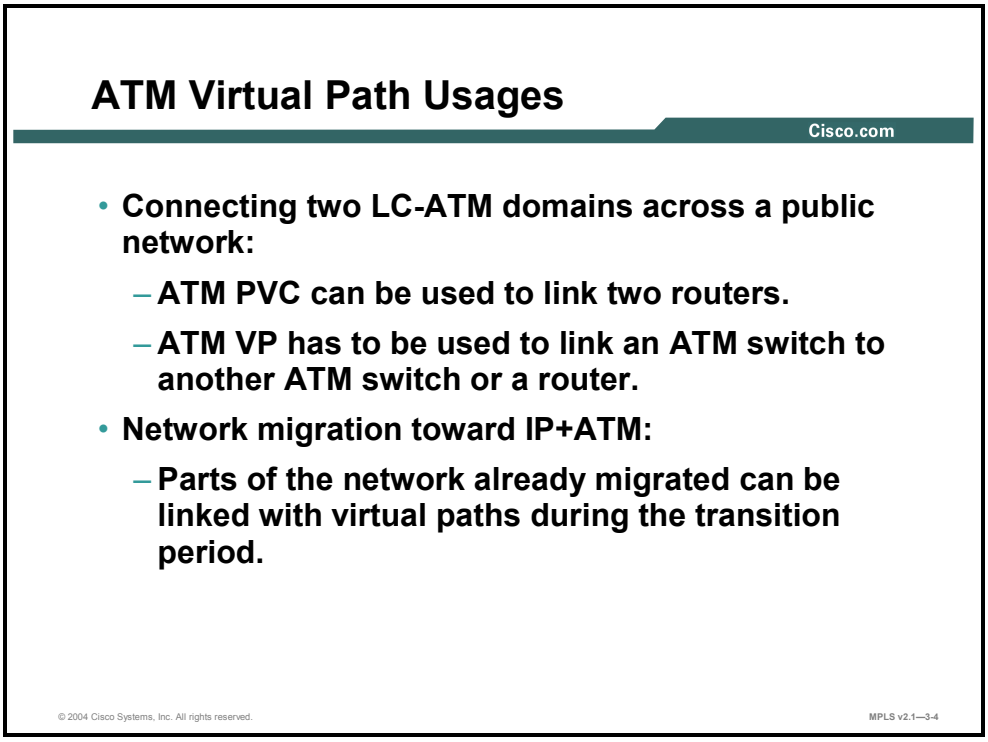

The following two options are available to enable two MPLS domains across a public ATM network:

- **Virtual circuit:** Frame-mode MPLS has to be used because ATM switches in the path do not support MPLS. Only routers support frame-mode MPLS. Switches cannot use frame-mode MPLS and, therefore, cannot use virtual circuits.
- **Virtual path:** Cell-mode MPLS can be used between routers or switches on both ends of the virtual path.

Virtual paths can also be used in the migration when sites are being reconnected to MPLS-enabled switches.

Virtual paths can be established from an MPLS-enabled switch to all devices connected to ATM switches that do not support MPLS. The network can then slowly be migrated toward IP+ATM without the need for an “overnight” full migration.

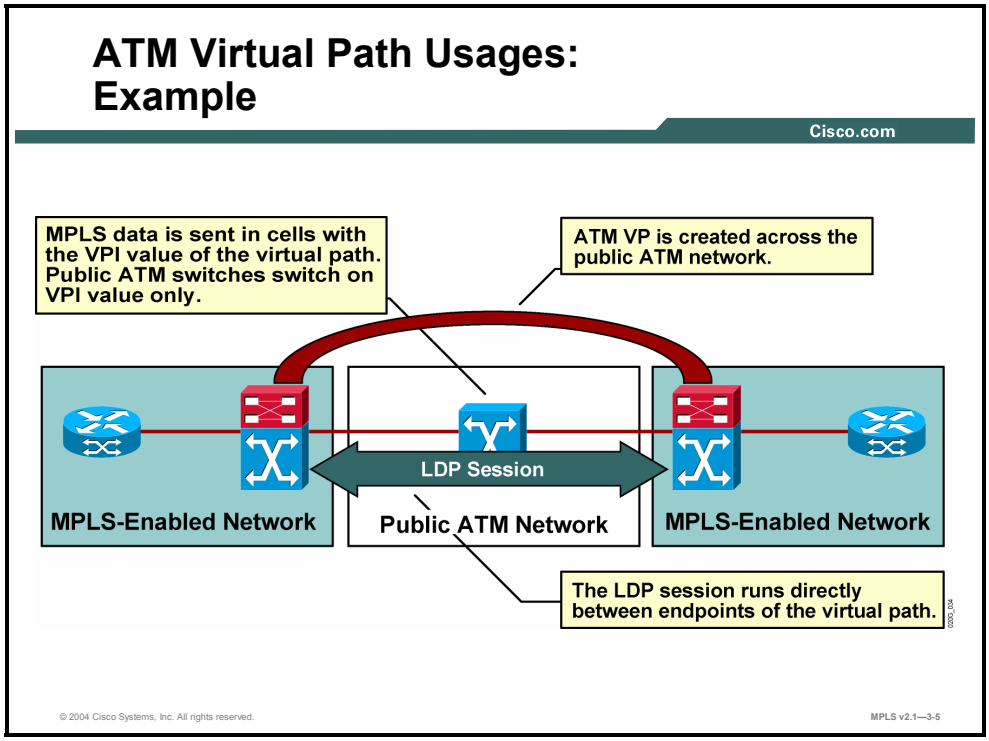

To enable cell-mode MPLS across a virtual path, the control virtual circuit *should use* the VPI of the virtual path.

A router or a switch will then establish an adjacency with a router or a switch on the other end of the virtual path.

It is *mandatory* that the same VPI be used on both ends of the path, because the VPI is part of the LDP virtual path range negotiation.

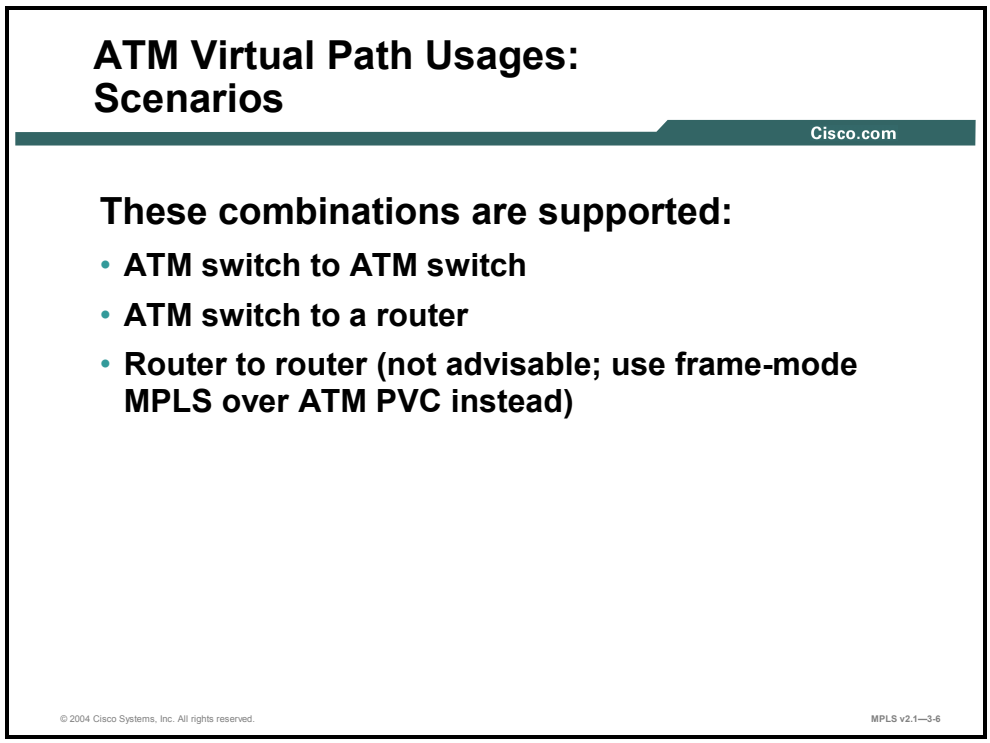

The following describes how a virtual path can be used to connect any pair of devices across a public ATM network:

- Switch to switch
- Switch to router
- Router to router (PVCs with frame-mode MPLS are usually used in this case.)

The first two options allow MPLS to run across a public ATM network.

The third option can also be used, but it has no advantage over using frame-mode MPLS across PVCs. However, the router-to-router solution requires a reservation of a large number of virtual circuits. (A virtual path carries 65,536 virtual circuits.)

Configuring MPLS over ATM Virtual Path—Switches

This topic describes how to configure MPLS over ATM VP for switches.

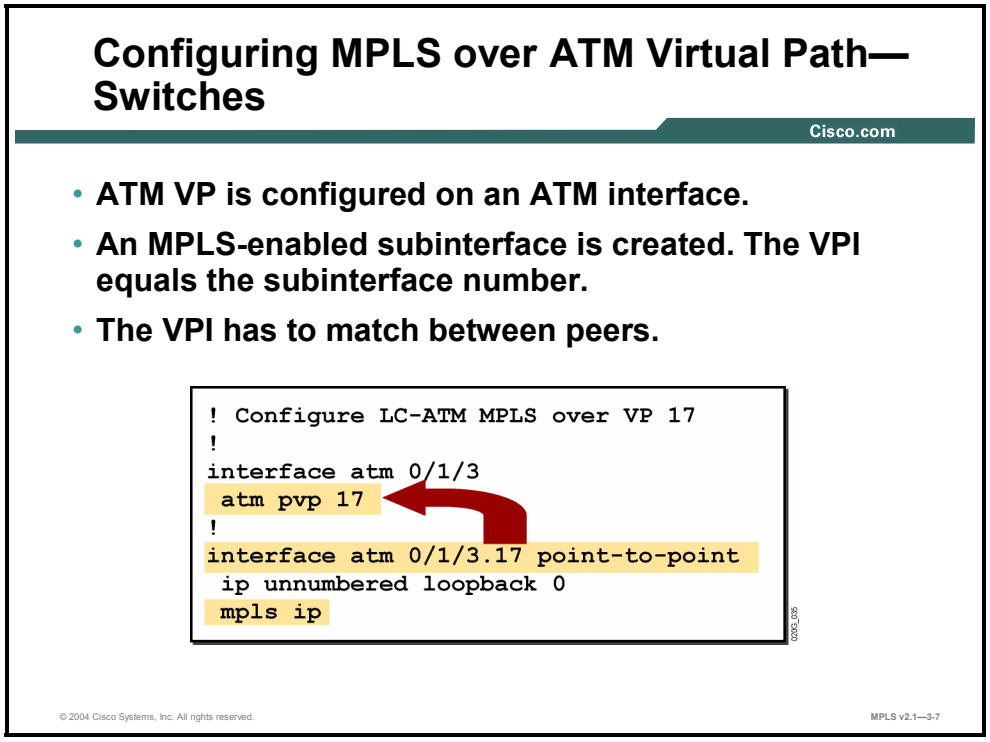

A subinterface is configured with the VPI, which equals the subinterface number and has cell-mode MPLS functionality.

Example: Configuring MPLS over ATM Virtual Path—Switches

In the figure, a virtual path with a VPI of 17 is created.

Note	The VPI has to match between peers.

Example: Configuration of Both MPLS-Enabled ATM Switches

This figure shows the configuration of both MPLS-enabled ATM switches connected by a virtual path across a public ATM network.

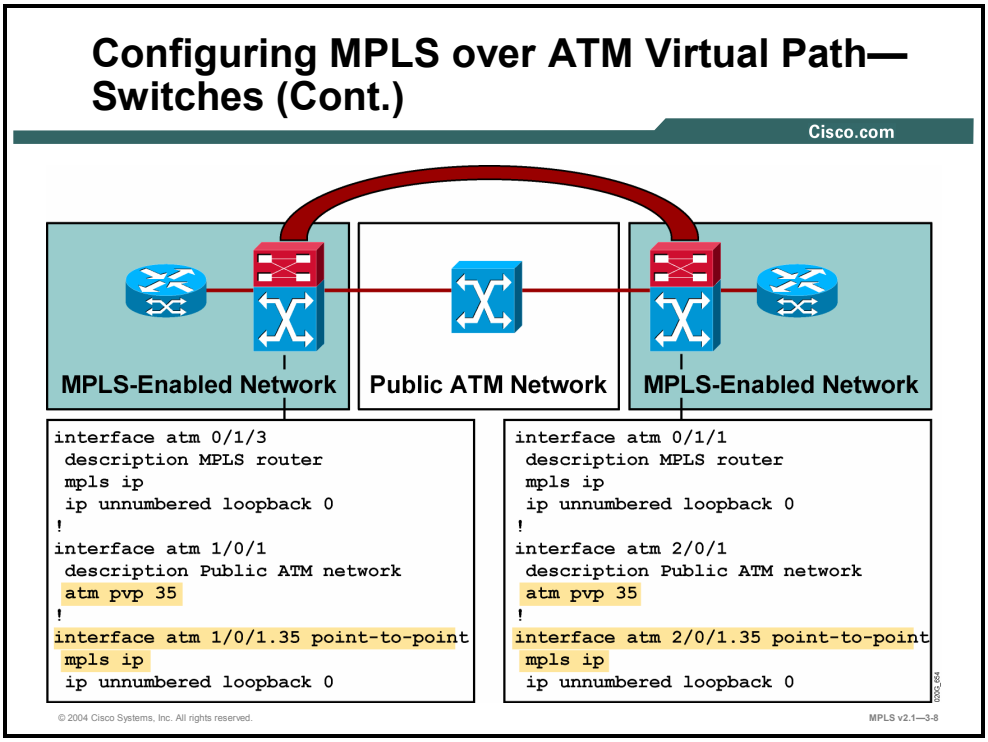

The VPI has to be the same on the first and last hop in the path. The ATM provider can use any VPI on any other link.

The example shows that the subinterface that is created, on both switches, has a subinterface number equal to the VPI.

Note	The example does not change the parameters of the control virtual circuit. PVCs will need to be established for the control virtual circuit (0/32).

Configuring MPLS over ATM Virtual Path—Routers

This topic describes how to configure MPLS over ATM VP for routers.

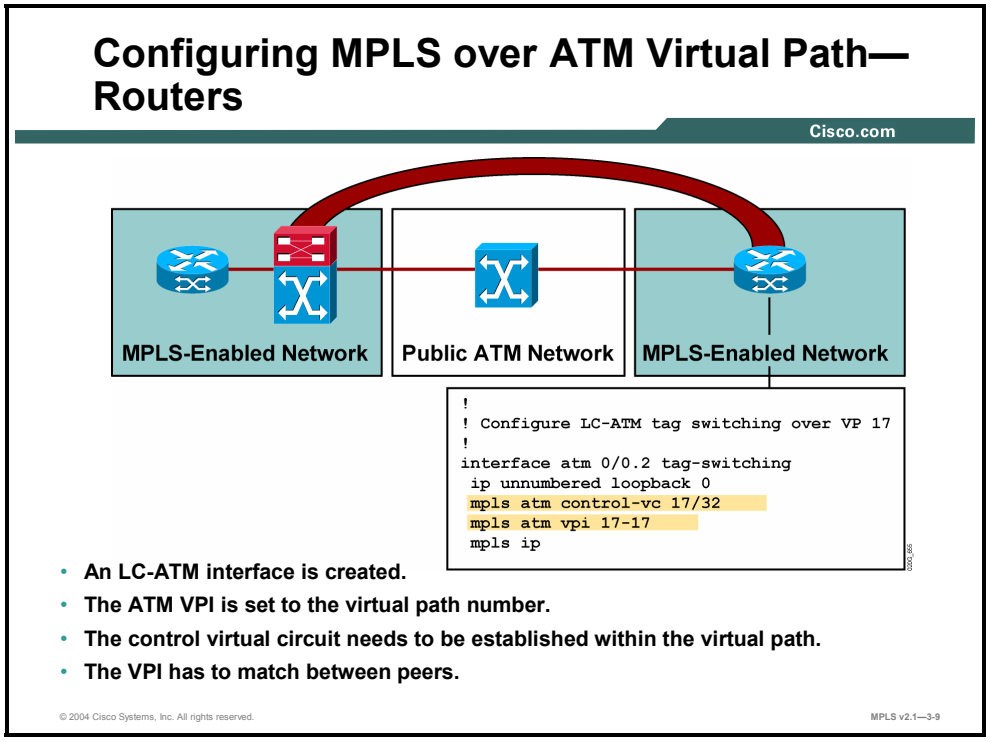

To simplify the provisioning of the connection across a public ATM network, you can also put the control virtual circuit into the virtual path.

Example: Configuring MPLS Over ATM Virtual Path—Routers

The figure shows how to change the control virtual circuit to use the same VPI used to establish the virtual path.

If the public network is forwarding cells for VPI=17, the control virtual circuit should be put into this virtual path (17/32) and the label range has to be set to use the same VPI (17-17).

Note	The VPI has to match between peers.

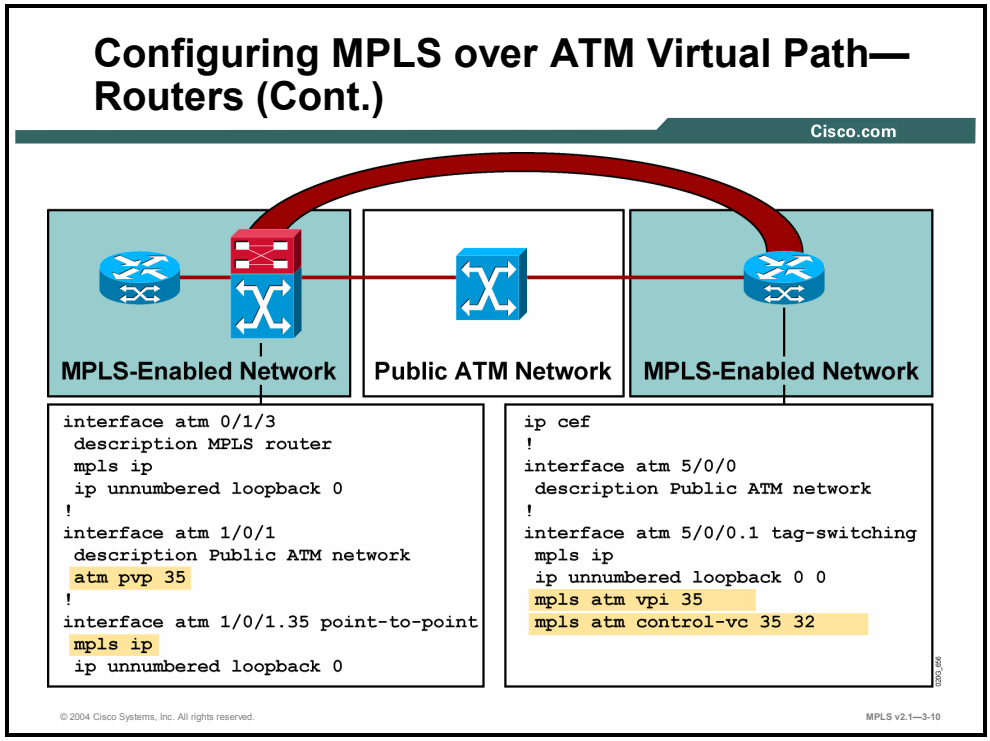

When you connect a router and a switch through a virtual path, you need to set only the parameters for the control virtual circuit and the label range on the router.

The router is unaware that the control virtual circuit is not terminated on the directly connected switch. The public ATM network simply forwards all cells based on the VPI to the other endpoint, where an MPLS-enabled switch continues forwarding based on VPI *and* VCI values.

Summary

This topic summarizes the key points discussed in this lesson.

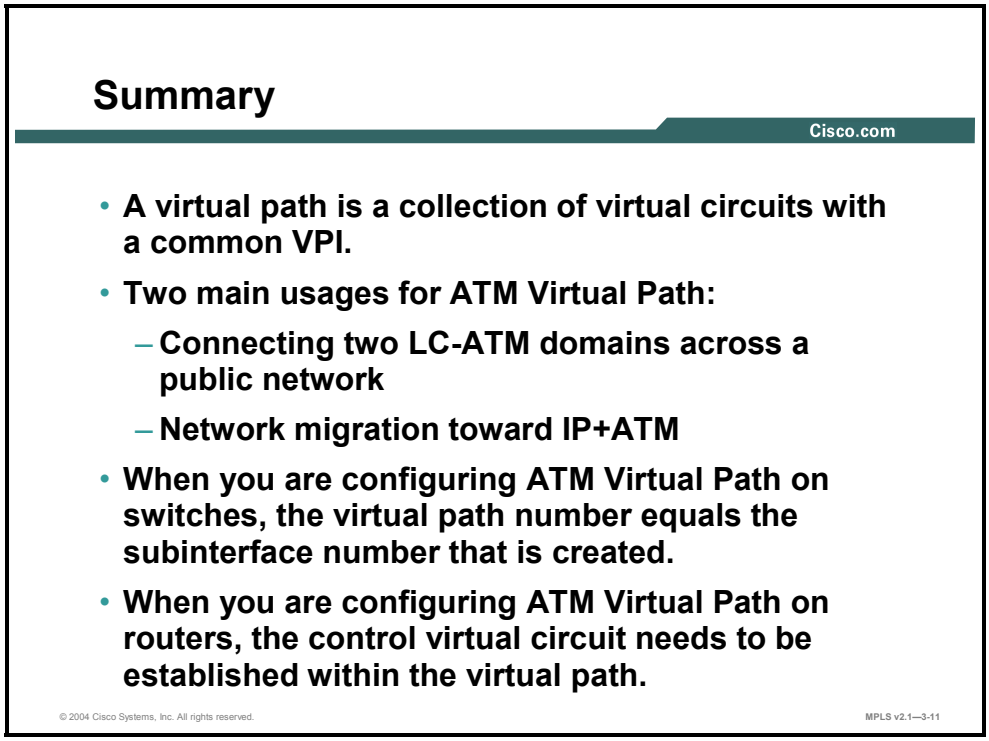

Lesson 7

Monitoring LC-ATM MPLS on Cisco IOS Platforms

Overview

This lesson describes the commands that are used to monitor LC-ATM functions, including command syntax, definitions, and examples.

It is important to understand the network that you have just configured. This lesson will help when you are looking at LC-ATM connections in your network and verifying that the network is running smoothly. The lesson will also help you to identify and isolate problems with the network.

Objectives

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

- Describe how to monitor specific LC-ATM label-switching functions
- Describe how to display summary information about all the entries in the ATM label-binding database
- Describe how to display current label bindings
- Describe how to display MPLS ATM capabilities negotiated by LDP
- Describe how to debug ATM LDP functions

How to Monitor Specific LC-ATM Label-Switching Functions

This topic describes how to monitor specific LC-ATM switching functions.

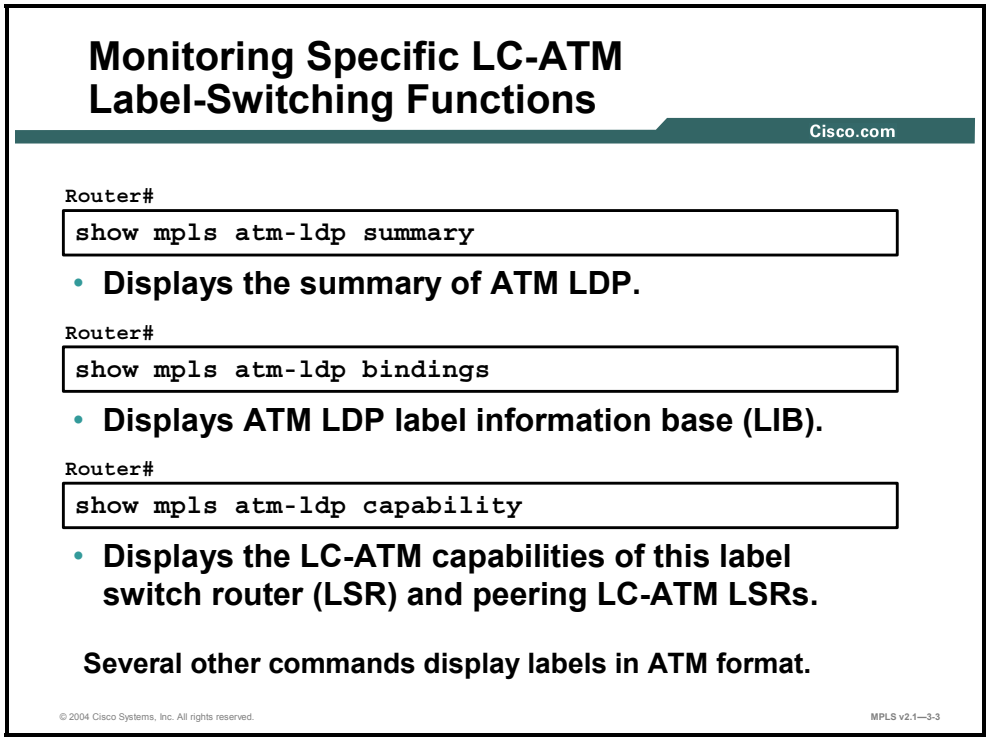

The commands are similar to **show mpls ldp** commands, except the **show mpls atm-ldp** commands display ATM specific parameters. Use a question mark to see all of the subcommands or use the **show mpls atm ldp** command.

show mpls atm-ldp summary

To display summary information about all of the entries in the ATM label-binding database, use the following command in privileged EXEC mode: **show mpls atm-ldp summary**.

show mpls atm-ldp bindings

To display specified entries from the ATM label-binding database, use the **show mpls atm-ldp bindings** command in privileged EXEC mode. The ATM label-binding database contains entries for LVCs on LC-ATM interfaces. The following illustrates this command: **show mpls atm-ldp bindings** [*network* {*mask* | *length*}] [**local-label** *vpi vci*] [**remote-label** *vpi vci*] [**neighbor** *interface*].

This table describes the syntax for the **show mpls atm-ldp bindings** command.

show mpls atm-ldp bindings Syntax Description

Parameter	Description
network	(Optional) Defines the destination network number.
mask	(Optional) Defines the network mask in the form A.B.C.D (destination prefix).
length	(Optional) Defines the mask length (1 to 32).
local-label *vpi vci*	(Optional) Selects the label values assigned by this router. (VPI range is 0 to 4095. VCI range is 0 to 65535.)
remote-label *vpi vci*	(Optional) Selects the label values assigned by the other router. (VPI range is 0 to 4095. VCI range is 0 to 65535.)
neighbor *interface*	(Optional) Selects the label values assigned by the neighbor on a specified interface.

show mpls atm-ldp capability

To display the MPLS ATM capabilities negotiated with LDP neighbors for LC-ATM interfaces, use the following **show mpls atm-ldp capability** command in privileged EXEC mode: **show mpls atm-ldp capability**.

How to Display Summary Information About ATM Entries

This topic describes how to display summary information about all of the entries in the ATM label-binding database.

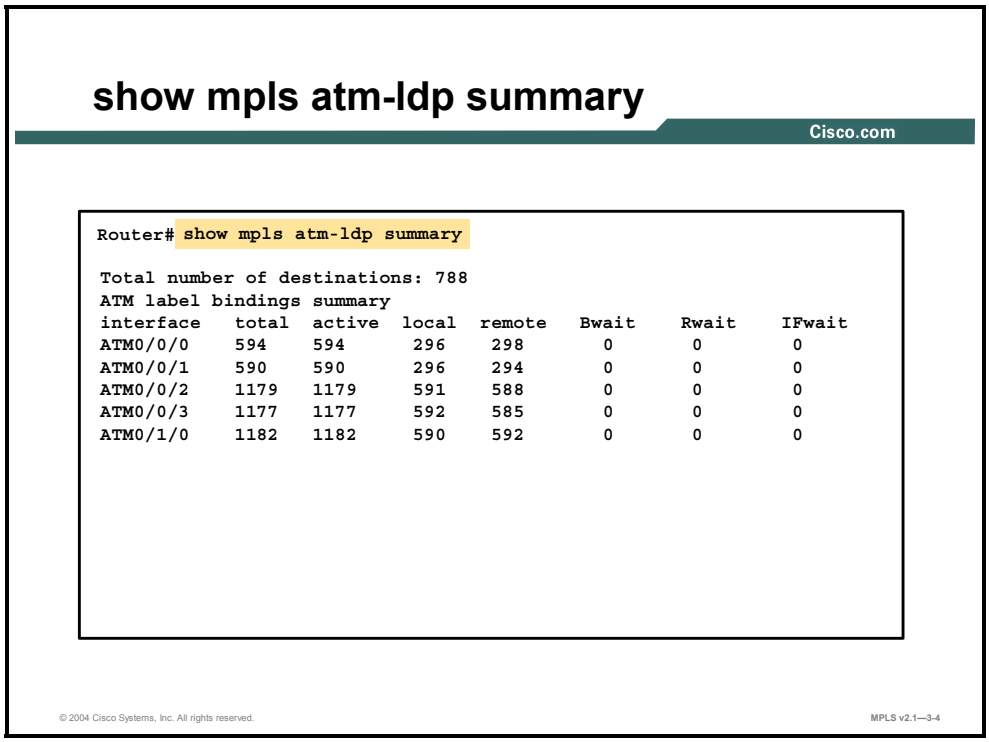

To display summary information about all of the entries in the ATM label-binding database, use the **show mpls atm-ldp summary** command in privileged EXEC mode.

This table describes the significant fields in the display.

show mpls atm-ldp summary Field Description

Field	Description
Total number of destinations	Number of known destination address prefixes.
interface	Name of an interface with associated ATM label bindings.
total	Total number of ATM labels on this interface.
active	Number of ATM labels in an “active” state, ready to use for data transfer.
local	Number of ATM labels assigned by this LSR on this interface.
remote	Number of ATM labels assigned by the neighbor LSR on this interface.
Bwait	Number of bindings that are waiting for a label assignment from the neighbor LSR.
Rwait	Number of bindings that are waiting for resources (VPI/VCI space) to be available on the downstream device.
IFwait	Number of bindings that are waiting for learned labels to be installed for switching use.

How to Display Current Label Bindings

This topic describes how to display current label bindings.

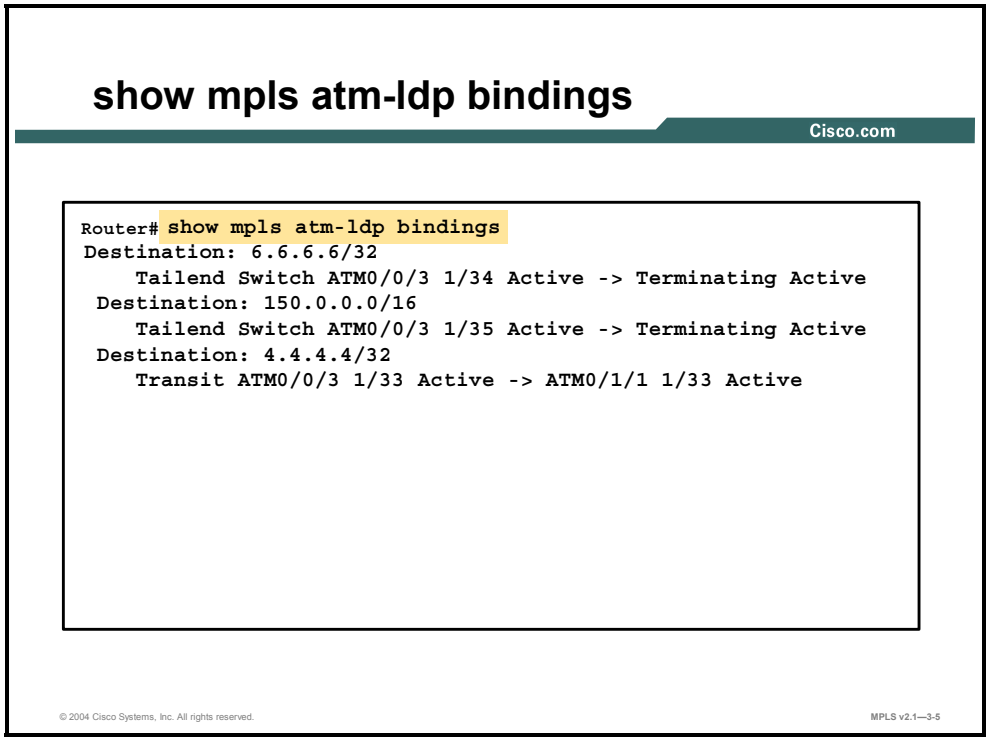

To display current label bindings, use the **show mpls atm-ldp bindings** command in privileged EXEC mode.

This table describes the significant fields in the display.

show mpls atm-ldp bindings Field Description

Field	Description
Destination	Destination (network and mask).
Headend Router Tailend Router Tailend Switch Transit	Indicates types of virtual circuits. Options include the following: ■ **Headend:** Virtual circuit that originates at this router ■ **Tailend:** Virtual circuit that terminates at this platform ■ **Transit:** Virtual circuit that passes through a switch
ATM0/0/3	Interface.
1/34	VPI/VCI.
Field	Description.
Active	Indicates the virtual circuit state. Options include the following: ■ **Active:** Set up and working ■ **Bindwait:** Waiting for a response ■ **Remote Resource Wait:** Waiting for resources (VPI/VCI space) to be available on the downstream device ■ **Parent Wait:** Transit virtual circuit input side waiting for output side to become active
VCD	Displays virtual circuit descriptor number.

How to Display MPLS ATM Capabilities by LDP

This topic describes how to display the MPLS ATM capabilities negotiated by LDP.

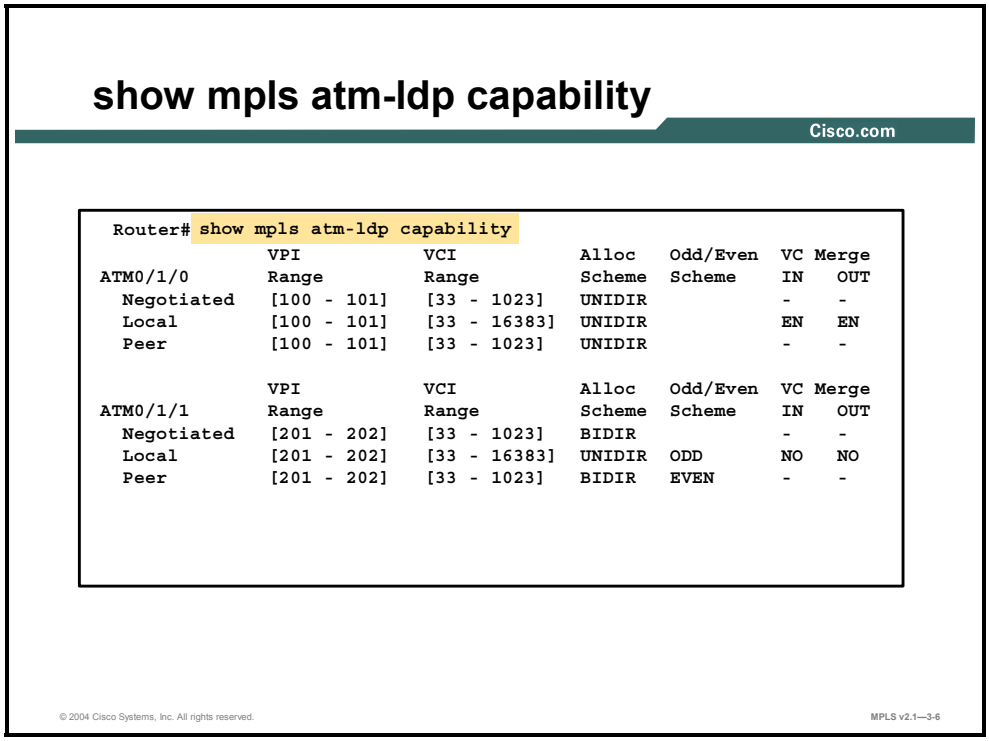

When two LSRs establish an LDP session, they negotiate parameters for the session, that is, a range of VPIs and VCIs that will be used as labels.

The **show mpls atm-ldp capability** command displays the MPLS ATM capabilities negotiated by LDP. The following explains each line of this command:

- The first line shows the negotiated (active) parameters.
- The second line shows the parameters proposed by this router.
- The third line shows the parameters proposed by the neighbor.

This table describes the significant fields in the display.

show mpls atm-ldp capability Field Description

Parameter	Description
VPI Range	Displays minimum and maximum number of VPIs supported on this interface.
VCI Range	Displays minimum and maximum number of VCIs supported on this interface.
Alloc Scheme	Indicates the applicable allocation scheme, as follows: ■ **UNIDIR:** Unidirectional capability indicates that the peer can, within a single VPI, support binding of the same VCI to different prefixes on different directions of the link. ■ **BIDIR:** Bidirectional capability indicates that within a single VPI, a single VCI can appear in one binding only. In this case, one peer allocates bindings in the even VCI space, and the other in the odd VCI space. The system with the lower LDP identifier assigns even-numbered VCIs. The negotiated allocation scheme is UNIDIR, but only if both peers have UNIDIR capability. Otherwise, the allocation scheme is BIDIR. **NOTE:** These definitions for “unidirectional” and “bidirectional” are consistent with normal ATM usage of the terms; however, they are exactly opposite from the definitions for them in the IETF LDP specification.
Odd/Even Scheme	Indicates whether the local device or the peer is assigning an odd- or even-numbered VCI when the negotiated scheme is BIDIR. This parameter does not display any information when the negotiated scheme is UNIDIR.
VC Merge	Indicates the type of VC merge support available on this interface. There are two possibilities, as follows: **IN:** Indicates the input interface merge capability. IN accepts the following values: ■ **EN:** The hardware interface supports VC merge, and VC merge is enabled on the device. ■ **DIS:** The hardware interface supports VC merge, and VC merge is disabled on the device. ■ **NO:** The hardware interface does not support VC merge. **OUT:** Indicates the output interface merge capability. OUT accepts the same values as the input merge side. The VC merge capability is meaningful only on ATM switches. This capability is not negotiated.
Negotiated	Indicates the set of options that both LDP peers have agreed to share on this interface. For example, the VPI or VCI allocation on either peer remains within the negotiated range.
Local	Indicates the options supported locally on this interface.
Peer	Indicates the options supported by the remote LDP peer on this interface.

Debugging Specific ATM LDP Functions

This topic describes how to debug ATM LDP issues.

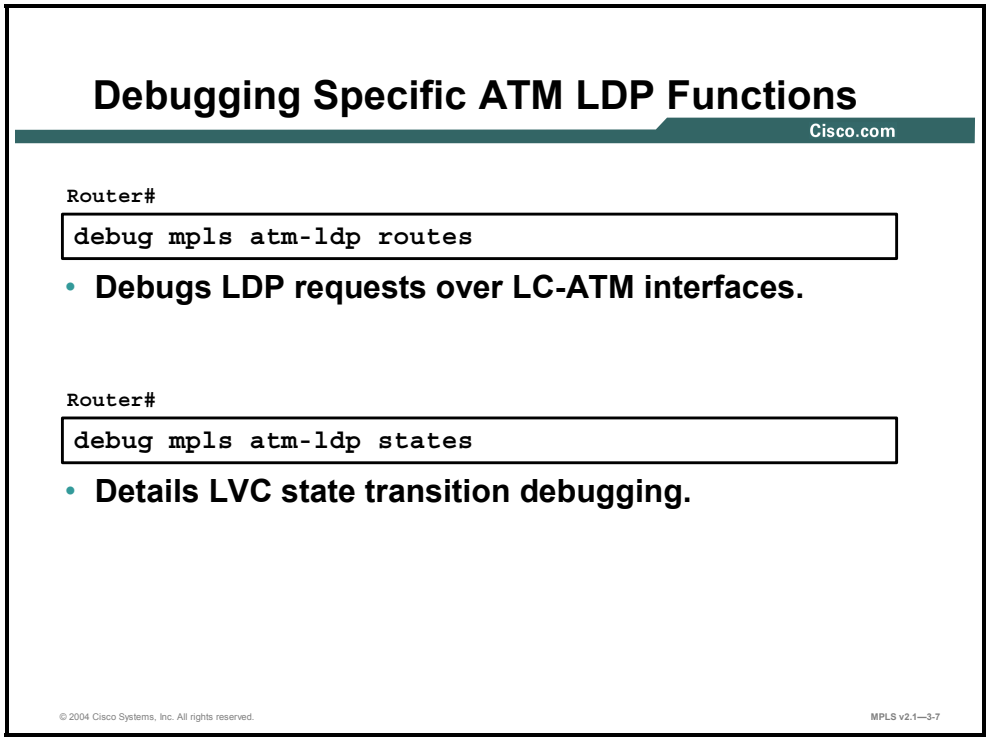

debug mpls atm-ldp routes

The **debug mpls atm-ldp routes** command displays information about the state of the routes for which VCI requests are being made.

When there are many routes and system activities (shutting down interfaces, learning new routes, and so on), the **debug mpls atm-ldp routes** command displays extensive information that might interfere with system timing. Most commonly, this interference affects normal LDP operation. To avoid this problem, increase the LDP hold time with the **mpls ldp holdtime** command.

debug mpls atm-ldp states

The **debug mpls atm-ldp states** command displays information about LVC state transitions as they occur.

When there are many routes and system activities (shutting down interfaces, learning new routes, and so on), the **debug mpls atm-ldp states** command displays extensive information that might interfere with system timing. Most commonly, this interference affects normal LDP operation. To avoid this problem, increase the LDP hold time with the **mpls ldp holdtime** command.

Summary

This topic summarizes the key points discussed in this lesson.

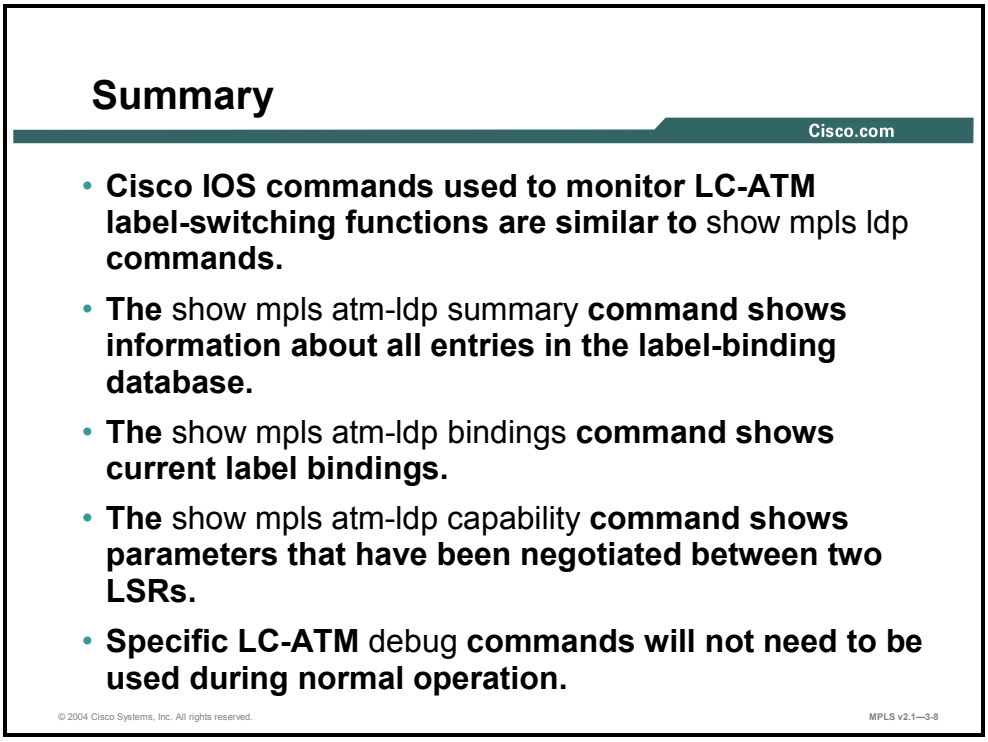

Module Summary

This topic summarizes the key points discussed in this module.

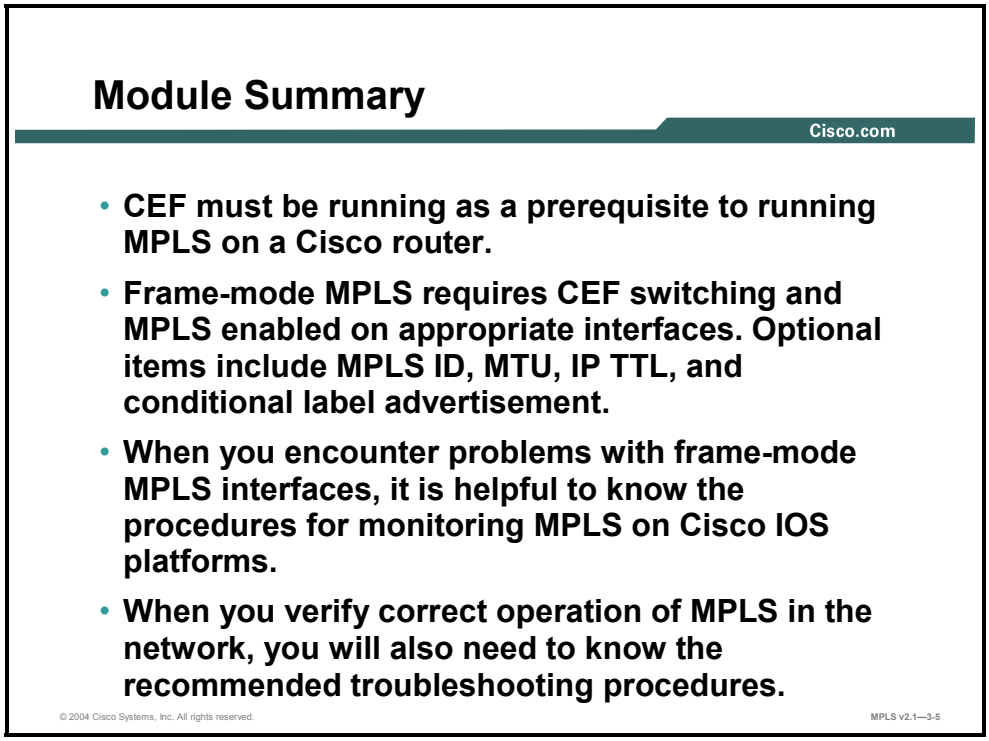

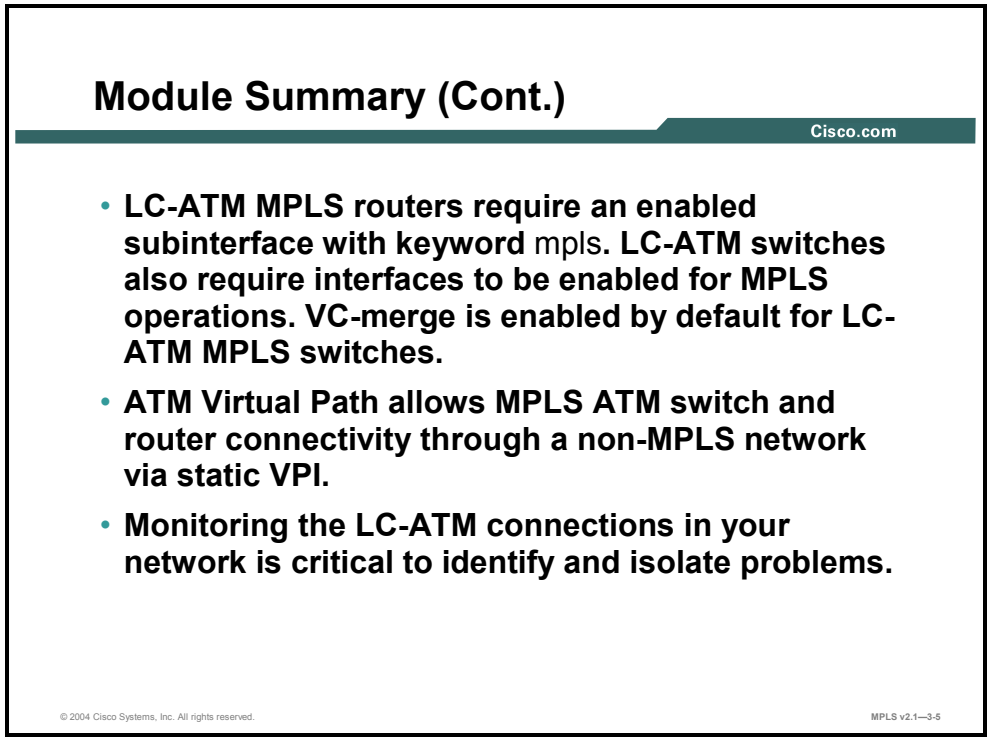

There are many detailed configuration, monitoring, and debugging guidelines when implementing frame-mode MPLS and cell-mode MPLS on Cisco IOS platforms. Advanced technologies, such as TTL propagation and label distribution, are also critical when switching implementations.

References

For additional information, refer to these resources:

- Search for “CEF switching” on Cisco.com for additional information.

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 ip cef**
C) Router#**show ip cef**
D) Router(config)#**show ip cef**

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

A) Router#**ip cef**
B) Router**>ip 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-B IP packets and MPLS stacks up to 3 levels deep. (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

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

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

Q9) Which of the following is NOT a mandatory step to enable MPLS? (Source: Configuring Frame-Mode MPLS on Cisco IOS Platforms)

A) Enable CEF switching.
B) Label the pool configuration.
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 following 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 of the following correctly implements the **show cef interface** command? (Source: Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms)

A) router>**show cef interface**
B) router#**show cef interface**
C) router(config)#**show cef interface**
D) router(config-router))#**show 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. The most possible cause would be which of the following? (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.

Q22) A _____ must be created on an LC-ATM router to support MPLS. (Source: Configuring LC-ATM MPLS)

Q23) On Cisco IOS platform routers, _____ mode MPLS is the default. (Source: Configuring LC-ATM MPLS)

Q24) Which VPI value do all LVCs use by default? (Source: Configuring LC-ATM MPLS)

A) 0
B) 1
C) 32
D) 100

Q25) For successful establishment of a label distribution session between an LC router and an ATM switch, both devices need to use the same of which item? (Source: Configuring LC-ATM MPLS)

A) IGP
B) VRI/VDI
C) subinterface number
D) label distribution protocol

Q26) Which command sets the threshold that will prevent forwarding loops in the setting up of label switch paths across an ATM region? (Source: Configuring LC-ATM MPLS)

A) **mpls vpi**
B) **mpls atm vpi**
C) **mpls maxhops**
D) **mpls ldp maxhops**

Q27) Which two of the following statements are correct? (Choose two.) (Source: Configuring LC-ATM MPLS)VC merge is enabled by default on all ATM switches.
B) VC merge is disabled by default on all ATM switches.
C) Disabling VC merge results in the ability to interleave cells, but an LVC must be created for every source-destination pair.
D) Disabling VC merge results in the ability to interleave cells, but an LVC will NOT be created for every source-destination pair.

Q28) What is a virtual path? (Source: Configuring LC-ATM MPLS over ATM Virtual Path)

A) a pool of MPLS labels
B) a collection of virtual circuits with a common VDI
C) a collection of virtual circuits with a common VPI
D) a collection of virtual circuits with a common VCI

Q29) Why is it mandatory that the VPI be used on both ends of the virtual path over a public ATM network? (Source: Configuring LC-ATM MPLS over ATM Virtual Path)

A) because the VPI value is part of the LDP virtual path range negotiation
B) because the VCI value is part of the LDP virtual circuit range negotiation
C) because the TTL value would not be able to be propagated
D) It is not mandatory, but only recommended.

Q30) Which two of the following statements are correct when describing the configuration of ATM Virtual Path between two ATM switches? (Choose two.) (Source: Configuring LC-ATM MPLS over ATM Virtual Path)

A) The virtual path number has to match between peers.
B) The virtual path number does not have to match between peers.
C) The MPLS-enabled subinterface number is the same as the virtual path number.
D) The MPLS-enabled subinterface number cannot be the same as the virtual path number.

Q31) Which two of the following statements are correct when describing the configuration of ATM Virtual Path between two ATM routers? (Choose two.) (Source: Configuring LC-ATM MPLS over ATM Virtual Path)

A) The virtual path number has to match between peers.
B) The virtual path number does not have to match between peers.
C) The control virtual circuit cannot be established within the virtual path.
D) The control virtual circuit can be established within the virtual path.

Q32) Which command is NOT used to monitor LC-ATM label-switching functions? (Source: Monitoring LC-ATM MPLS on Cisco IOS Platforms)

A) **show mpls atm-ldp labels**
B) **show mpls atm-ldp bindings**
C) **show mpls atm-ldp summary**
D) **show mpls atm-ldp capability**

Q33) Which command provides summary information about all entries in the label-binding database? (Source: Monitoring LC-ATM MPLS on Cisco IOS Platforms)

A) **show mpls atm-ldp bindings**
B) **show mpls atm-ldp summary**
C) **show mpls atm-ldp capability**
D) **show mpls atm-ldp labels-summary**

Q34) Which command shows the current label bindings? (Source: Monitoring LC-ATM MPLS on Cisco IOS Platforms)

A) **show mpls atm-ldp bindings**
B) **show mpls atm-ldp summary**
C) **show mpls atm-ldp capability**
D) **show mpls atm-ldp labels-bindings**

Q35) Which command shows the negotiated parameters between LSRs? (Source: Monitoring LC-ATM MPLS on Cisco IOS Platforms)

A) **show mpls atm-ldp lsr**
B) **show mpls atm-ldp bindings**
C) **show mpls atm-ldp summary**
D) **show mpls atm-ldp capability**

Q36) Which of the following is used to debug an LC-ATM issue? (Source: Monitoring LC-ATM MPLS on Cisco IOS Platforms)

A) **debug mpls atm-ldp lsrs**
B) **debug mpls atm-ldp routes**
C) **debug mpls atm-ldp nodes**
D) **debug mpls atm-ldp switches**

Module Self-Check Answer Key

Q1) A

Q2) C

Q3) D

Q4) A

Q5) C

Q6) D

Q7) 1512

Q8) D

Q9) 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

Q22) subinterface

Q23) frame-

Q24) B

Q25) D

Q26) D

Q27) A, C

Q28) C

Q29) A

Q30) A, C

Q31) A, D

Q32) A

Q33) B

Q34) A

Q35) D

Q36) B

Module 4

MPLS Virtual Private Network 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 features, benefits, and drawbacks of overlay VPN and peer-to-peer VPN
- Describe the characteristics of the different VPN topology categories
- Describe the major architectural components of MPLS VPNs
- Identify the routing requirements for MPLS VPNs
- Describe how packets are forwarded in an MPLS VPN environment

Lesson 1

Introducing Virtual Private Networks

Overview

This lesson explains the concept of VPNs and the terminology introduced by MPLS VPN architecture. The lesson also looks at why VPNs were first introduced.

It is important to understand the background of VPNs, because moving forward, you should be able to determine the need for a VPN and explain how MPLS VPNs can help save time and money for a customer.

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 how VPNs replace the connectivity of traditional router-based networks
- Identify the major network elements in a VPN
- Describe how virtual circuits are used in switched WANs to create a VPN

Traditional Router-Based Network Connectivity

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

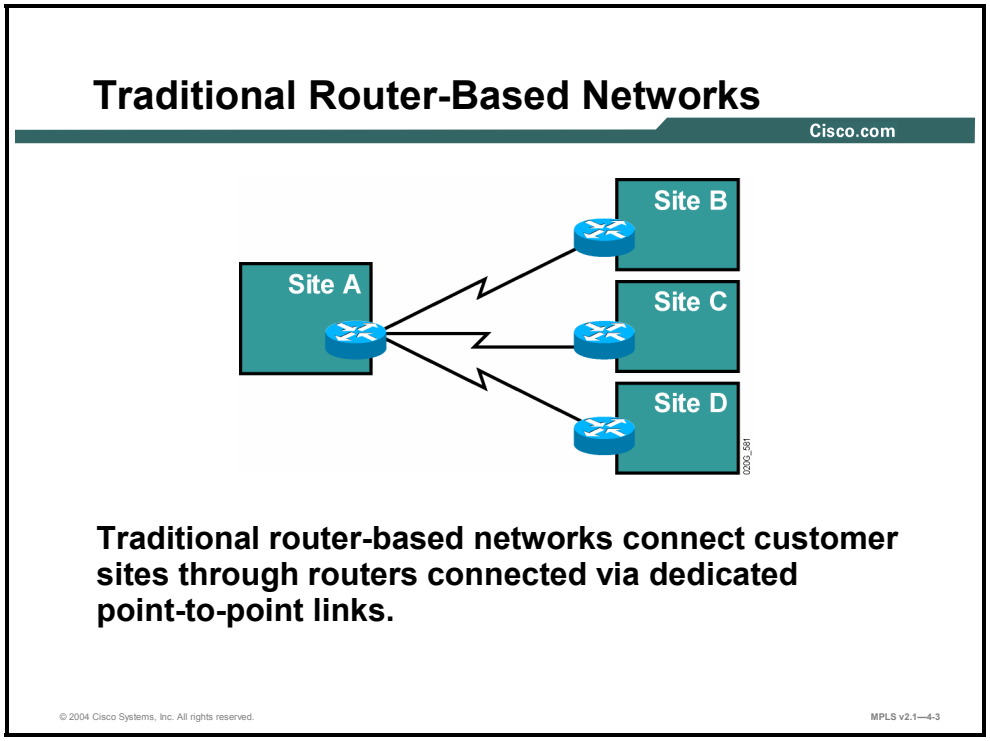

Traditional router-based networks were implemented with dedicated point-to-point links connecting customer sites. The cost of this approach was comparatively high for the following reasons:

- 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 Virtual Private Networks

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

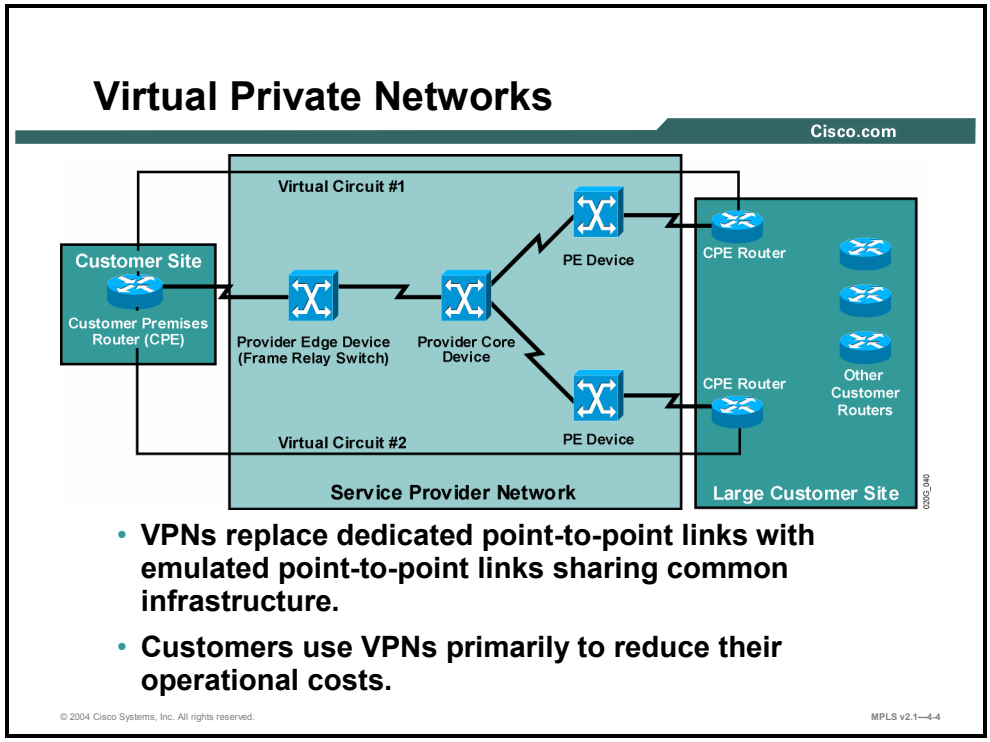

VPNs were introduced very early in the history of data communications with technologies such as X.25 and Frame Relay, which use virtual circuits to establish the end-to-end connection over a shared service provider infrastructure. The following technologies, although sometimes considered legacy technologies and obsolete, 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: Virtual Private Networks

The figure shows the statistical sharing, where the customer premises equipment (CPE) router on the left has one physical connection to the service provider and two virtual circuits 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.

What Are VPN Network Elements?

This topic identifies the major network elements in a VPN.

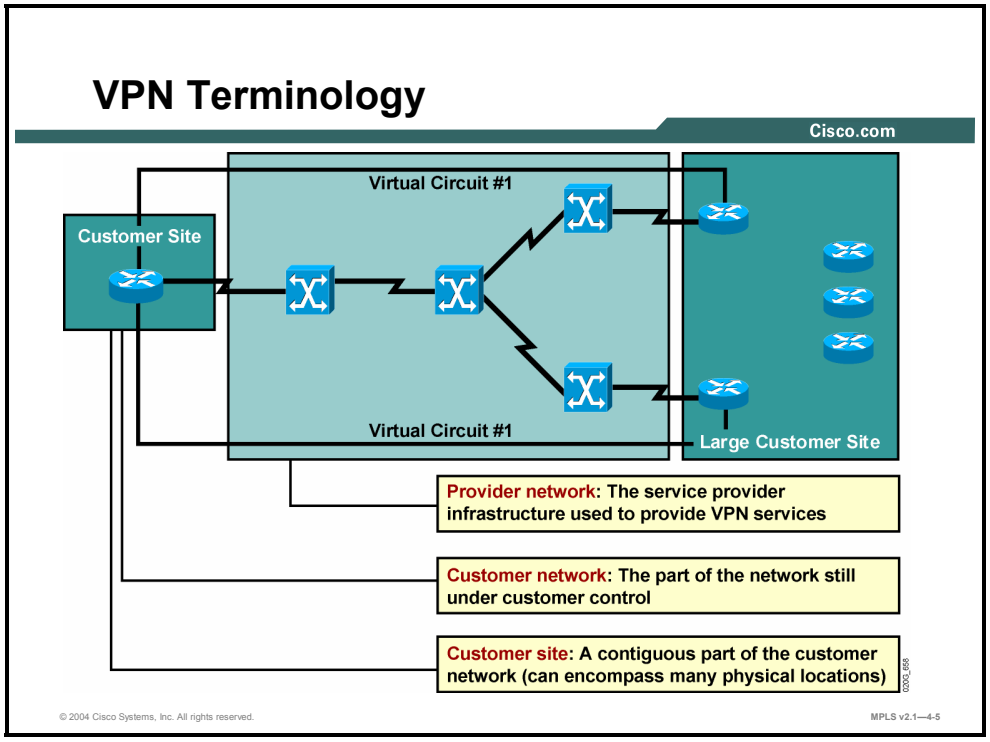

There are many conceptual models and terminologies describing various VPN technologies and implementations. The terminology is generic enough to cover nearly any VPN technology or implementation and is thus extremely versatile.

The major parts of an overall VPN solution are always the following:

- **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).

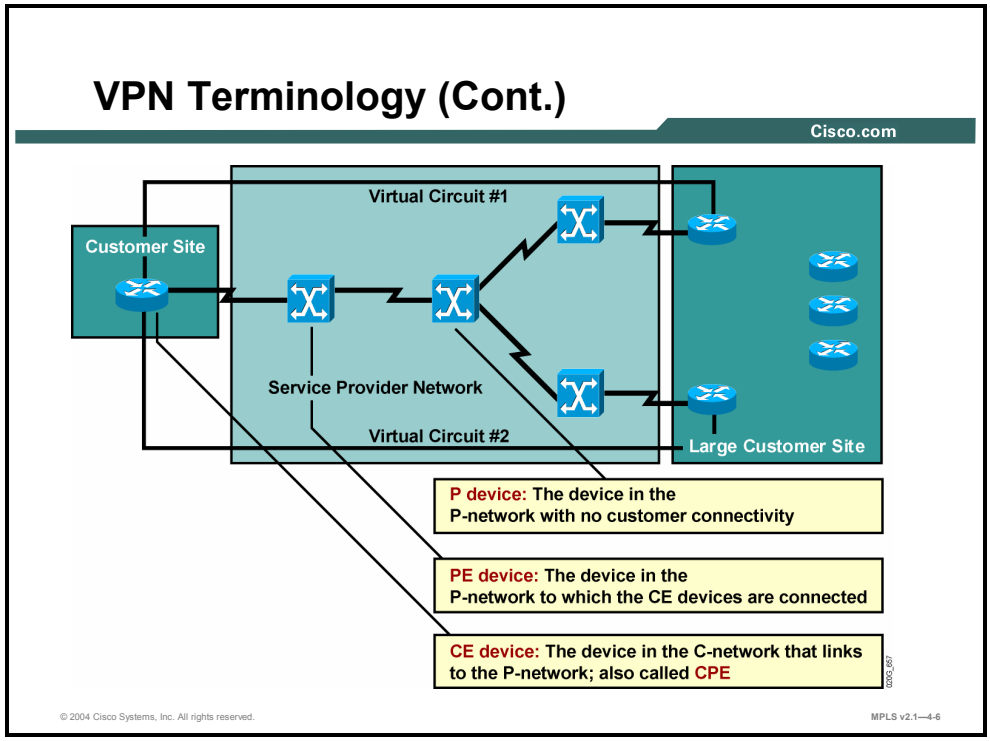

The following describes the devices that enable the overall VPN solution, which are named based on their position in the network:

- 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 provider edge (PE) devices. In traditional switched WAN implementations, these devices would be Frame Relay or X.25 edge switches.
- 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.

Note If the connecting device is not a router but, for example, a packet assembler/disassembler (PAD), it is still called a CE device.

How Are Virtual Circuits Used in Switched WANs?

This topic describes how virtual circuits are used in switched WANs to create a VPN.

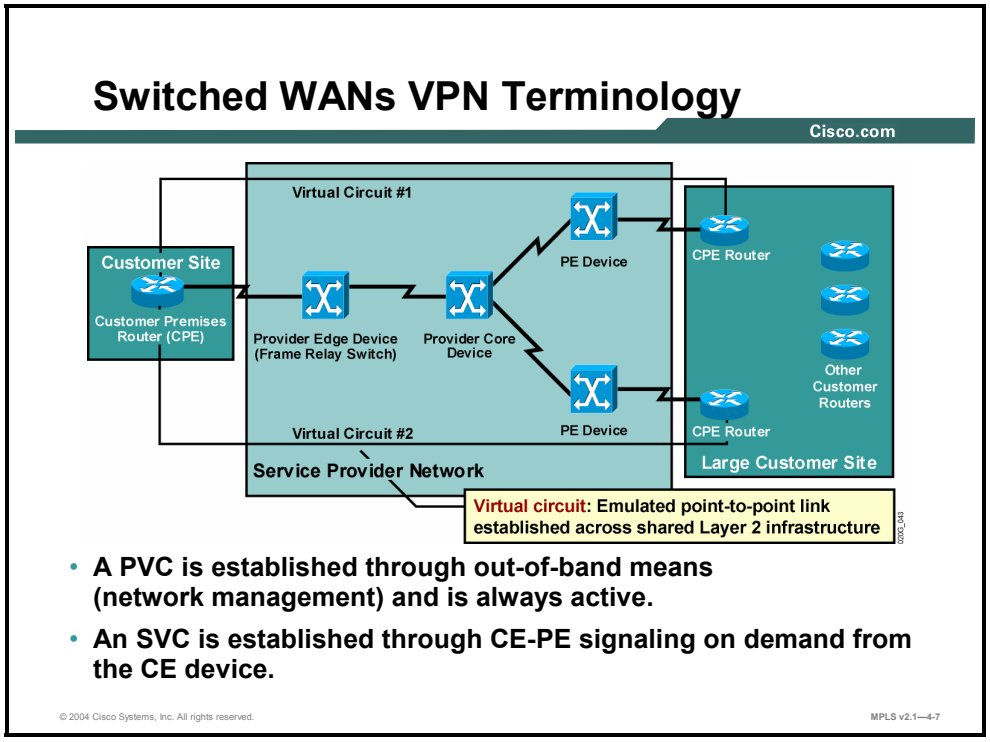

Switched WAN technologies introduced the virtual circuit, an emulated point-to-point link established across the Layer 2 infrastructure (for example, a Frame Relay network). Virtual circuits are further differentiated into permanent virtual circuits (PVCs), which are preestablished by means of network management or manual configuration, and switched virtual circuits (SVCs), which are established on demand through a call setup request from the CE device.

Summary

This topic summarizes the key points discussed in this lesson.

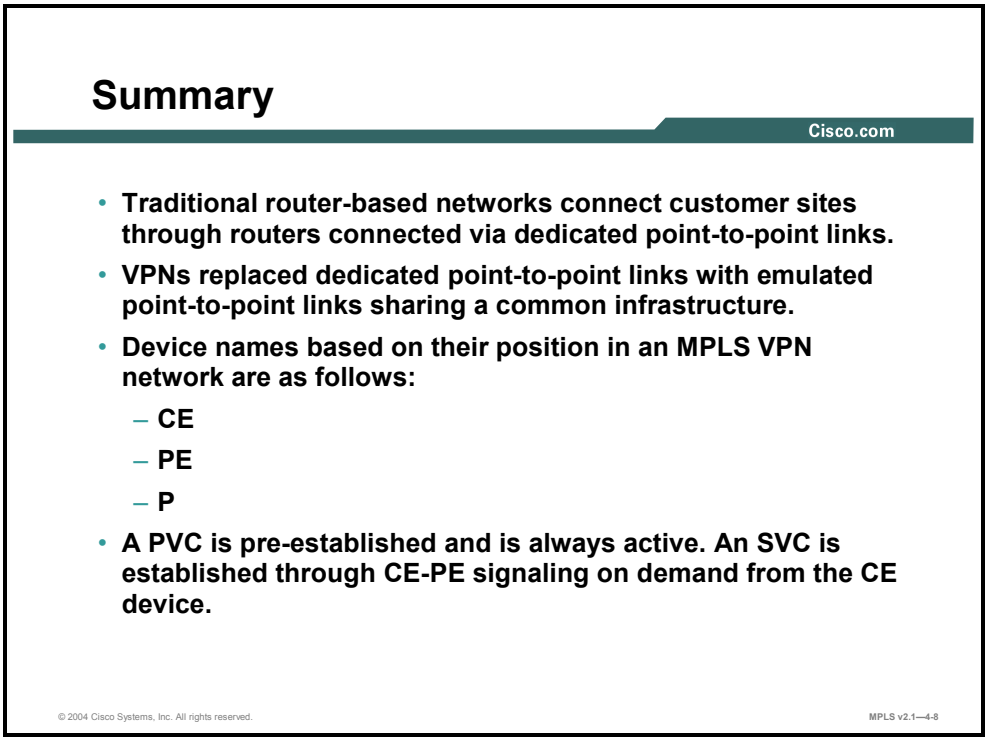

Lesson 2

Introducing Overlay and Peer-to-Peer VPNs

Overview

This lesson explains the differences between the overlay and peer-to-peer VPN models, how they are implemented, and the benefits and drawbacks of each implementation. The lesson also discusses the various virtual networking concepts.

It is important to understand the different types of VPNs, and how each one is used. This understanding 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 describe the differences between overlay VPNs and peer-to-peer VPNs, explaining their implementation, benefits, and drawbacks. This ability includes being able to meet these objectives:

- Identify the two major VPN implementation technologies
- Describe the implementation techniques for overlay VPNs
- Describe the implementation techniques for peer-to-peer VPNs
- Describe the benefits of each type of VPN model
- Describe the drawbacks of each VPN model
- Describe the drawbacks of the traditional peer-to-peer VPN model

What Are the VPN Implementation Technologies?

This topic describes the two major VPN implementation technologies.

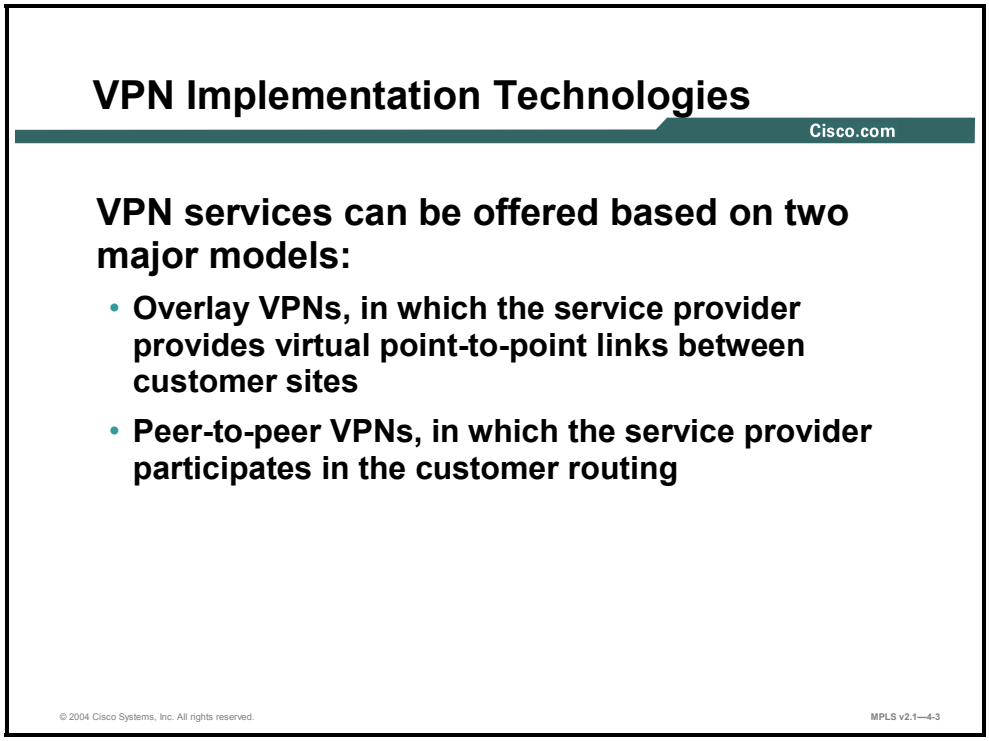

Traditional VPN implementations were all based on the overlay model, in which the service provider sold virtual circuits between customer sites as a replacement for dedicated point-to-point links. The overlay model had a number of drawbacks, which are identified in this lesson. To overcome these drawbacks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

What Are the Overlay VPN Implementation Techniques?

This topic describes the implementation techniques for overlay VPNs.

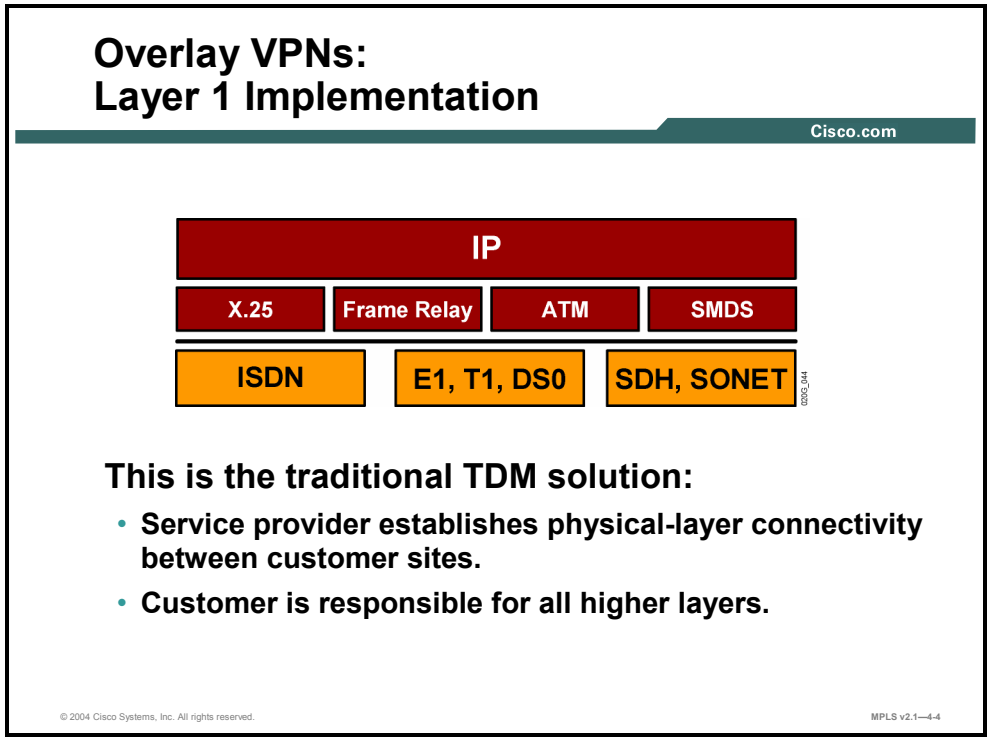

In the Layer 1 overlay VPN implementation, the service provider sells Layer 1 circuits (bit pipes) implemented with technologies such as ISDN, digital service zero (DS0), E1, T1, Synchronous Digital Hierarchy (SDH), or SONET. The customer is responsible for Layer 2 encapsulation between customer devices and the transport of IP data across the infrastructure.

Overlay VPNs: Layer 2 Implementation

Cisco.com

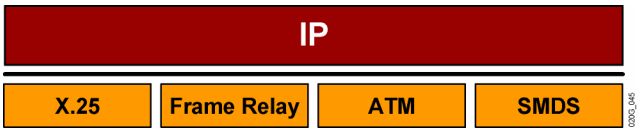

This is the traditional switched WAN solution:

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

A Layer 2 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.

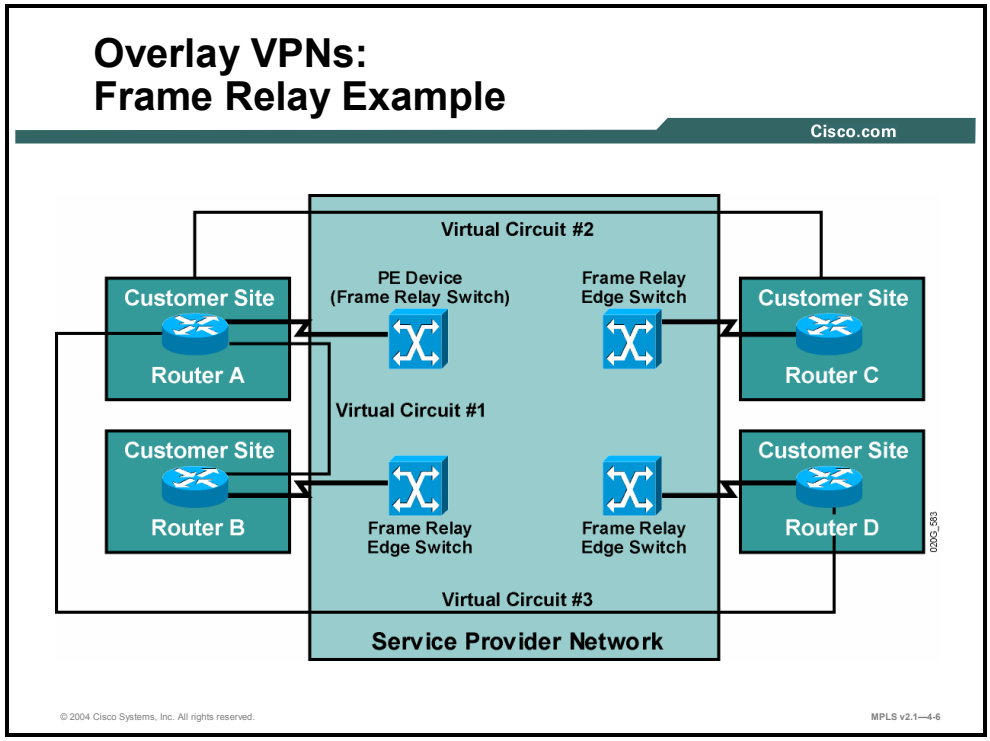

Example: Overlay VPS—Frame Relay

The figure shows a typical overlay VPN implemented by a Frame Relay network. The customer needs to connect three sites to site A (central site, or hub) and orders connectivity between site A (hub) and site B (spoke), between site A and site C (spoke), and between site A and site D (spoke). The service provider implements this request by providing two PVCs across the Frame Relay network.

Note	The implementation displayed in this example does not provide full connectivity. Data flow between spoke sites is through the hub.

Overlay VPNs: IP Tunneling

Cisco.com

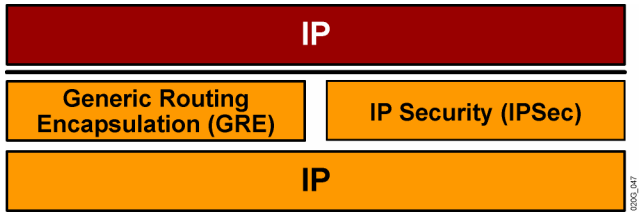

VPN is implemented with IP-over-IP tunnels:

- **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

Cisco.com

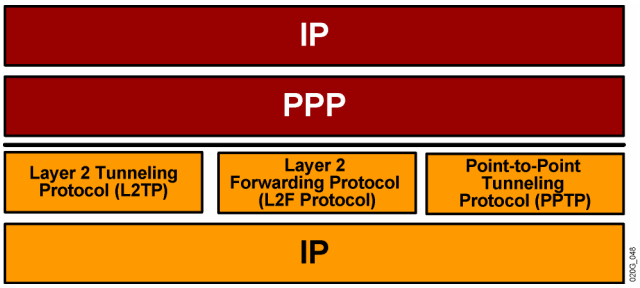

- **VPN is implemented with PPP-over-IP tunnels.**
- **VPN is usually used in access environments (dialup, digital subscriber line).**

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.

The following 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

Cisco.com

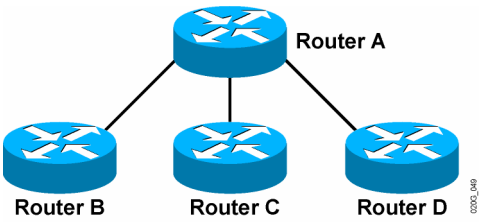

- The service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
- The service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

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.

What Are the Implementation Techniques for Peer-to-Peer VPNs?

This topic describes the implementation techniques for peer-to-peer VPNs.

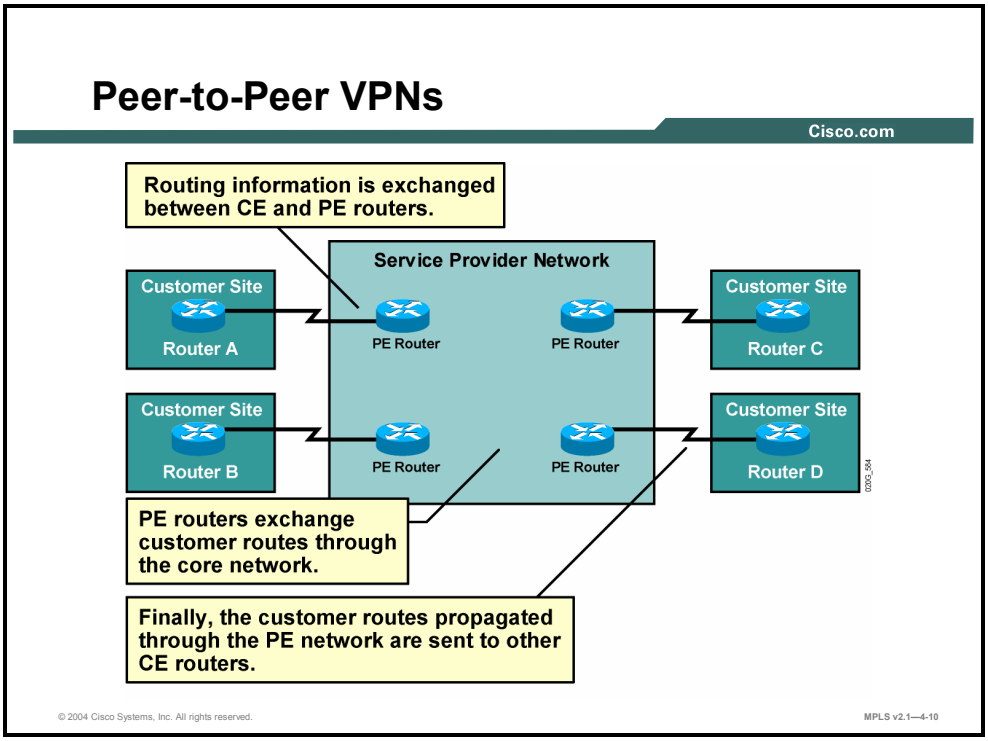

The overlay VPN model has a number of drawbacks, most significantly the need for customers to establish point-to-point links or virtual circuits between sites. The formula to calculate how many point-to-point links or virtual circuits are needed in the worst case is ([*n*][*n*-1])/2, where *n* is the 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. To overcome this drawback 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.

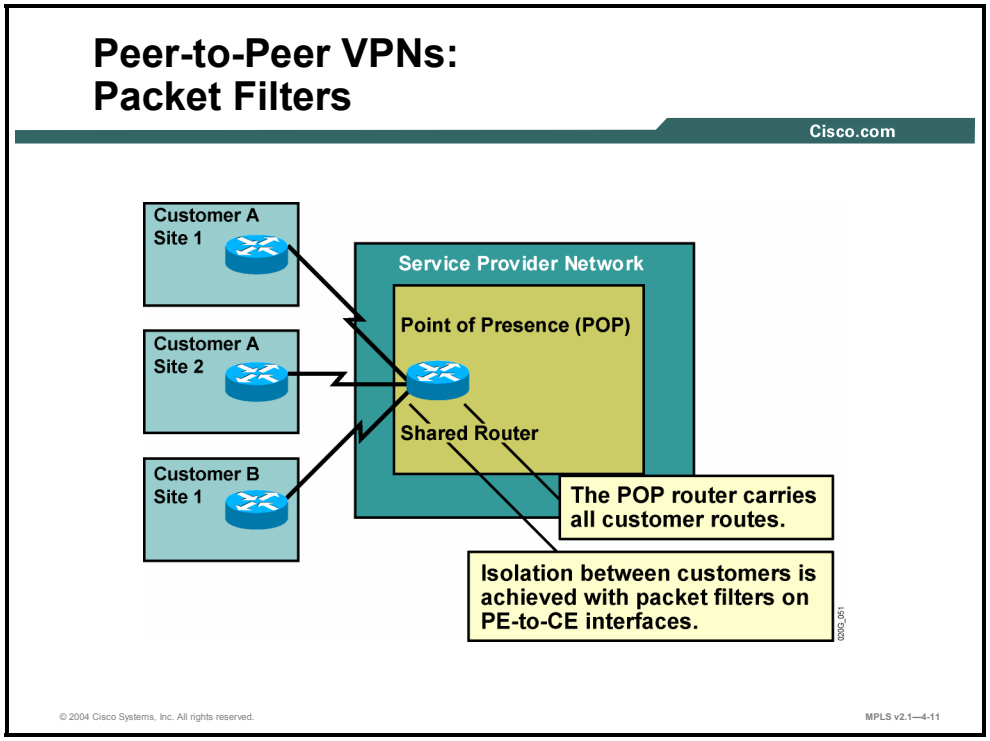

The first peer-to-peer VPN solutions appeared with the widespread deployment of IP in service provider networks. Architectures similar to that of the Internet were used to build them. Special provisions were taken into account to transform the architecture, which was targeted toward public backbones (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.

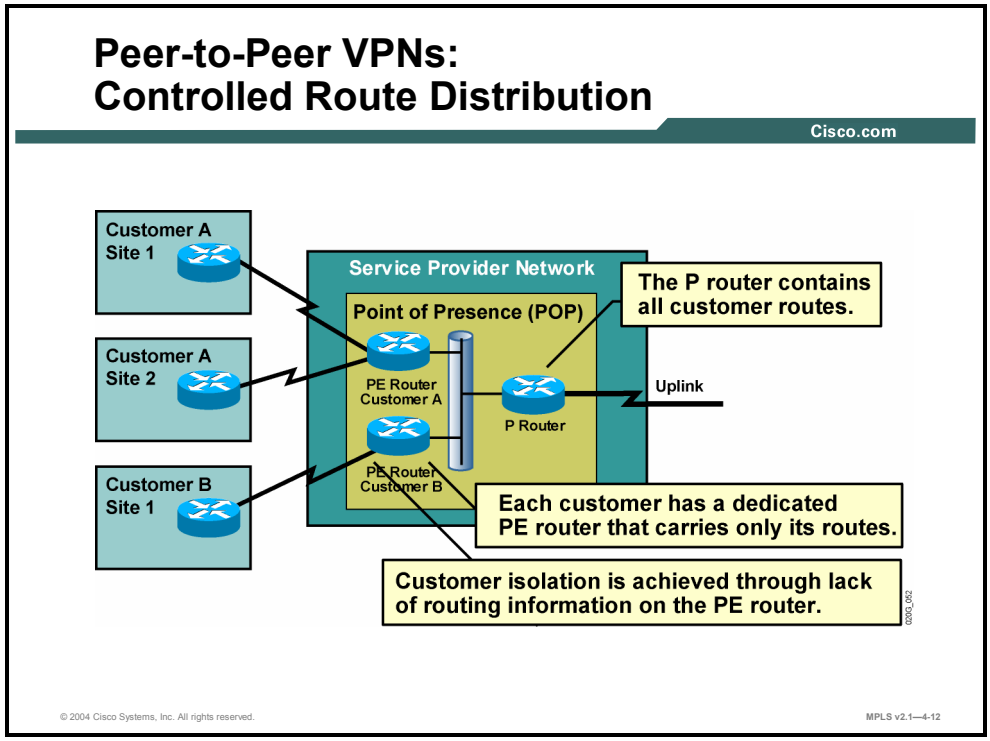

Maintaining packet filters is a mundane and error-prone task. Some service providers have thus implemented more innovative solutions based on controlled route distribution. In this approach, the customer has a dedicated PE router. The core service P routers contain all customer routes, and the dedicated 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. 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 VPN Implementations?

This topic describes the benefits of each type of MPLS VPN implementation.

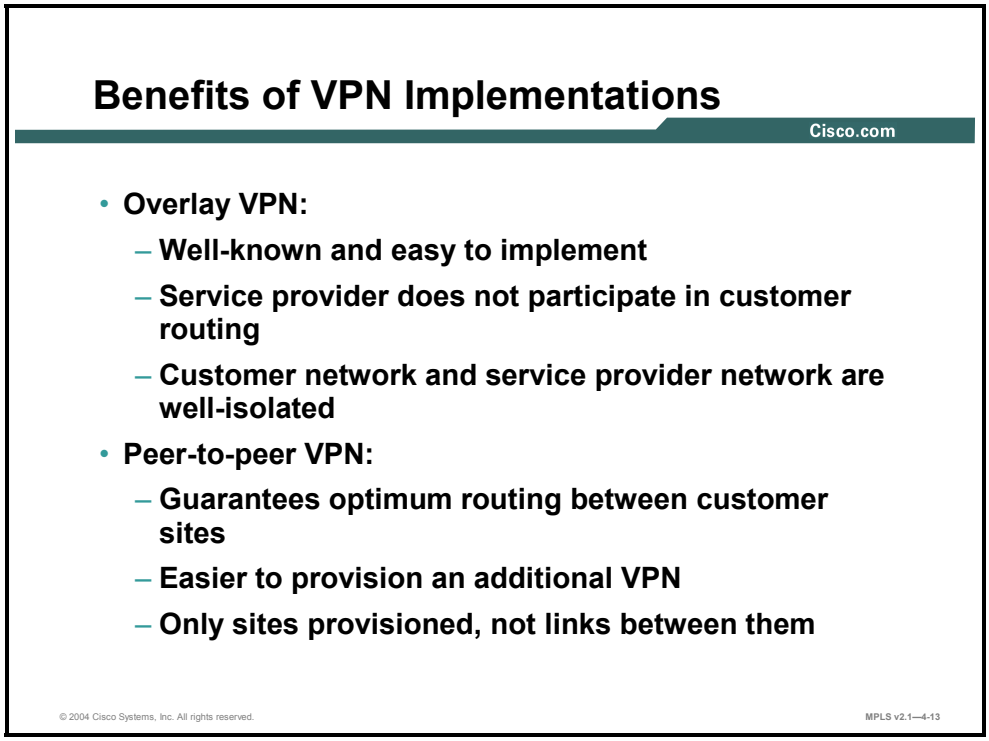

Each VPN model has a number of benefits. For example, overlay VPNs have the following advantages:

- 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 the following:

- 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 VPN Implementations?

This topic describes the drawbacks of each VPN implementation model.

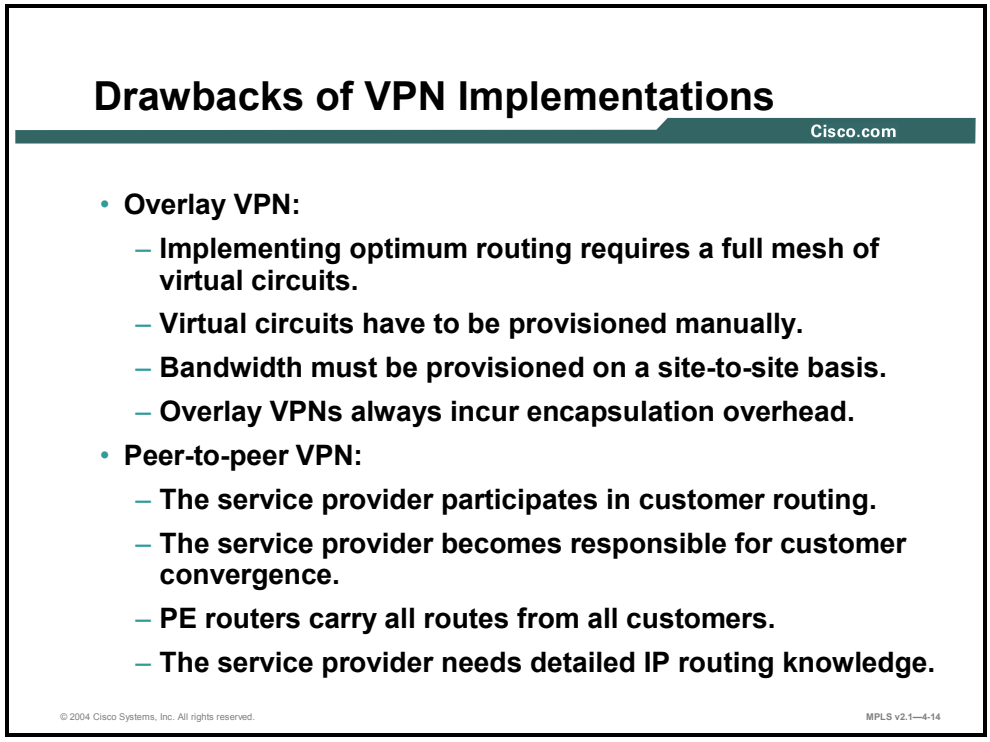

Each VPN model also has a number of drawbacks. Overlay VPNs have the following disadvantages:

- 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 (B) to 80 B per transported datagram.

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

- 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.

What Are the Drawbacks of Traditional Peer-to-Peer VPNs?

This topic describes the drawbacks of the traditional peer-to-peer VPN implementation model.

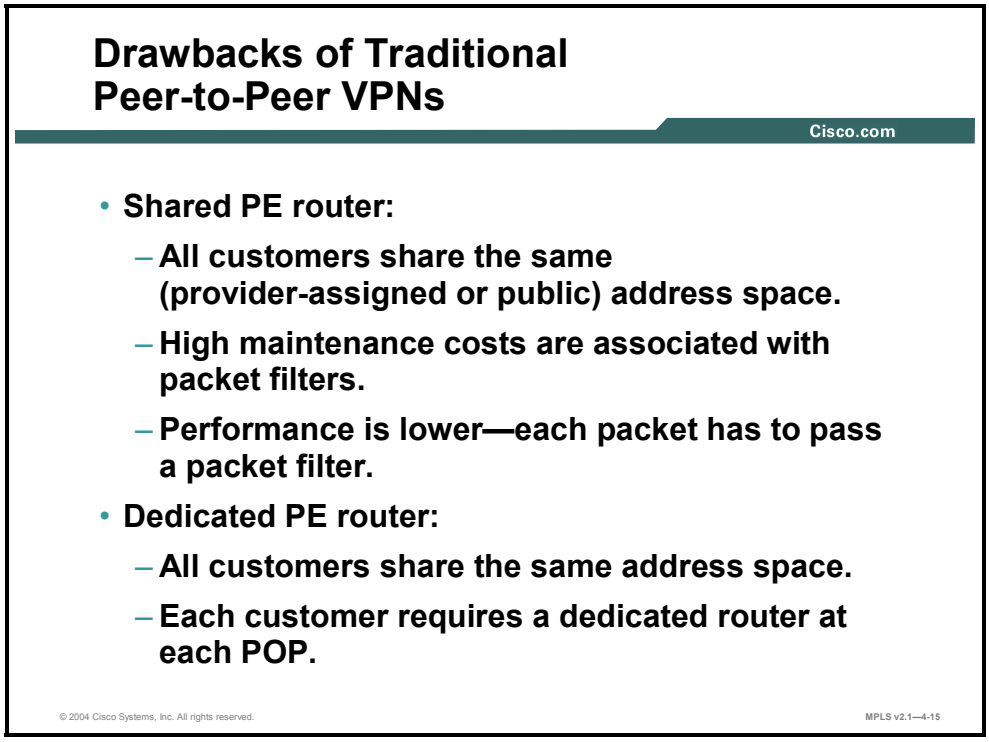

Pre-MPLS VPN implementations or peer-to-peer VPNs all share a common drawback. Customers have to share the same global address space, either using their own public IP addresses or relying on provider-assigned IP addresses. In both cases, connecting a new customer to a peer-to-peer VPN service usually requires IP renumbering inside the 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 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 POP. Thus, this approach is usually used if the service provider has only a small number of large customers.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- The two major VPN models are overlay VPN and peer-to-peer VPN.
- Overlay VPNs can be implemented using Layer 1, Layer 2, and Layer 3 technologies.
- Traditional peer-to-peer VPNs are implemented using IP routing technology.
- Overlay VPNs use well-known technologies and are easy to implement. Peer-to-peer VPNs guarantee optimum routing between customer sites.

Summary (Cont.)

Cisco.com

- Overlay VPN virtual circuits must be provisioned manually. Peer-to-peer VPNs require that the service provider participate in customer routing.
- Both shared PE router and dedicated PE router implementations of peer-to-peer VPNs require customers to share a common address space.

Lesson 3

Categorizing VPNs

Overview

This lesson explains the different VPN topology categories, taking a closer look at each topology type and how VPNs can be categorized based on business need or connectivity requirement.

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 categories of the overlay VPN topology
- Describe the characteristics of the hub-and-spoke overlay VPN topology
- Describe the characteristics of the partial mesh overlay VPN topology
- Identify the major components of the VPN business category
- Describe the characteristics of the extranet component of the VPN business category
- Identify the major components of the VPN connectivity category
- Describe the characteristics of the central services extranet component of the VPN connectivity category
- Describe the characteristics of the managed network component of the VPN connectivity category

What Are the Overlay VPN Categories?

This topic identifies the major categories of the overlay VPN topology.

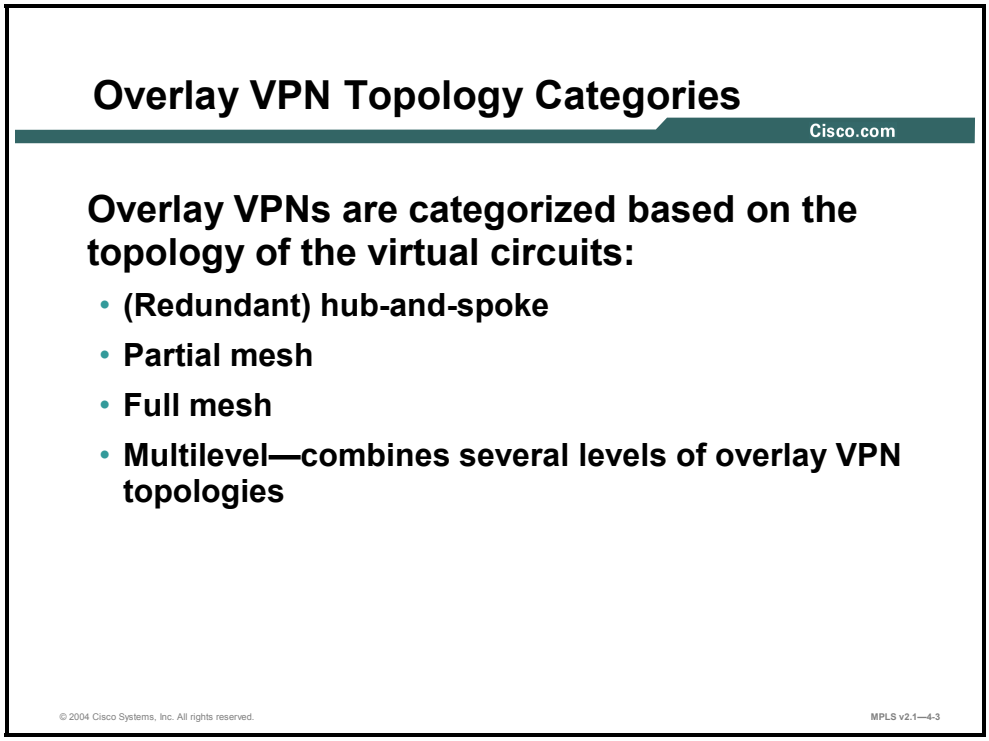

The oldest VPN category is based on the topology of point-to-point links in an overlay VPN implementation. Some VPN categories are as follows:

- **Hub-and-spoke:** The hub-and-spoke topology is the ultimate reduction of the partial mesh topology; many sites (spokes) are connected only with the central site(or sites), or hub (hubs), with no direct connectivity between the spokes. To prevent single points of failure, the hub-and-spoke topology is sometimes extended to a *redundant* hub-and-spoke topology.
- **Full mesh:** The full mesh topology provides a dedicated virtual circuit between any two CE routers in the network.
- **Partial mesh:** The partial mesh topology reduces the number of virtual circuits, usually to the minimum number that provides optimum transport between major sites.

Large networks usually deploy a layered combination of these technologies. Here are some examples:

- Partial mesh in the network core
- Redundant hub-and-spoke topology for larger branch offices (spokes) connected to distribution routers (hubs)
- Simple hub-and-spoke topology for noncritical remote locations (for example, home offices)

What Is the Hub-and-Spoke Overlay VPN Topology?

This topic describes the characteristics of the hub-and-spoke overlay VPN topology category.

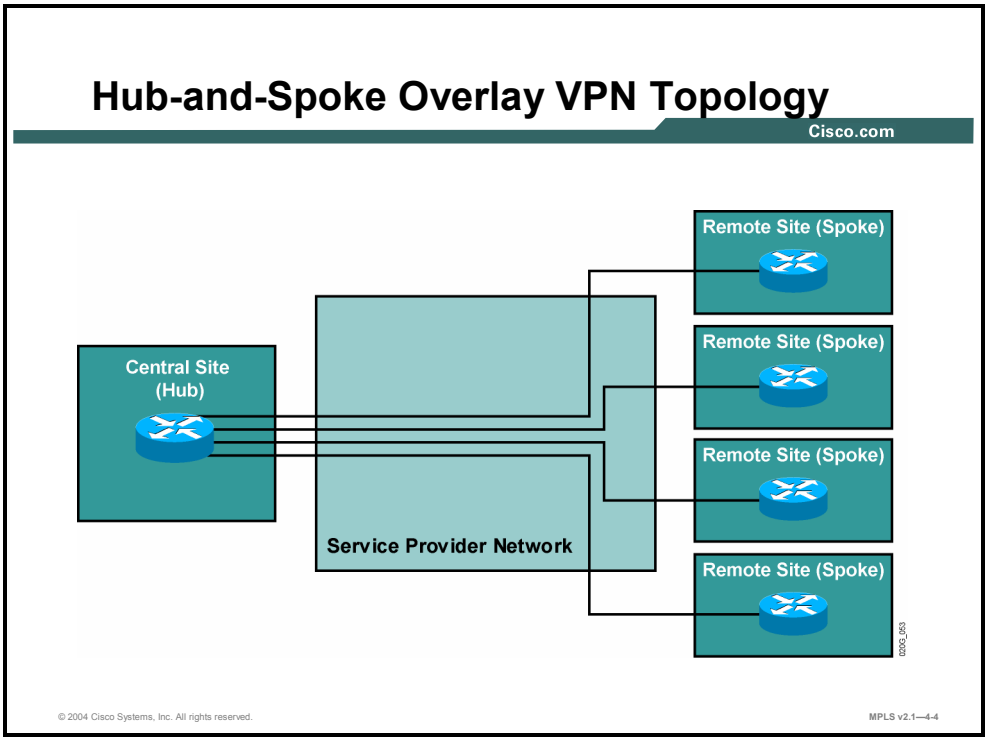

The hub-and-spoke topology is the simplest overlay VPN topology—all remote sites are linked with a single virtual circuit to a central CE router. The routing is also extremely simple—static routing or a distance vector protocol such as Routing Information Protocol (RIP) is more than adequate. If a dynamic 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.

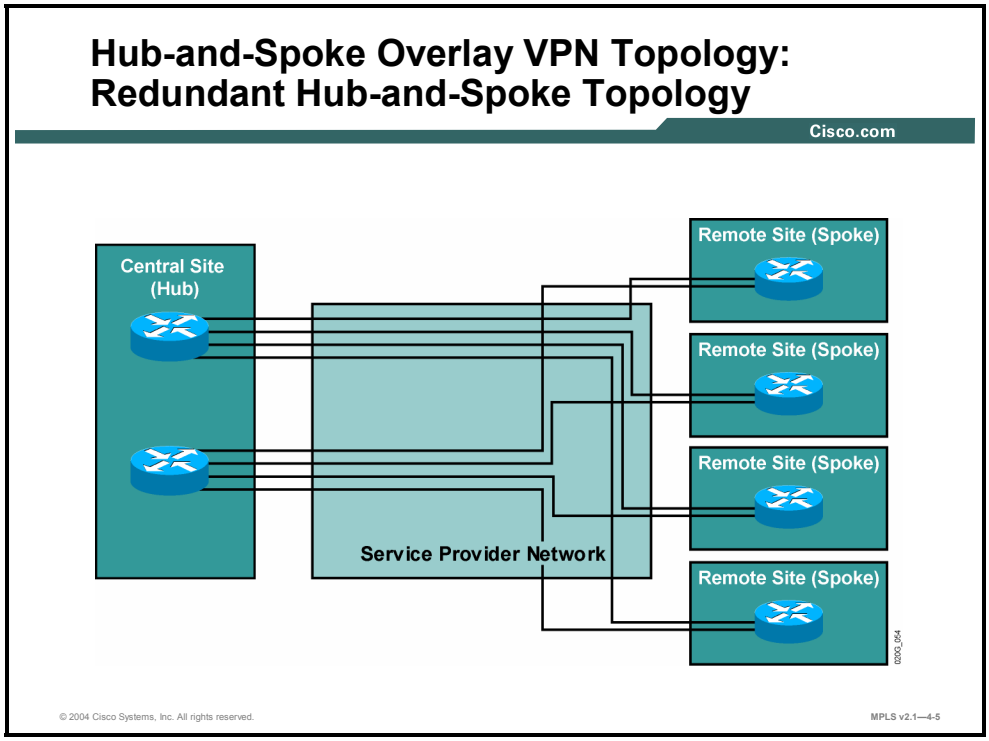

A typical redundant hub-and-spoke topology introduces central site redundancy (more complex topologies might also introduce router redundancy at spokes).

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 with backup configuration.

What Is the Partial Mesh Overlay VPN Topology?

This topic describes the characteristics of the partial mesh overlay VPN topology.

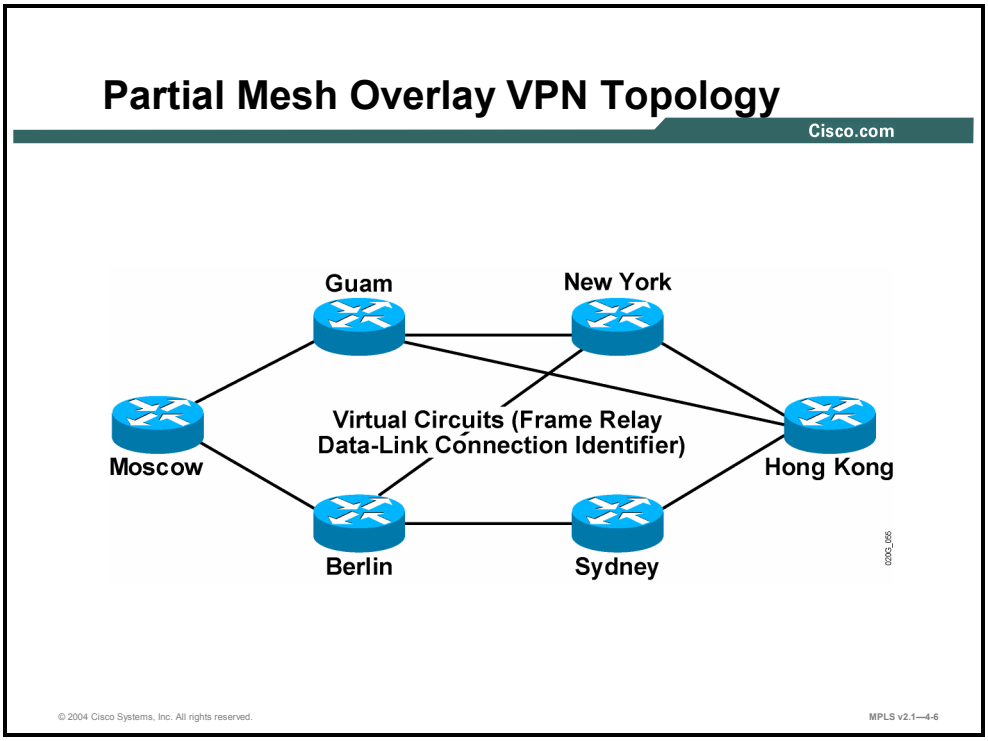

A partial mesh topology is used in environments where cost or complexity factors prevent a full mesh topology between customer sites. The virtual circuits in a partial mesh topology can be established based on the following wide range of criteria:

- Traffic pattern between sites
- Availability of physical infrastructure
- Cost considerations

What Are the VPN Business Categories?

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

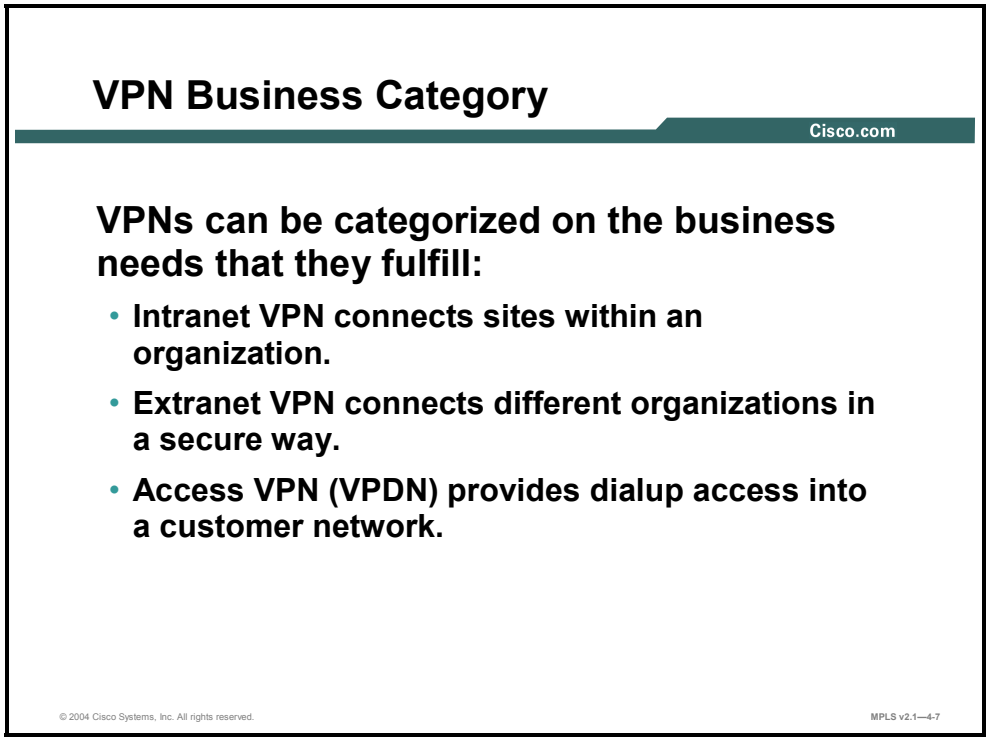

Here is a list of some very popular VPN categories that classify VPNs based on the business needs that they fulfill:

- **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 the extranet component of the VPN business category.

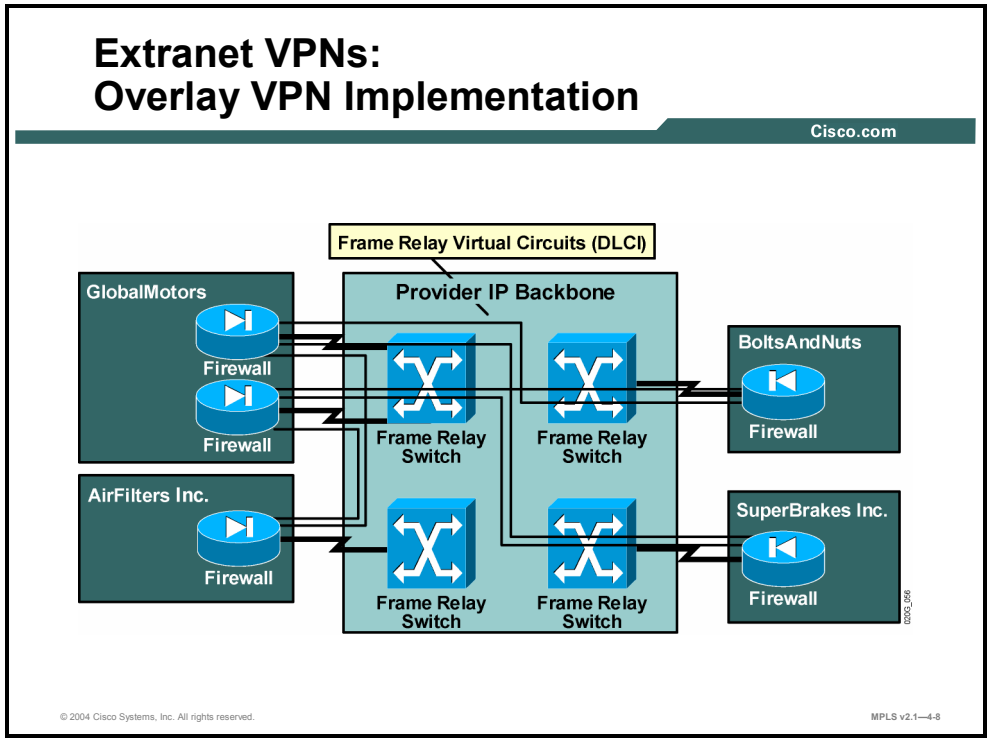

In an overlay implementation of an extranet, organizations are linked with dedicated virtual circuits.

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 the following 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.

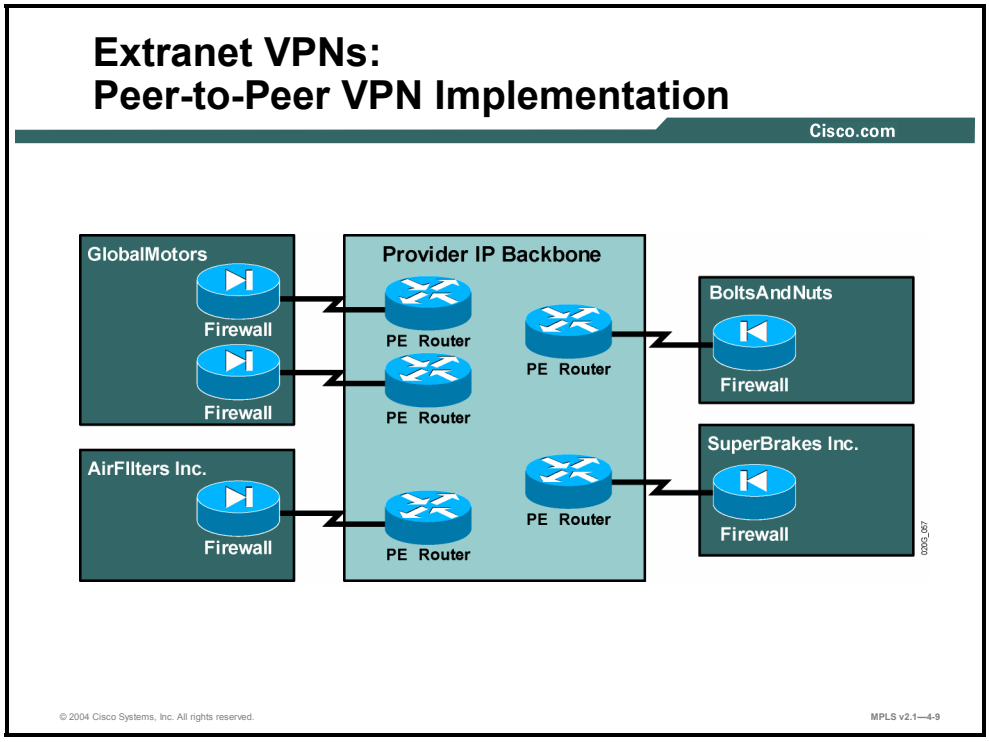

Peer-to-peer VPN implementation of an extranet VPN is very simple compared with overlay VPN implementation—all sites are connected to the P-network, and optimum routing between sites is enabled by default.

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 Is the VPN Connectivity Category?

This topic identifies the major components of the VPN connectivity category.

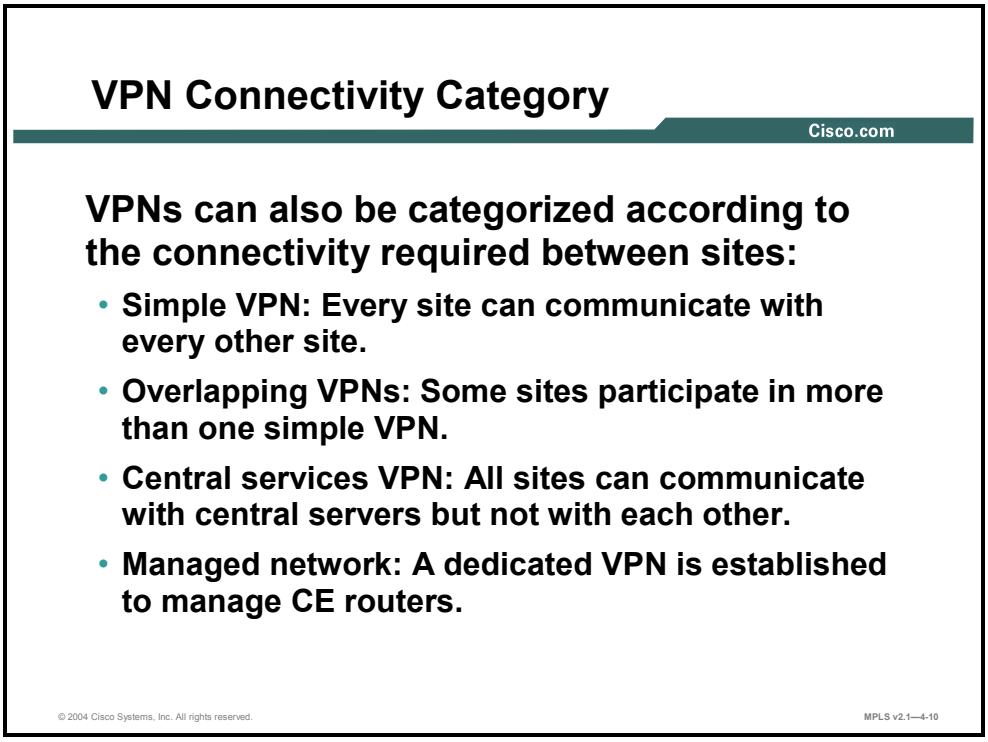

The VPNs discussed so far have usually been very simple in terms of connectivity, as described here:

- 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.

The following describes a number of advanced VPN topologies with more complex connectivity requirements:

- Overlapping VPNs, 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 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 the central services extranet component of the VPN connectivity category.

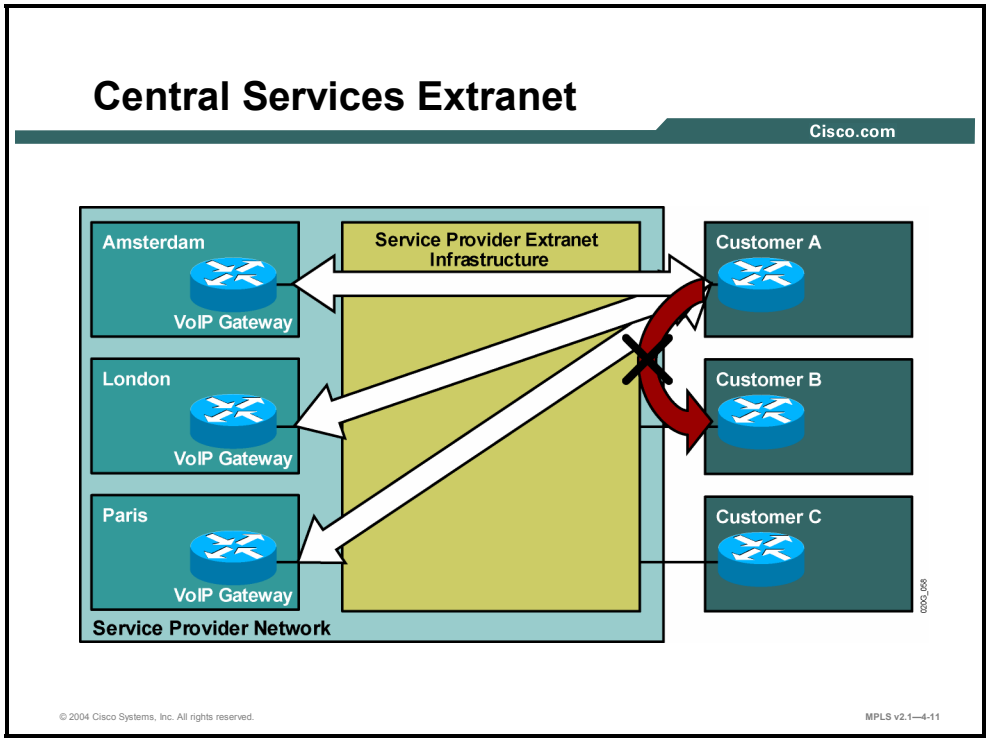

A central services extranet can implement international Voice over IP (VoIP) service.

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.

Example: Hybrid Implementation

The network diagram shows an interesting scenario in which peer-to-peer VPN and overlay VPN implementation can be used together to provide end-to-end service to the customer.

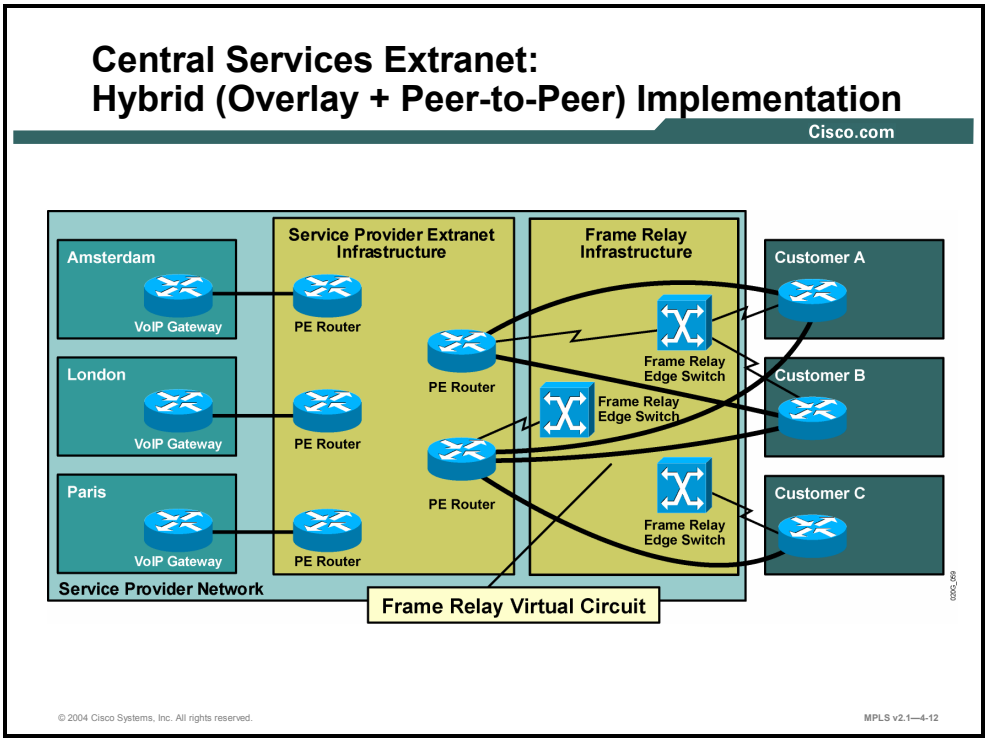

The VoIP service is implemented with a central services extranet topology, which is in turn implemented with a peer-to-peer VPN. Connectivity between PE routers in the peer-to-peer VPN and customer routers is implemented with an overlay VPN based on Frame Relay. The PE routers of the peer-to-peer VPN and the CE routers act as CE devices of the Frame Relay network.

What Is a Managed Network Implementation?

This topic describes the characteristics of the managed network component of the VPN connectivity category.

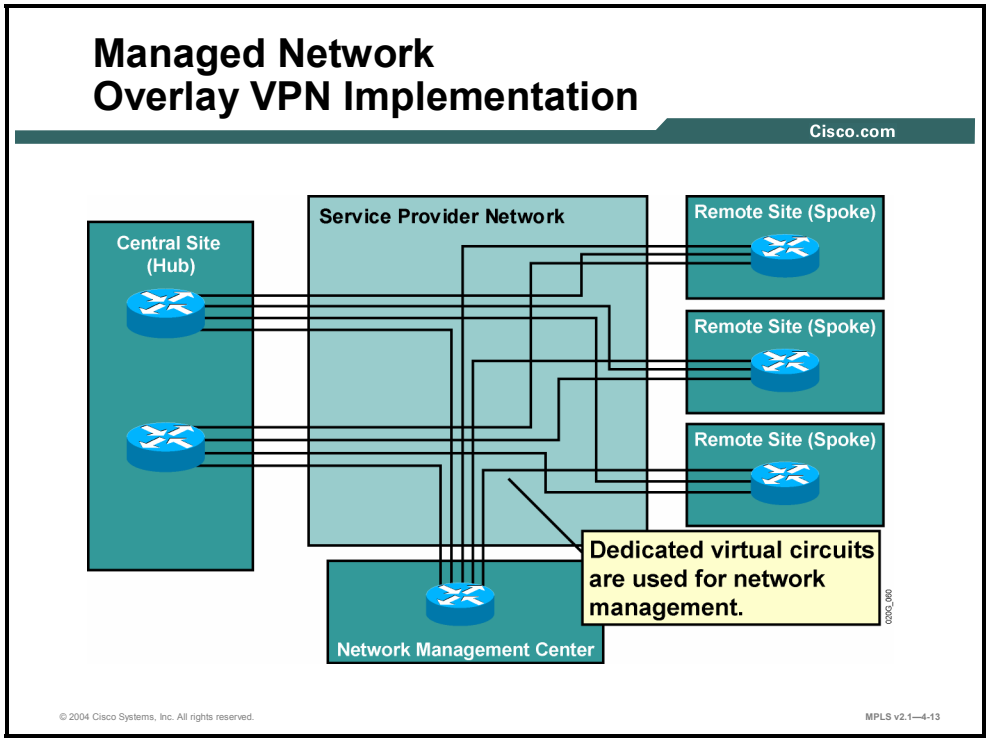

A managed network VPN is traditionally implemented in combination with overlay VPN services. Dedicated virtual circuits are deployed between any managed CE router and the central network management system (NMS) router to which the NMS is connected.

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.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- There are four major VPN topologies: hub-and-spoke, partial mesh, full mesh, and multilevel.
- In the hub-and-spoke topology, all remote sites are linked with a central CE router via virtual circuits. More than one virtual circuit is used in this topology.
- A partial mesh topology is used in environments where cost or complexity factors prevent a full mesh topology between customer sites.
- There are three VPN business categories: intranet VPN, extranet VPN, and access VPN.

Summary (Cont.)

Cisco.com

- In an extranet VPN, organizations are linked with dedicated virtual circuits.
- There are four VPN connectivity categories: simple VPN, overlapping VPN, central service VPN, and managed network.
- A central services extranet enables customers to access common servers for services.
- Managed networks allow customer CE devices to be owned and managed by the service provider.

Lesson 4

Introducing MPLS VPN Architecture

Overview

This lesson explains the MPLS 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 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 P-network
- Describe the features of route distinguishers
- Describe the features of route targets
- Describe how complex VPNs have redefined the meaning of VPNs
- Describe the impact of complex VPN topologies on virtual routing tables

What Is the MPLS VPN Architecture?

This topic describes the features of the MPLS VPN architecture.

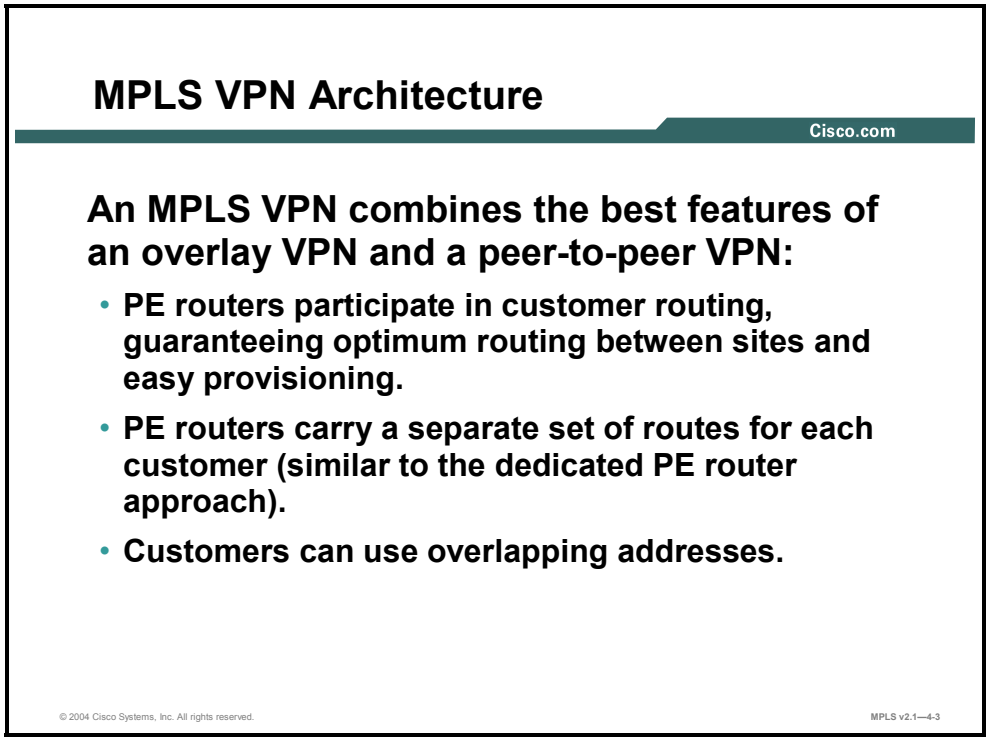

The MPLS VPN architecture offers service providers a peer-to-peer VPN architecture that combines the best features of overlay VPNs (support for overlapping customer address spaces) with the best features of peer-to-peer VPNs. The following describes these characteristics:

- PE routers participate in customer routing, guaranteeing optimum routing between customer sites.
- PE routers carry a separate set of routes for each customer, resulting in perfect isolation between customers.
- Customers can use overlapping addresses.

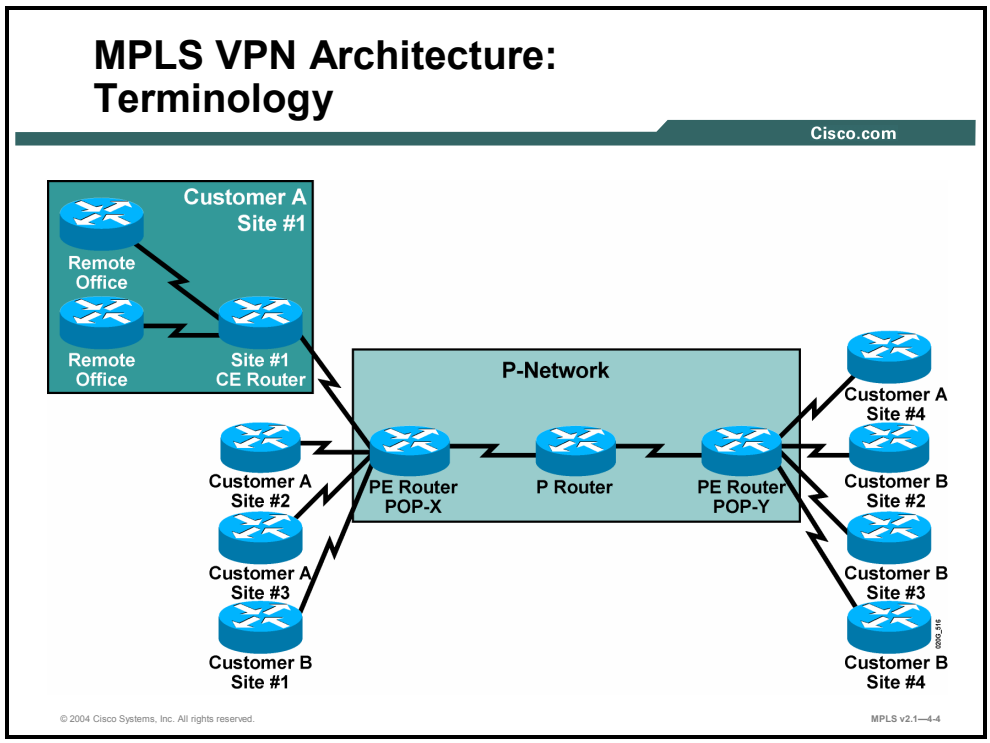

MPLS VPN terminology divides the overall network into a customer-controlled part (the C-network) and a provider-controlled part (the P-network). Contiguous portions of the C-network are called sites and are linked with the P-network via CE routers. The CE routers are connected to the PE routers, which serve as the edge devices of the P-network. The core devices in the P-network, the P routers, provide transit transport across the provider backbone and do not carry customer routes.

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.

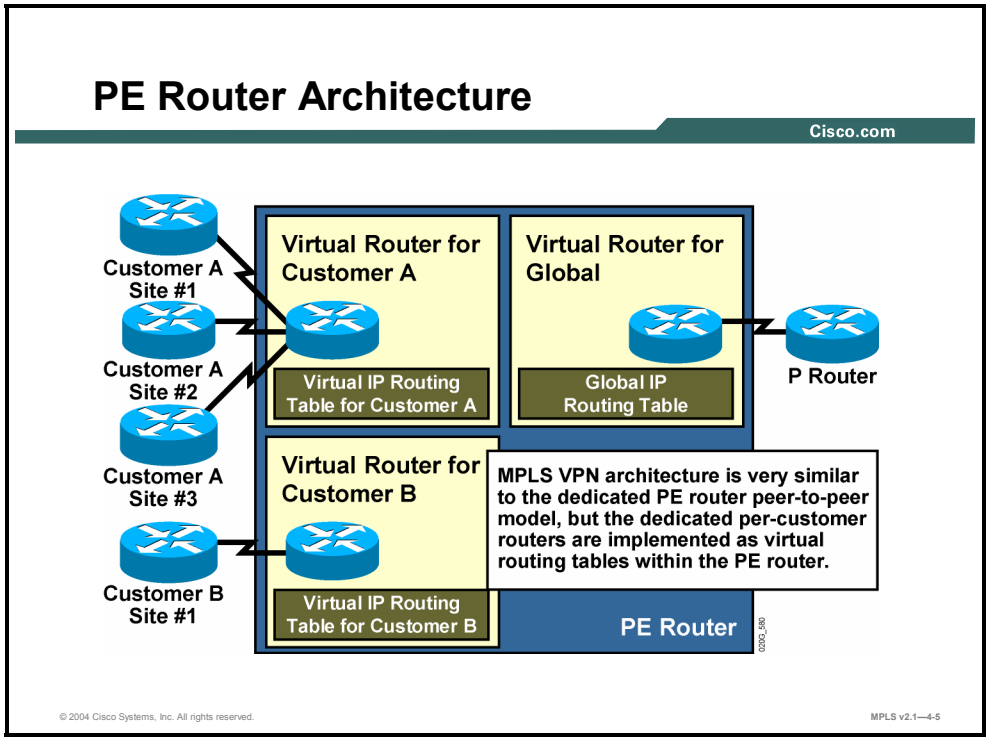

The architecture of a PE router in an MPLS VPN is very similar to the architecture of a POP in the dedicated PE router peer-to-peer model. The only difference is that the whole architecture is condensed into one physical device. Each customer is assigned an independent routing table (virtual routing table) that corresponds to the 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. 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 Propagation Across the P-Network?

This topic describes the different methods of propagating routing information across the P-network.

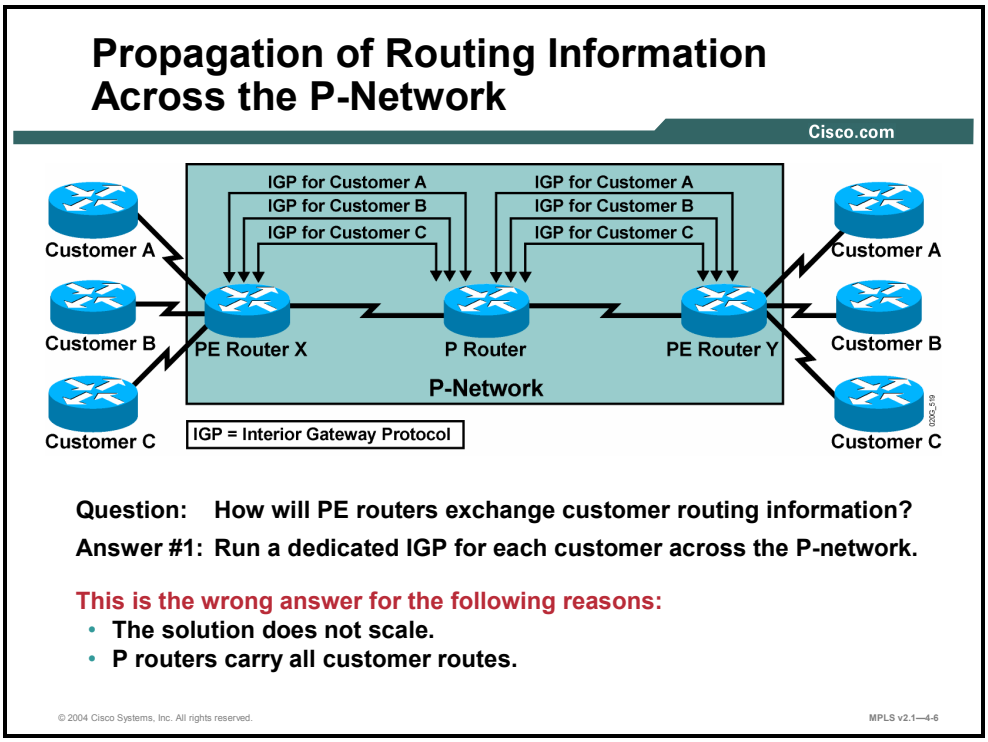

Although virtual routing tables provide isolation between customers, the data from these routing tables still needs to be exchanged between PE routers to enable data transfer between sites attached to different PE routers. Therefore, a routing protocol is needed that will transport all customer routes 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. Both require a per-customer routing protocol be run between PE routers. In one implementation, the P routers participate in customer routing and pass the customer routing information between PE routers. In the other implementation, the PE routers are connected via point-to-point tunnels, for example IPSEC, thereby hiding the customer routing from the P routers.

This solution, although very simple to implement (and often used by some customers), 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.

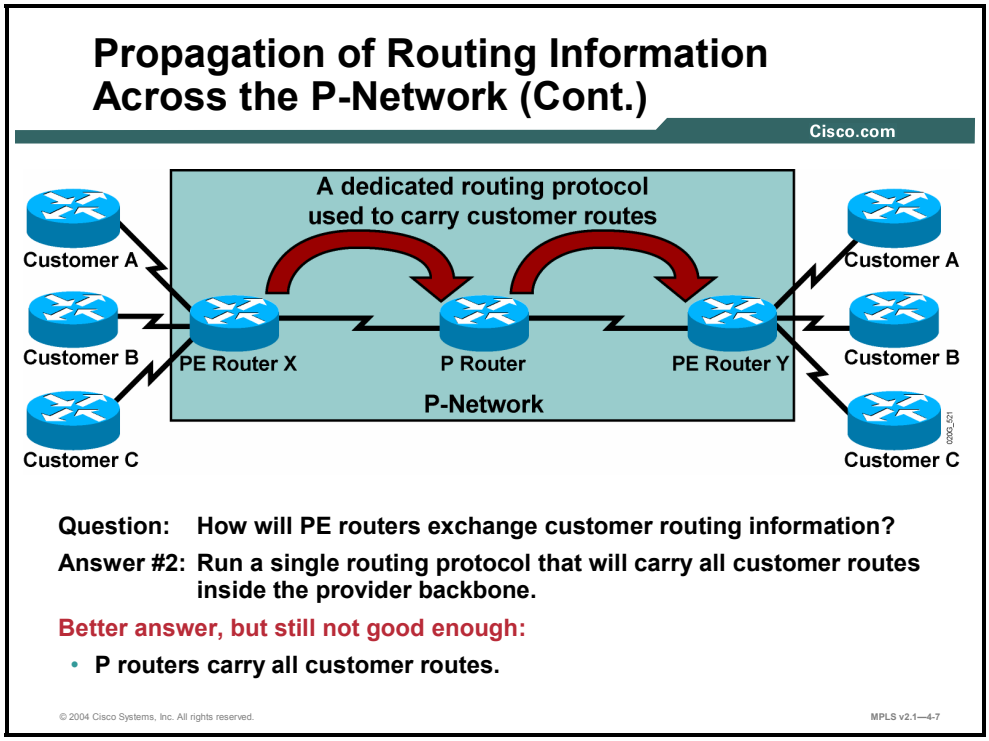

A better approach to the route propagation problem is to deploy a single routing protocol that can exchange all customer routes across the P-network. Although this approach is better than the previous one, the P routers are still involved in customer routing; therefore, the proposal retains some of the same scalability issues of the previous one.

Propagation of Routing Information Across the P-Network (Cont.)

Cisco.com

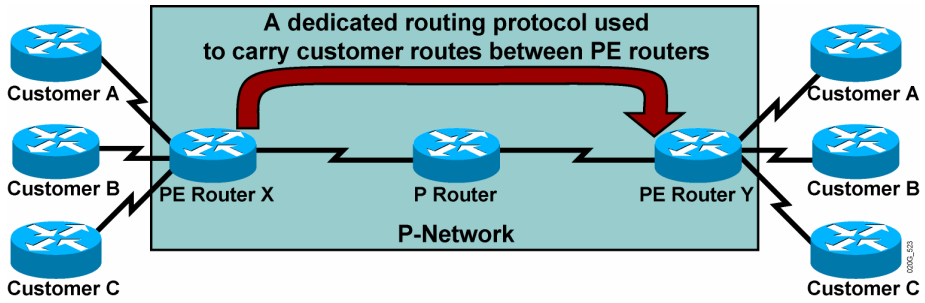

Question: How will PE routers exchange customer routing information?

Answer #3: Run a single routing protocol that will carry all customer routes between PE routers. Use MPLS labels to exchange packets between PE routers.

The best answer:

- P routers do not carry customer routes; the solution is scalable.

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.

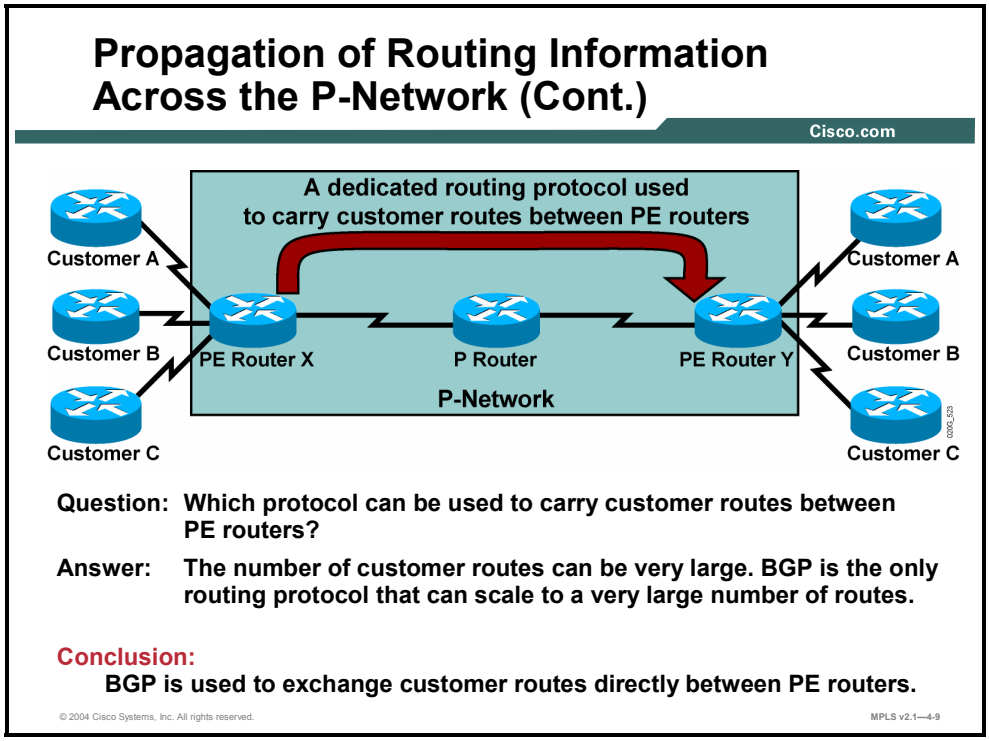

The next design decision to be made is the choice of the routing protocol running between PE routers. Given that the total number of customer routes is expected to be very large, the only well-known protocol with the required scalability is BGP. In fact, BGP is used in MPLS VPN architecture to transport customer routes directly between PE routers.

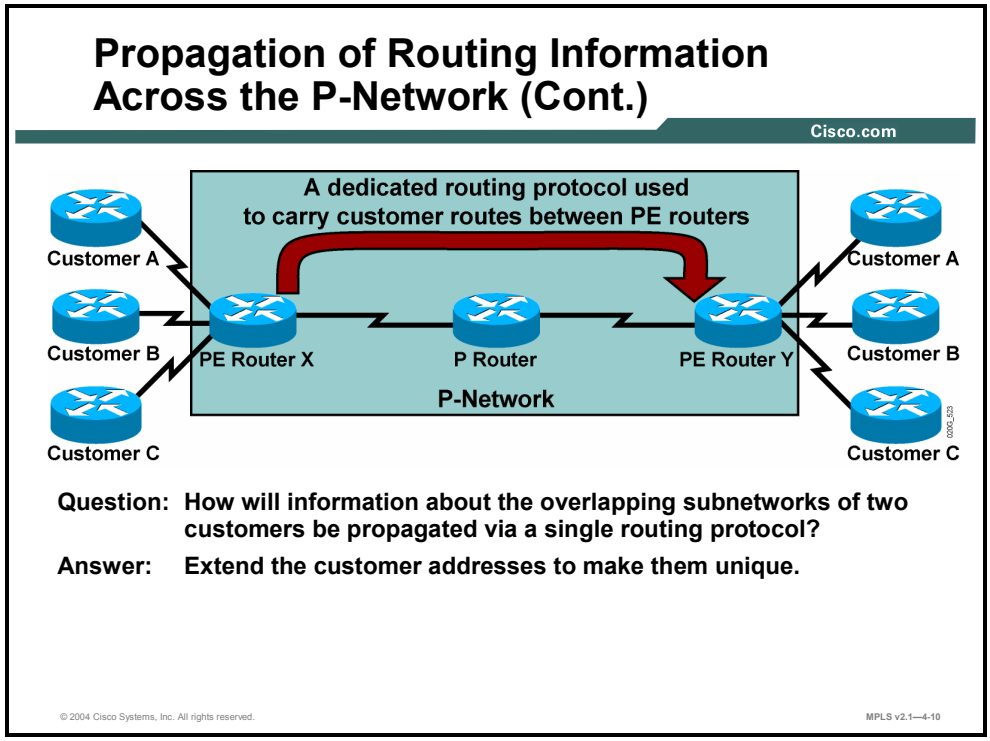

MPLS VPN architecture differs in an important way from traditional peer-to-peer VPN solutions—the support of overlapping customer address spaces.

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 nonunique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE routers.

What Are Route Distinguishers?

This topic describes the features of RDs.

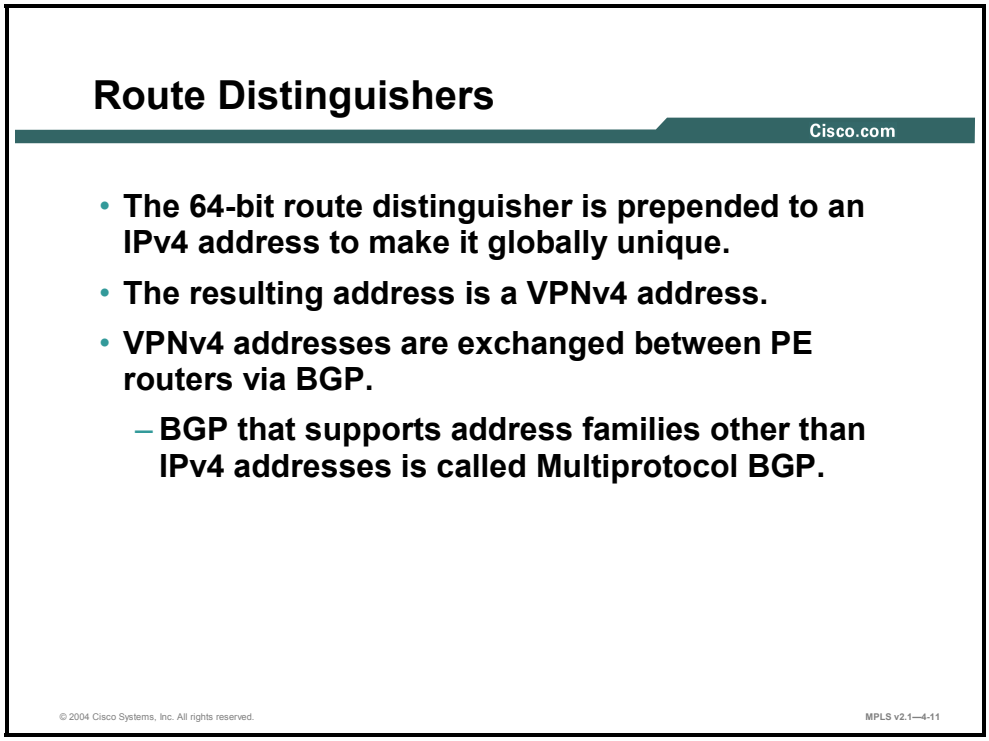

The RD is used only to transform nonunique 32-bit customer IP version 4 (IPv4) addresses into unique 96-bit VPNv4 addresses (also called VPN IPv4 addresses).

VPNv4 addresses are exchanged only between PE routers; they are never used between CE routers. BGP between PE routers must therefore support the exchange of traditional IPv4 prefixes and the exchange of VPNv4 prefixes. A BGP session between PE routers is consequently called a Multiprotocol BGP (MP-BGP) session.

Note	Initial MPLS VPN implementation in Cisco IOS software supports only MPLS VPN services within a single autonomous system (AS). In such a scenario, the BGP session between PE routers is always an IBGP session.

Route Distinguishers (Cont.)

Cisco.com

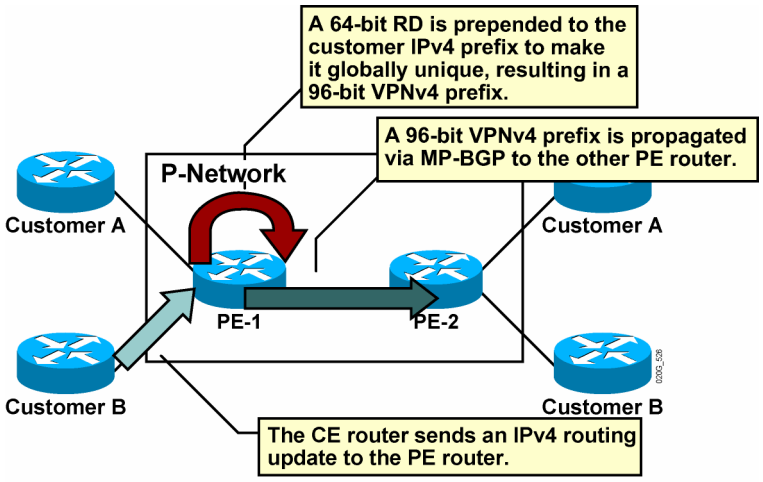

Customer route propagation across an MPLS VPN network is done using the following process:

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.

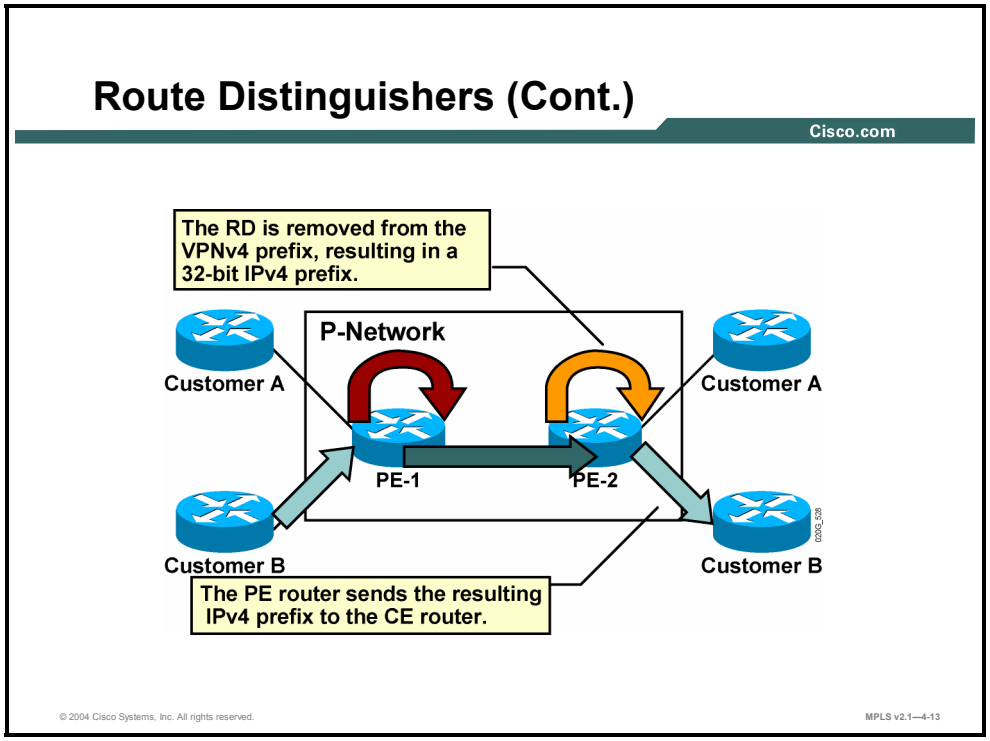

Step 4 The receiving PE routers strip the RD from the VPNv4 prefix, resulting in an IPv4 prefix.

Step 5 The IPv4 prefix is forwarded to other CE routers within an IPv4 routing update.

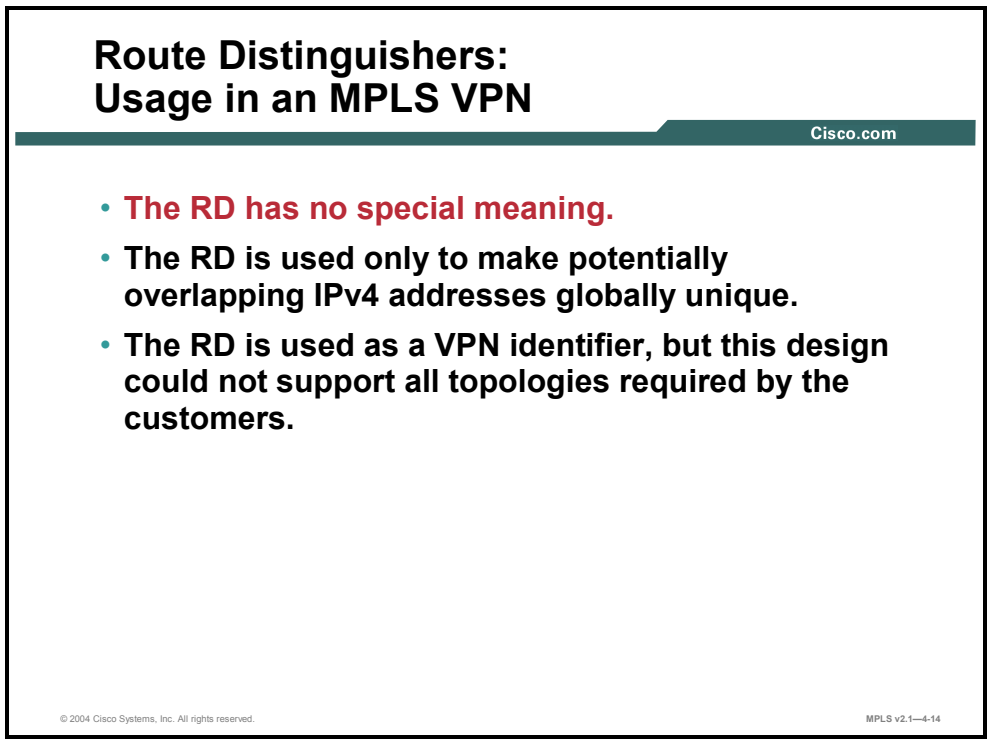

The RD has no special meaning or role in MPLS VPN architecture; its only function is to make overlapping IPv4 addresses globally unique.

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 CPE 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.

What Are Route Targets?

This topic describes the features of RTs.

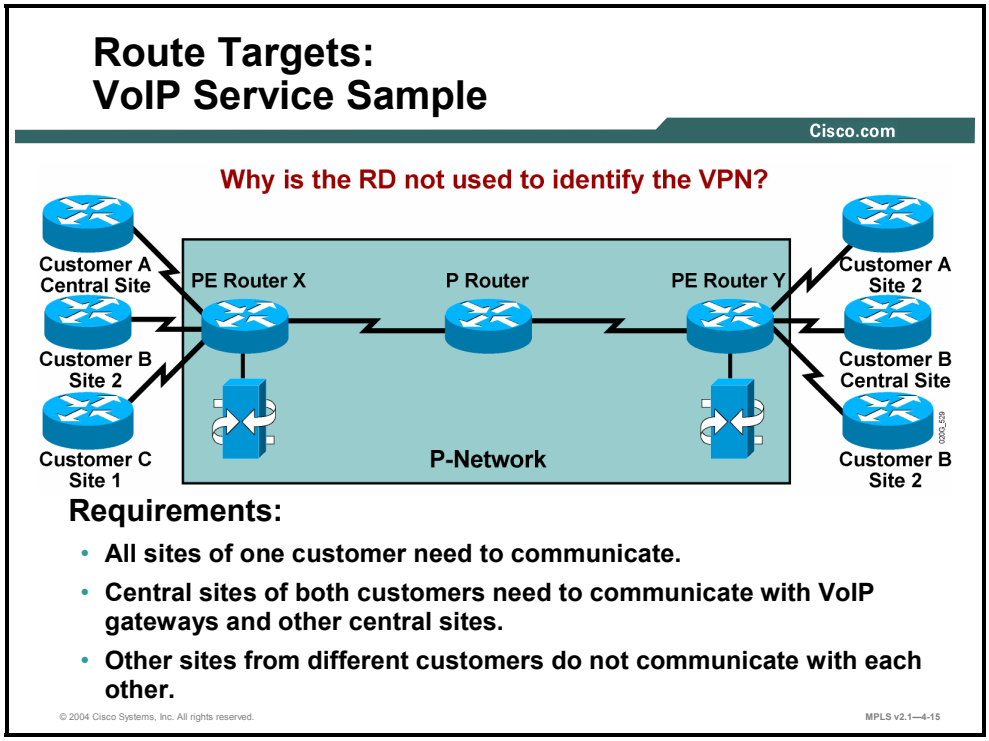

To illustrate the need for a more versatile VPN indicator than the RD, consider the VoIP service.

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.

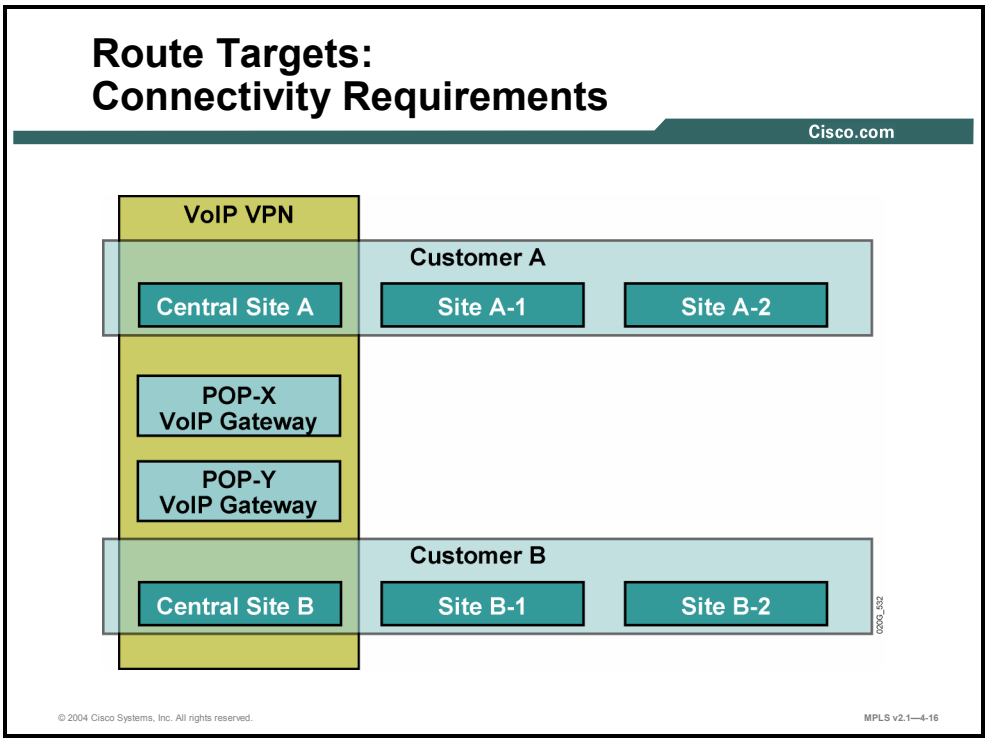

Three VPNs are needed to implement the desired connectivity: two customer VPNs and a shared VoIP VPN. Central customer sites participate in the customer VPN and in the VoIP VPN.

Route Targets: Why Are They Needed?

Cisco.com

- Some sites have to participate in more than one VPN.
- The RD cannot identify participation in more than one VPN.
- RTs were introduced in the MPLS VPN architecture to support complex VPN topologies.
  - A different method is needed in which a set of identifiers can be attached to a route.

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.

RTs were introduced into the MPLS VPN architecture to support this requirement.

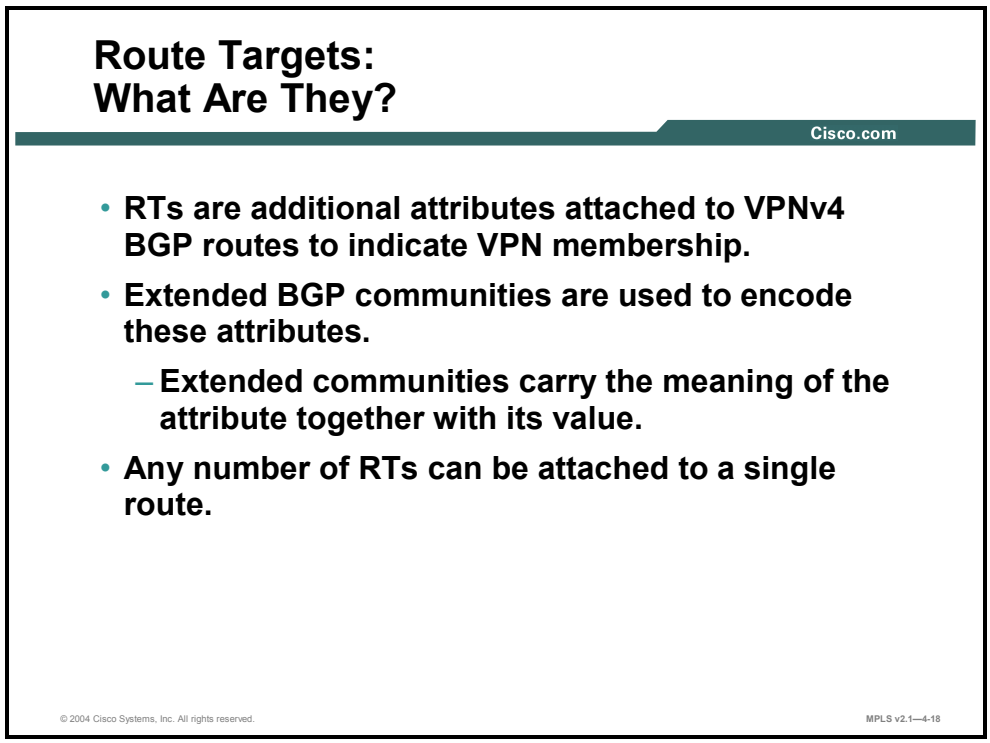

RTs are attributes that are attached to a VPNv4 BGP route to indicate its VPN membership. The extended BGP communities of routing updates are used to carry the RT of that update, thus identifying to which VPN the update belongs.

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.

Route Targets: How Do They Work?

Cisco.com

- **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.

What Is the New Meaning of VPNs?

This topic describes how complex VPNs have redefined the meaning of VPNs.

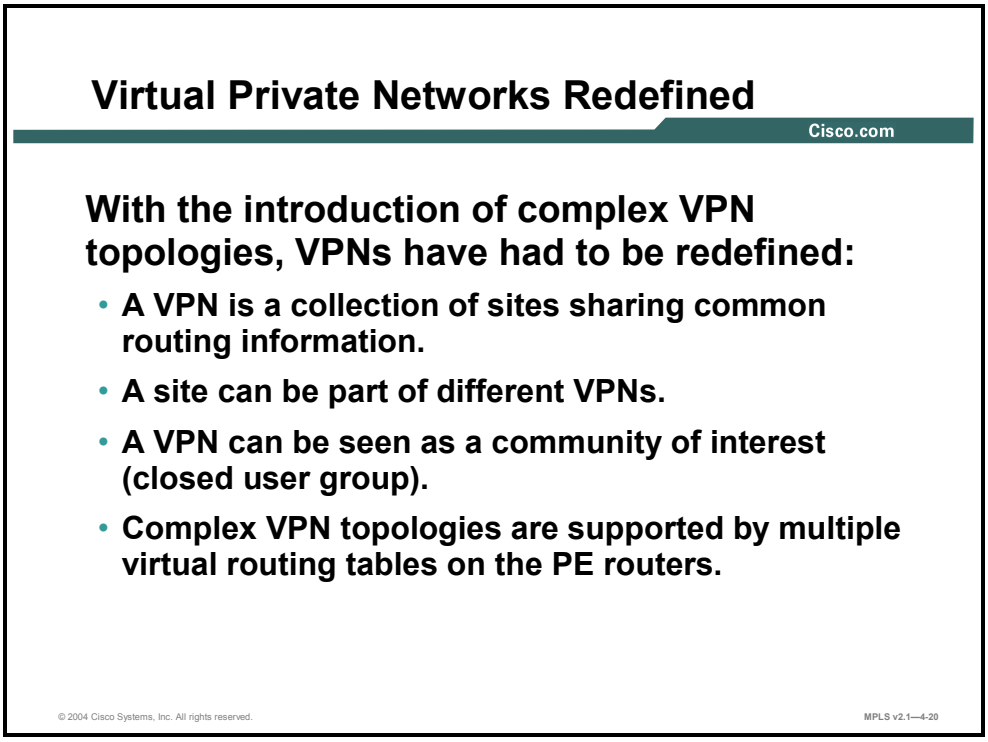

With the introduction of complex VPN topologies, the definition of a VPN has needed to be changed. A VPN is simply a collection of sites sharing common routing information. In traditional switched WAN terms (for example, in X.25 terminology), such a concept would be called a closed user group (CUG).

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.

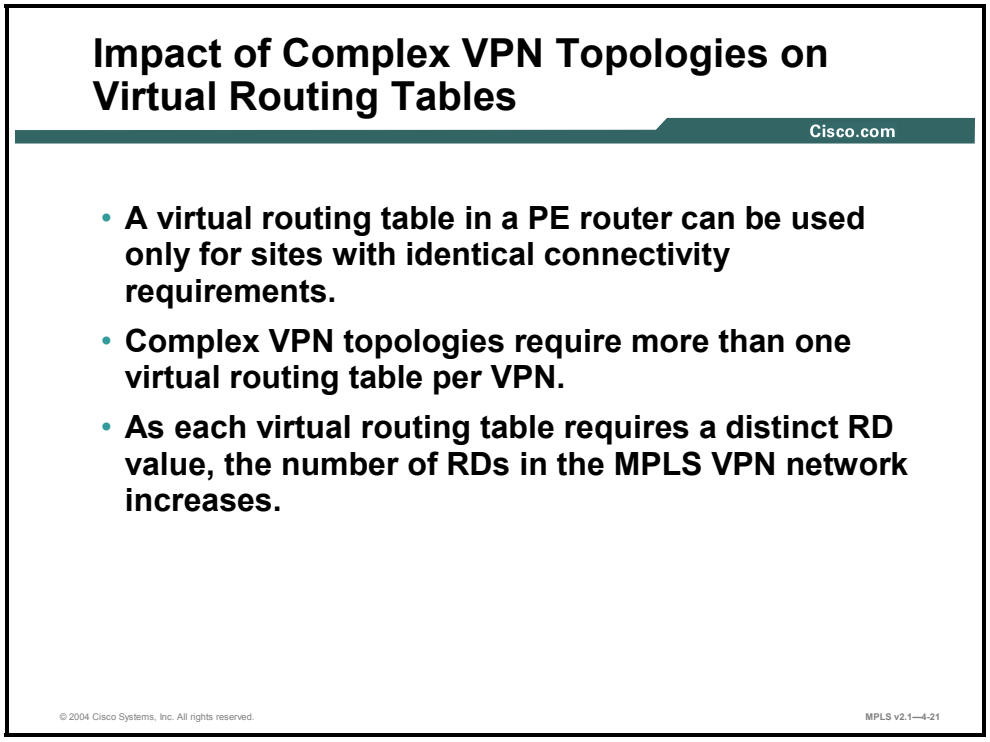

A single virtual routing table can be used only for sites with identical connectivity requirements. Complex VPN topologies, therefore, require more than one virtual routing table per VPN.

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).

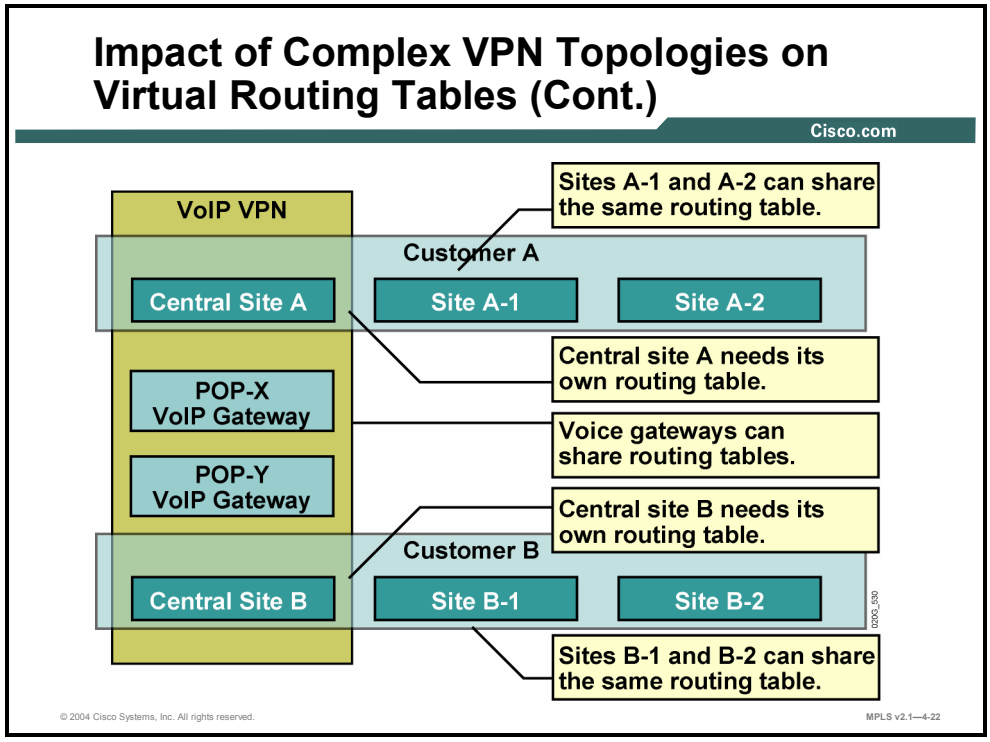

The virtual routing table needs of this service are as follows:

- All sites of customer A (apart from the central site) can share the same virtual routing table because they belong to a single VPN.
- The same is true for all sites of customer B (apart from the central site).
- The VoIP gateways participate only in the VoIP VPN and can belong to a single virtual routing table.
- 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.
- 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 discussed in this lesson.

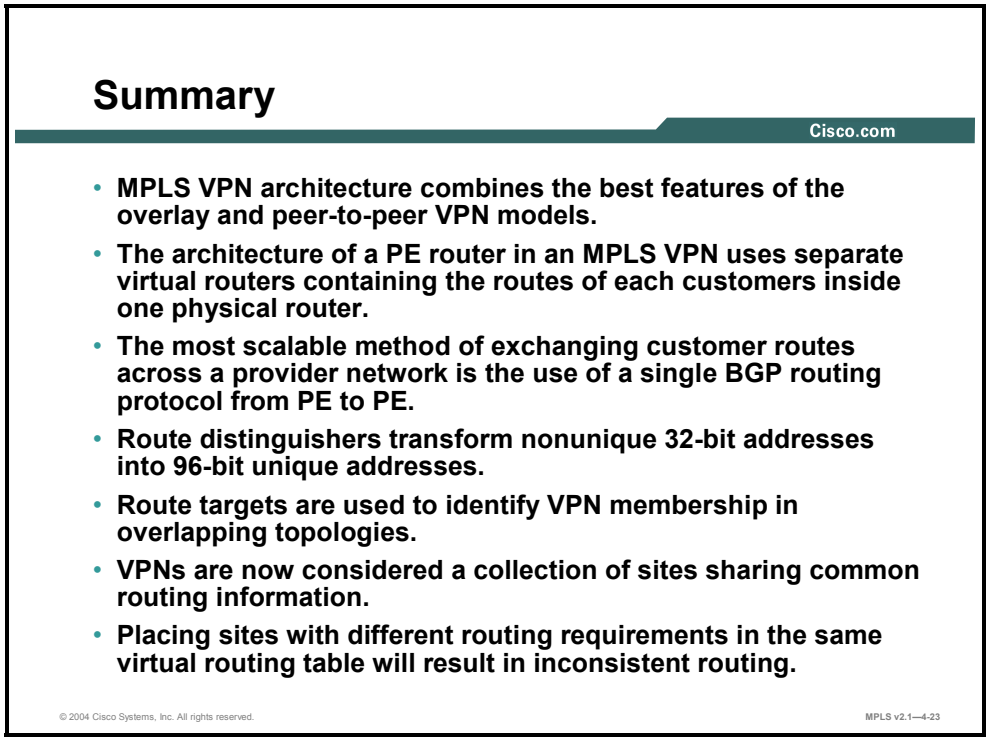

Lesson 5

Introducing the MPLS VPN Routing Model

Overview

This lesson explains the routing requirements for MPLS VPNs. The lesson offers address and routing perspectives from the customer and service provider side, and it discusses how routing tables appear on PE routers. This lesson also discusses MPLS VPN end-to-end information flow, 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

This topic describes the routing requirements for MPLS VPNs.

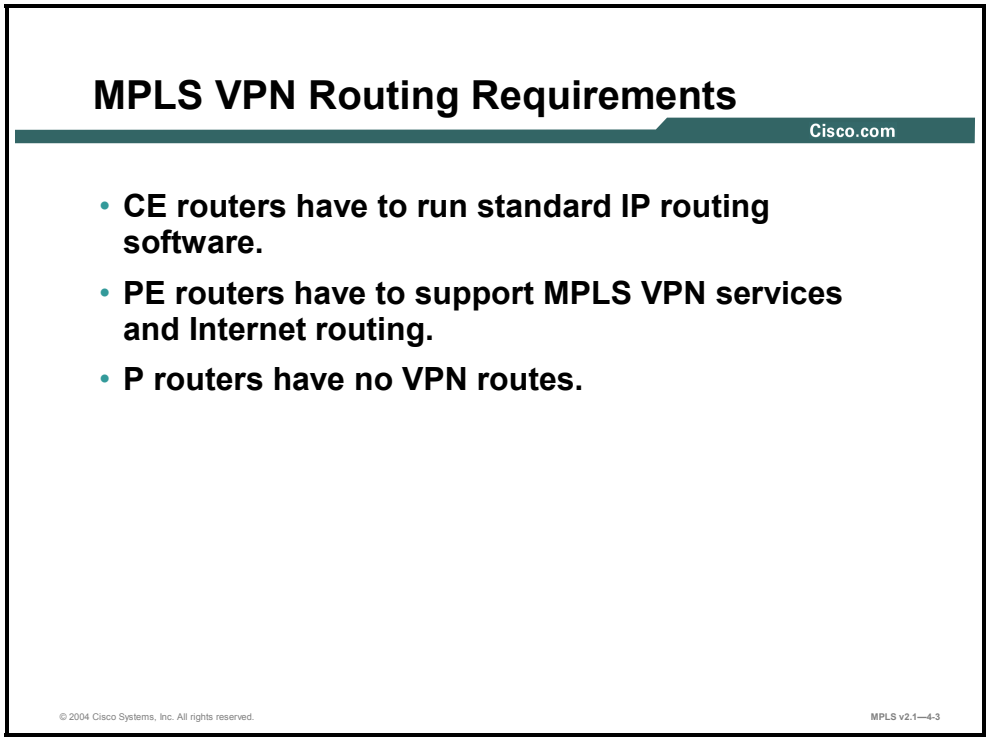

The designers of MPLS VPN technology were faced with the following routing requirements:

- 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, P routers must not carry VPN routes.

What Is the MPLS VPN Routing Model?

This topic describes the MPLS VPN routing model for CE routers, PE routers, and P routers.

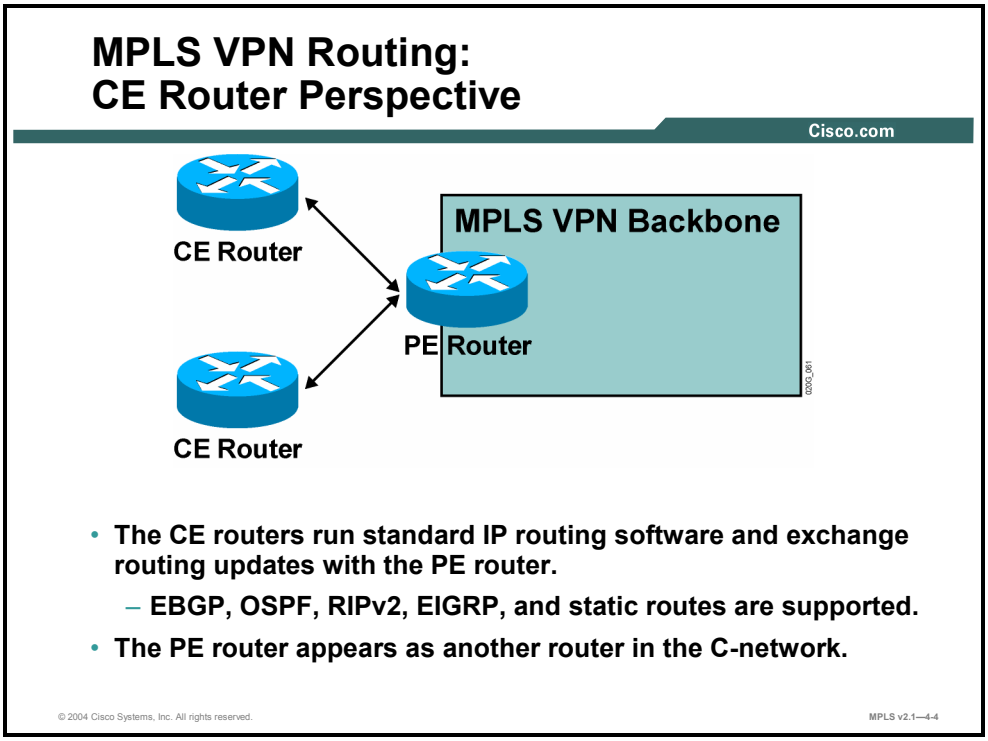

The MPLS VPN backbone should look like a standard corporate backbone to the CE routers. The CE routers run standard IP routing software and exchange routing updates with the PE routers, which appear to them as normal routers in the C-network.

Note	In Cisco IOS Release 12.2, the choice of routing protocols that can be run between a CE router and a PE router is limited to static routes, RIP version 2 (RIPv2), Open Shortest Path First (OSPF), and External Border Gateway Protocol (EBGP).

MPLS VPN Routing: Overall Customer Perspective

Cisco.com

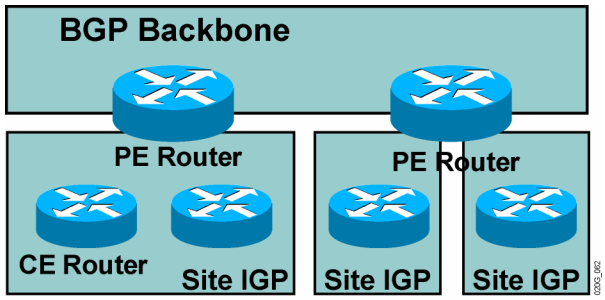

- **To the customer, the PE routers appear as core routers connected via a BGP backbone.**
- **The usual BGP and IGP design rules apply.**
- **The P routers are hidden from the customer.**

From the customer perspective, the MPLS VPN backbone looks like an intracompany 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

Cisco.com

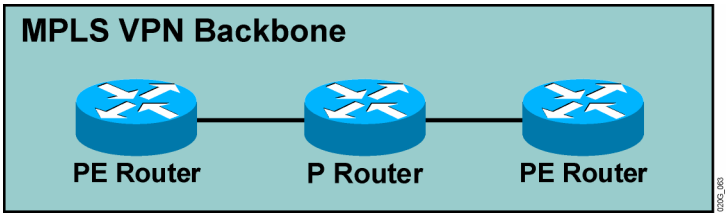

- **P routers do not participate in MPLS VPN routing and do not carry VPN routes.**
- **P routers run backbone IGP with the PE routers and exchange information about global subnetworks (core links and loopbacks).**

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.

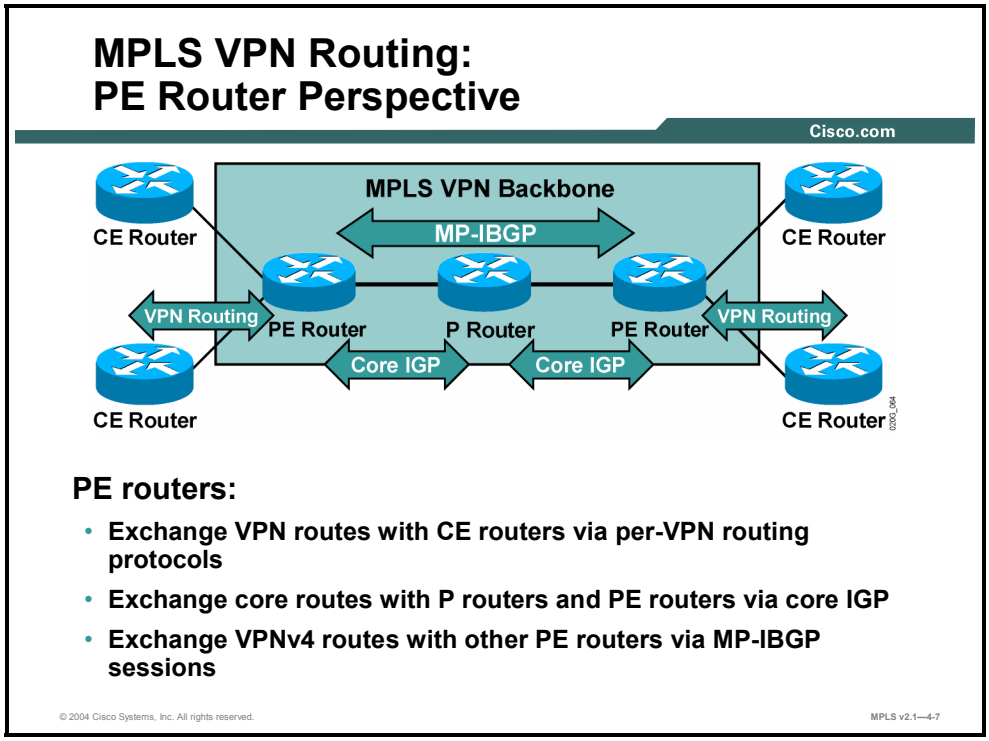

The PE routers are the only routers in MPLS VPN architecture that see all routing aspects of the MPLS VPN. PE routers are able to do the following:

- PE routers exchange IPv4 VPN routes with CE routers via various routing protocols running in the virtual routing tables.
- PE routers exchange VPNv4 routes via 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.

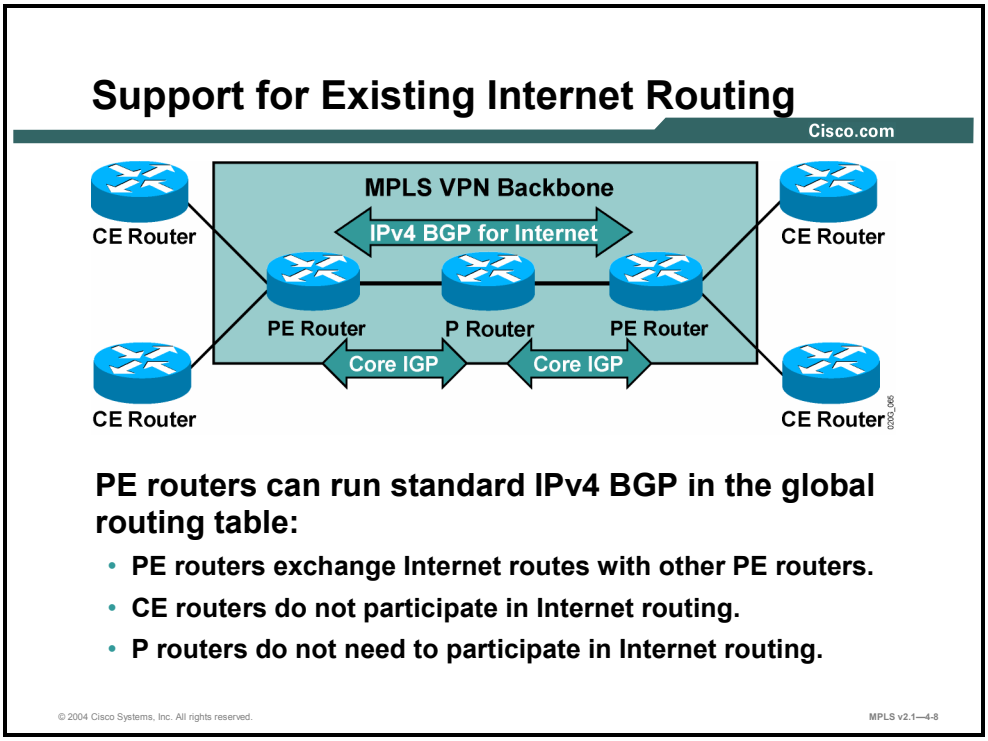

The routing requirements for PE routers also extend to supporting Internet connectivity—PE routers have to exchange Internet routes with other PE routers. The CE routers cannot participate in Internet routing if the Internet routing is performed in global address space. The P routers could participate 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.

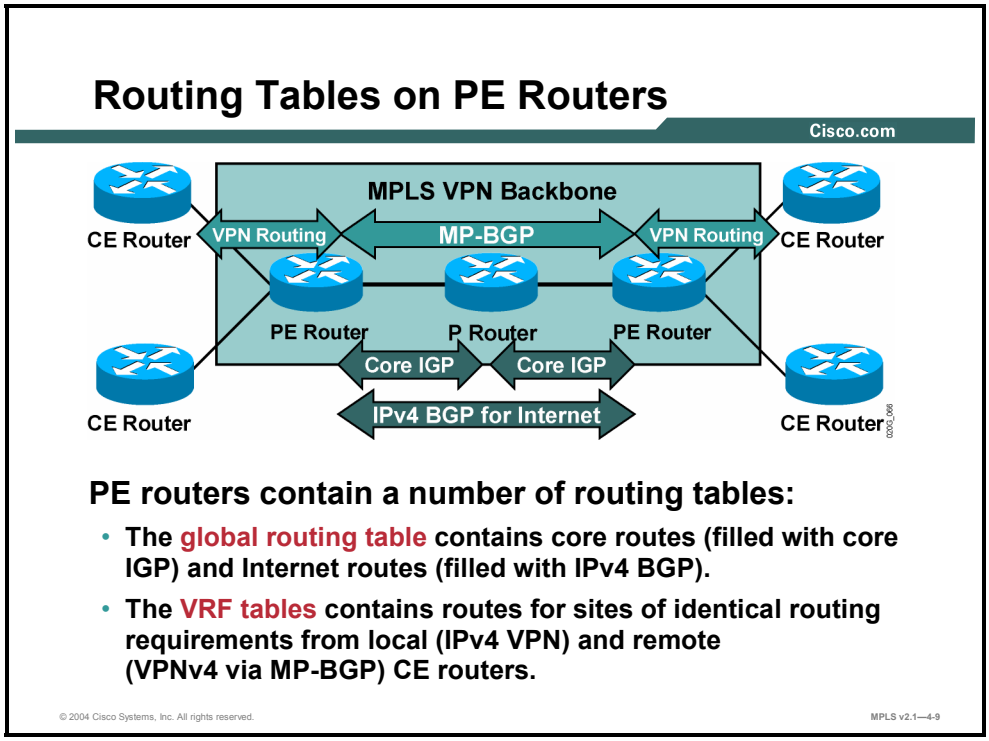

The PE routers fulfill various routing requirements imposed on them by using a number of IP routing tables, such as the following:

- 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 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.

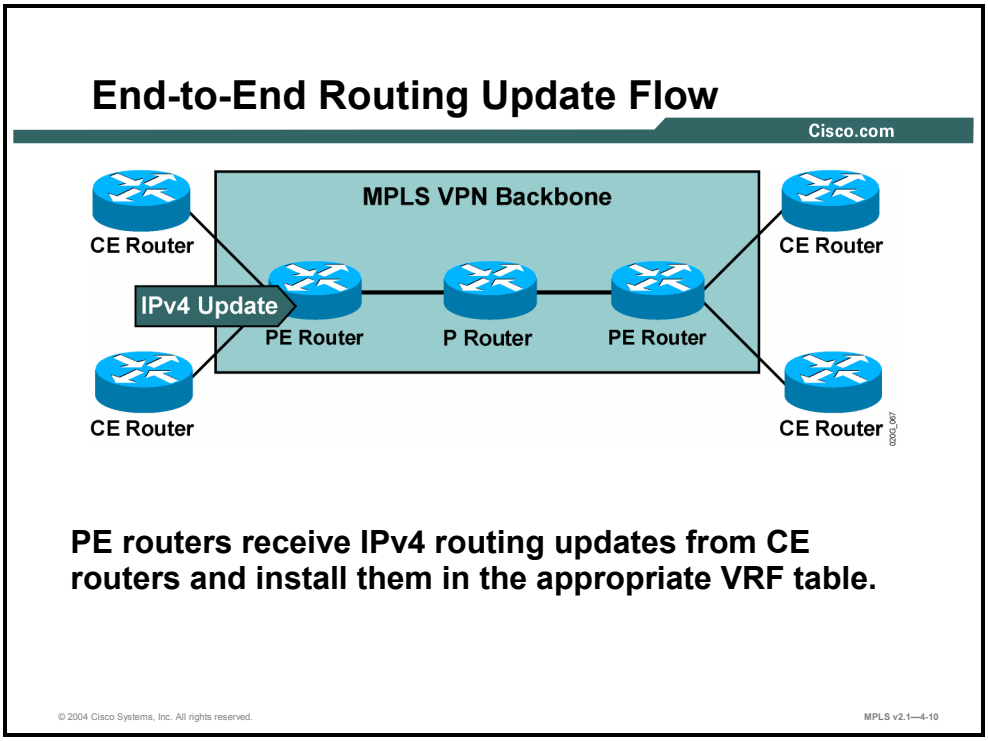

These figures provide an overview of end-to-end routing information flow in an MPLS VPN network.

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.

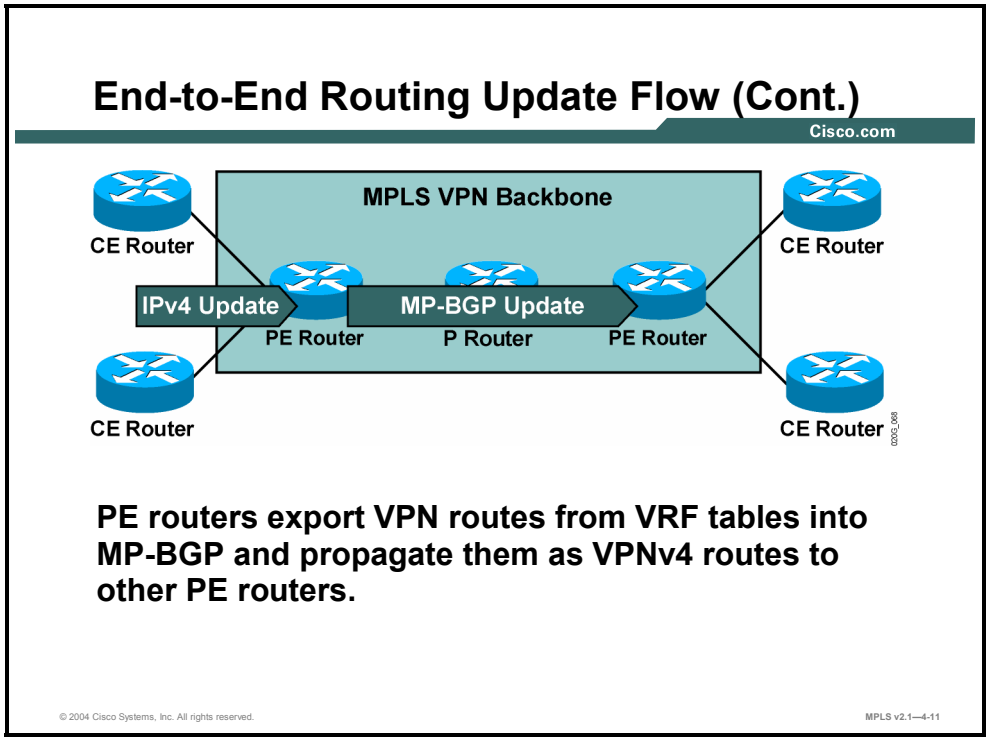

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers.

Initial MPLS VPN implementation in Cisco IOS software (Cisco IOS Releases 12.0 T and 12.1) supports MPLS VPN services only within the scope of a single AS. The MP-BGP sessions between the PE routers are therefore IBGP sessions and are subject to the IBGP split-horizon rules. Thus, either a full mesh of MP-IBGP sessions is required between PE routers, or route reflectors need to be used to reduce the full mesh IBGP requirement.

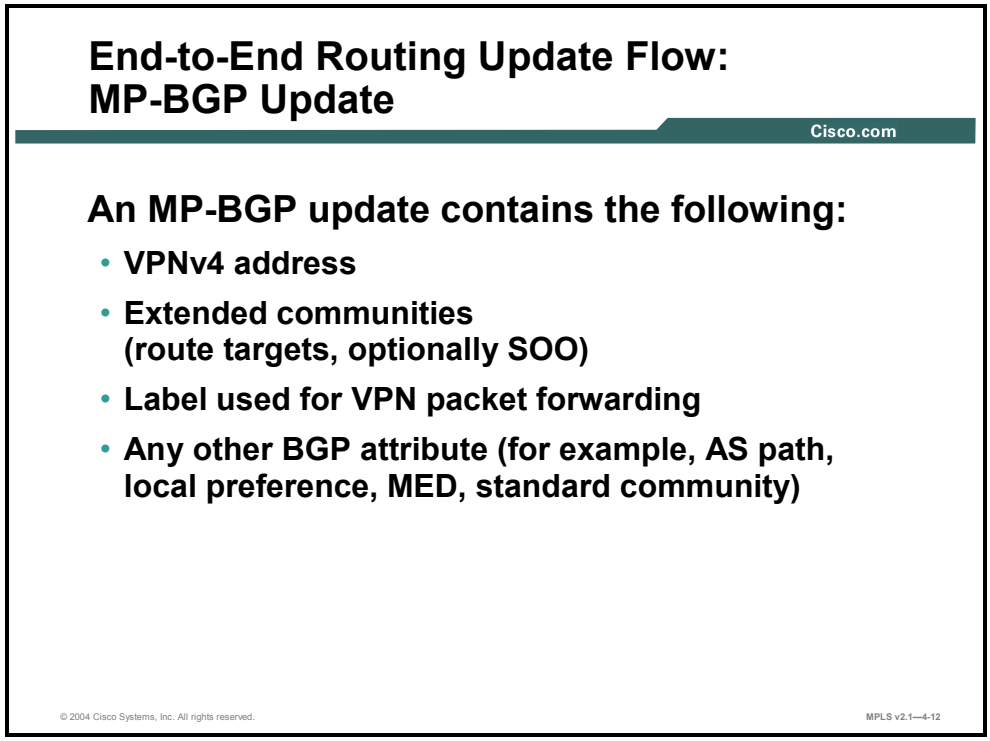

An MP-BGP update exchange between PE routers contains the following:

- VPNv4 address
- Extended BGP communities (RTs required; Site of Origin, or SOO, optional)
- Label used for VPN packet forwarding (The “Forwarding MPLS VPN Packets” lesson explains how the label is used.)
- 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.)

Cisco.com

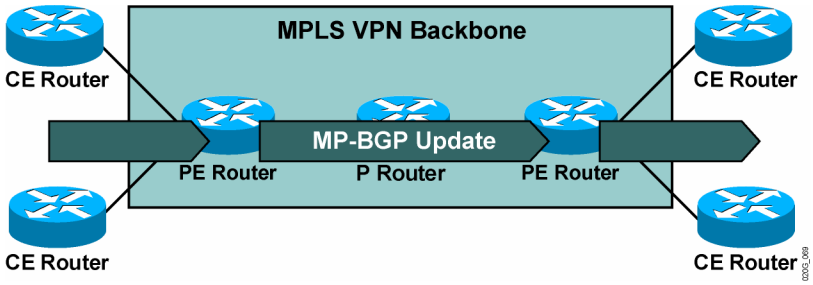

- **The receiving PE router imports the incoming VPNv4 routes into the appropriate VRF based on route targets attached to the routes.**
- **The routes installed in the VRFs are propagated to the CE routers.**

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.

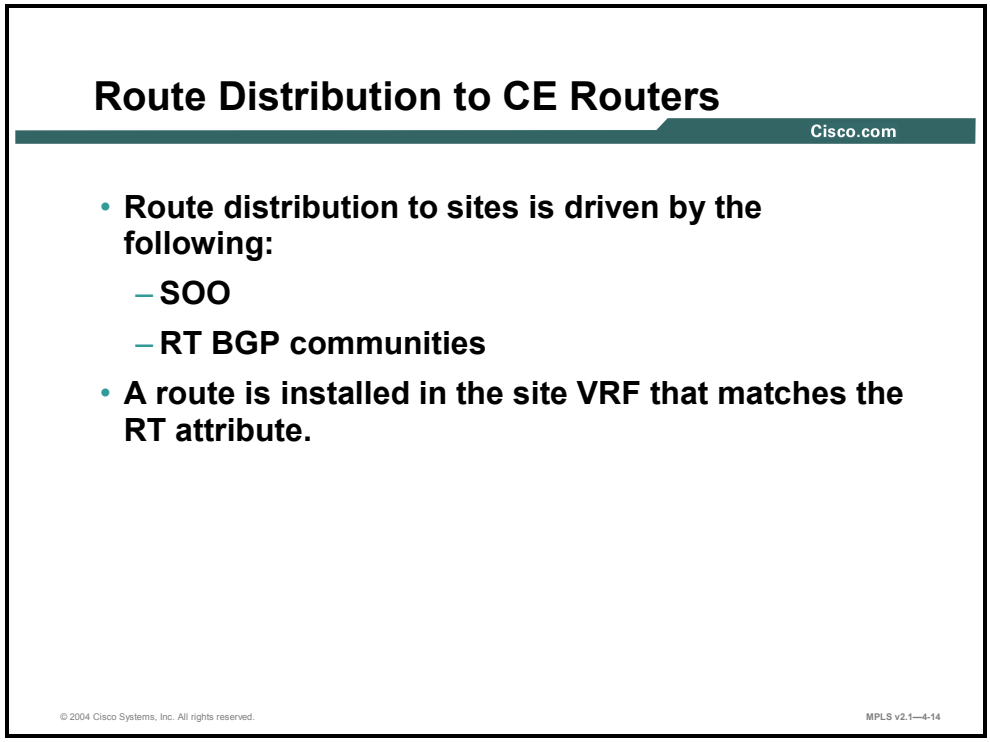

VPNv4 routes are installed into VRFs on the receiving PE router; the incoming VPNv4 route is imported into the VRF only if at least one RT attached to the route matches at least one import RT configured in the VRF.

The SOO attribute attached to the VPNv4 route controls 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. The RTs attached to a route and the import RTs configured in the VRF drive the import of the routes to the CE router.

Summary

This topic summarizes the key points discussed in this lesson.

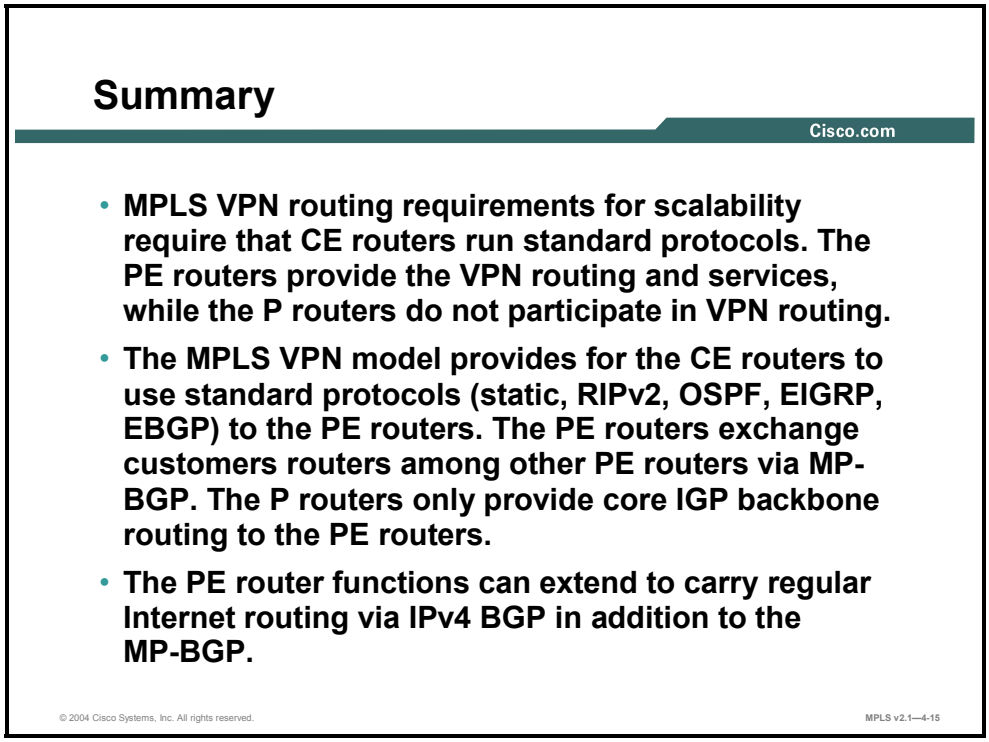

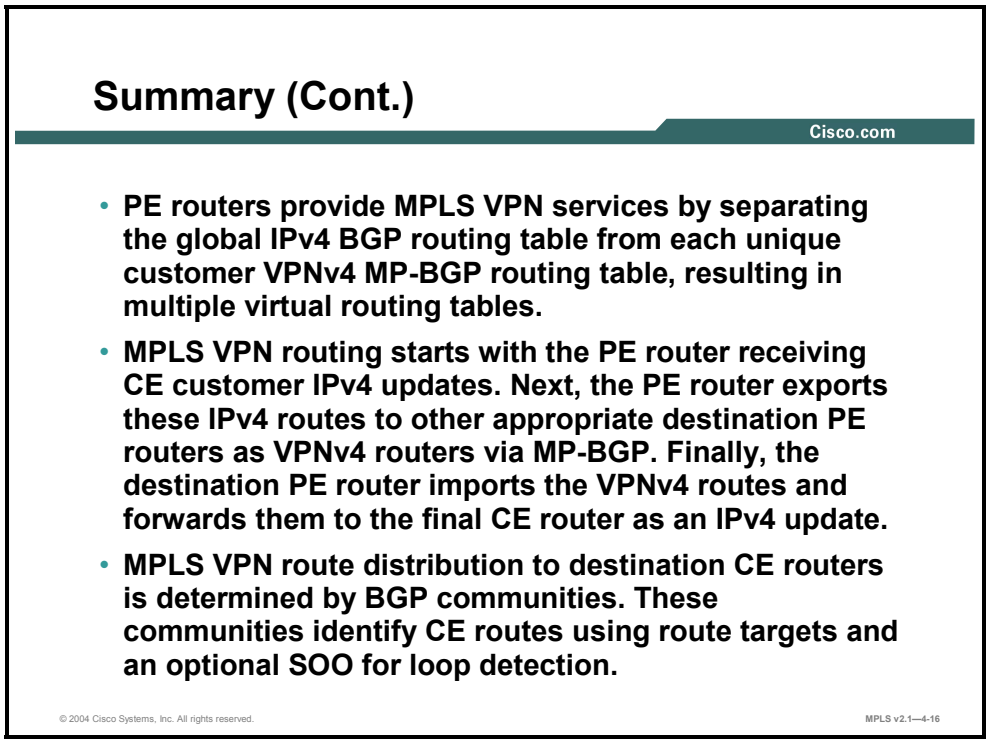

Lesson 6

Forwarding MPLS VPN Packets

Overview

This lesson explains how forwarding across an MPLS 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 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.

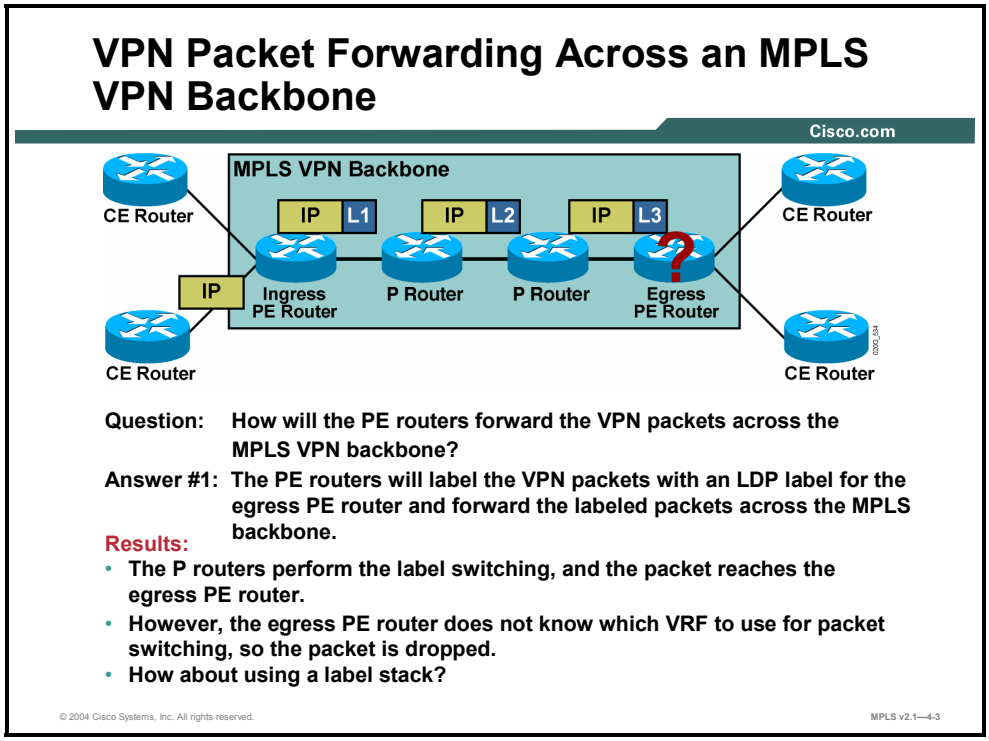

An MPLS-oriented approach to MPLS VPN packet forwarding across the MPLS VPN backbone would be to label the customer packet with the label assigned by Label Distribution Protocol (LDP) for the egress PE router. The core routers consequently would never see the customer IP packet; instead, the core routers would see just a labeled packet targeted toward the egress PE router. The core routers would perform simple label-switching operations, finally delivering the customer packet to the egress PE router. Unfortunately, the customer IP packet would contain no VPN or 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, therefore, have to drop the packet.

VPN Packet Forwarding Across an MPLS VPN Backbone (Cont.)

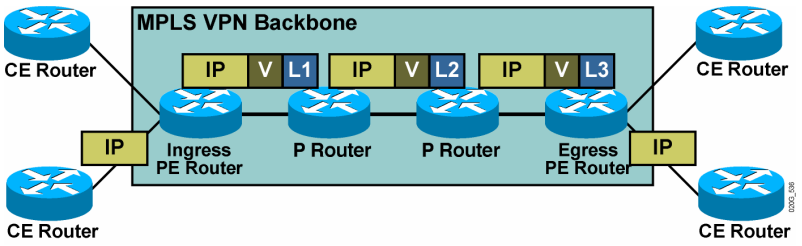

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #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, and the packet reaches the egress PE router.
- 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 the router 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 are used in Cisco IOS software. The second label in the stack points toward an outgoing interface whenever the 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 the following 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.

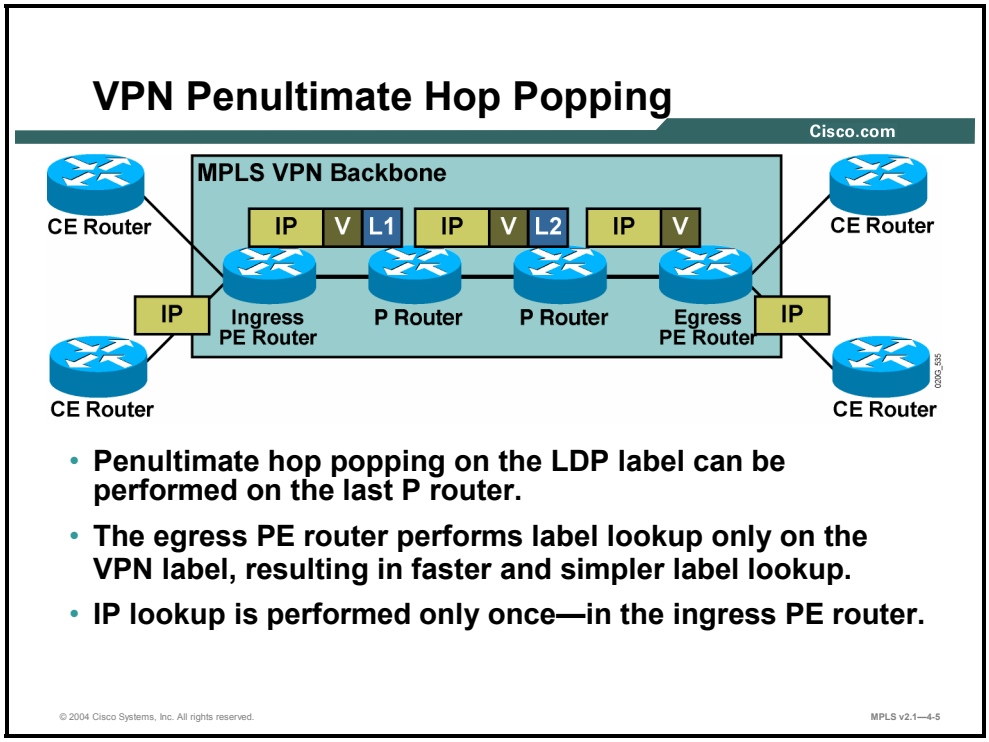

PHP (the removal of the top label in the stack on the hop prior to the egress router) can be performed in frame-based MPLS networks. In these networks, the last P router in the label switched path (LSP) tunnel pops the LDP label (as previously requested by the egress PE router through LDP), and the PE 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.

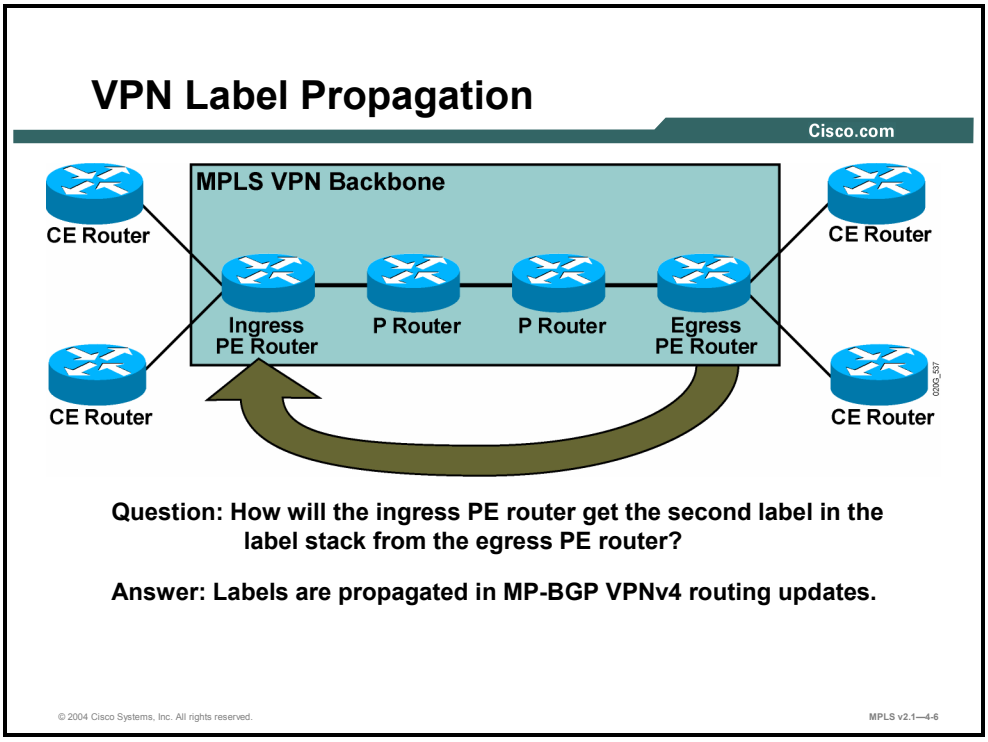

The previous figures showed that an MPLS label stack with the second label is required for proper MPLS VPN operation. This label was allocated by the egress PE router. This label has to be propagated from the egress PD router to the ingress PE routers to enable proper packet forwarding. MP-BGP was chosen as the propagation mechanism. Every MP-BGP update thus carries a label assigned by the egress PE router together with the 96-bit VPNv4 prefix.

Example: VPN Label Propagation Between PE Routers

The figure illustrates VPN label propagation between PE routers.

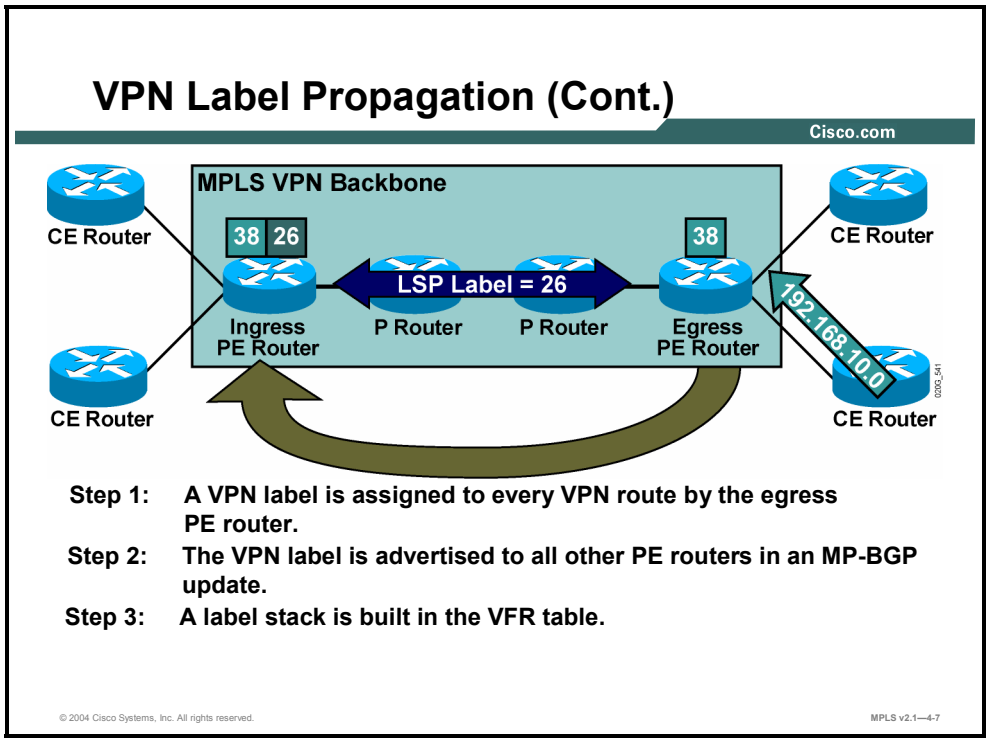

These steps describe the label propagation between PE routers.

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.

The VPN labels assigned locally by the PE router can be inspected with the **show mpls forwarding vrf xxx** command (where "xxx" is the name of the VRF).

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 tags** 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).

For example, the VPN label for destination 192.188.10.0 is 38 and was assigned by the egress PE router.

Note Like many Cisco IOS software show commands, the **show ip bgp vpnv4 all tags** command uses the old terminology labels called "tags."

Step 3 The ingress PE router has two labels associated with a remote VPN route: a label for the 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 detail** command. The *tags imposed* part of the printout 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.

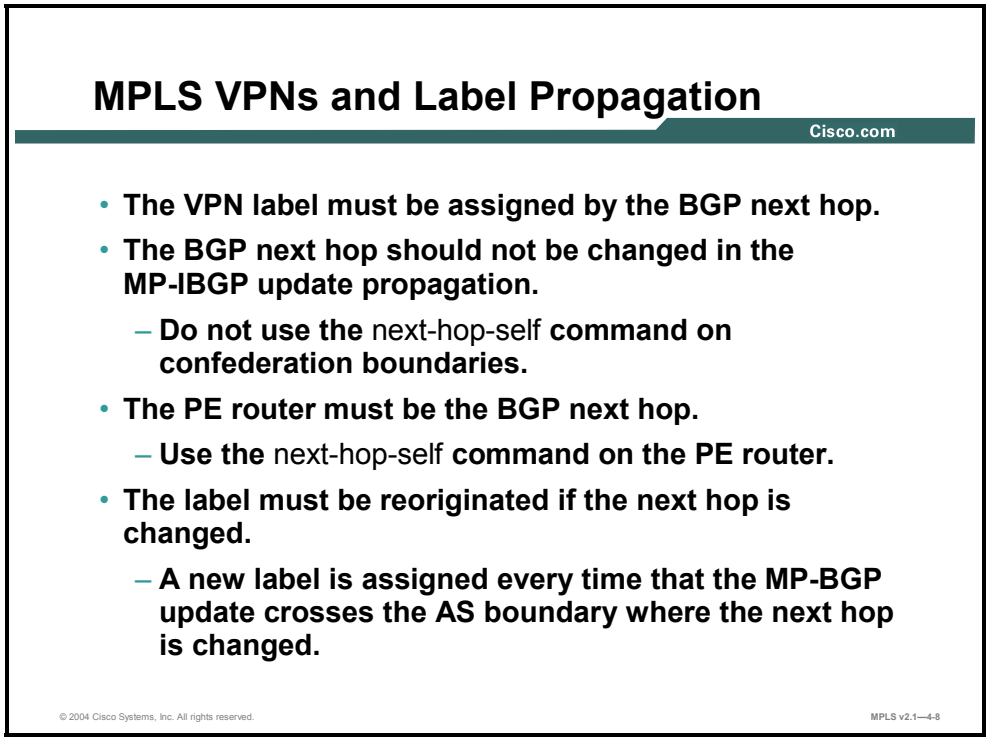

MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same router as the one that assigned the second label in the label stack. The following describes three scenarios that can cause the BGP next hop to be different from the IP address of the PE router assigning the VPN label:

- If the customer route is received from the CE router via an 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 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 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.

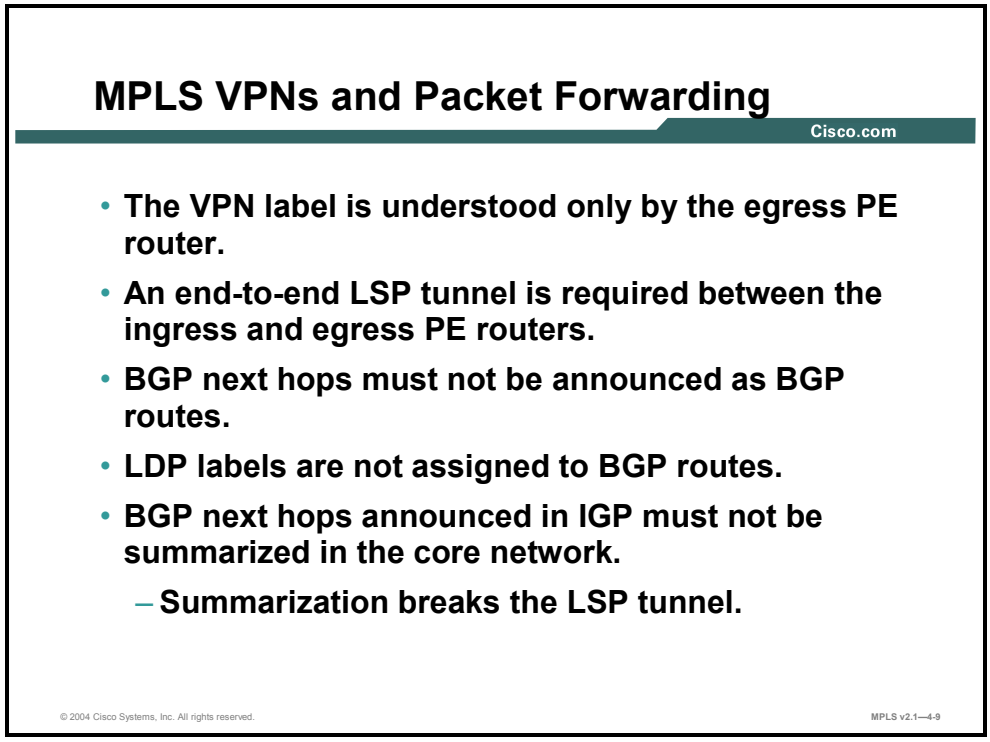

For successful propagation of MPLS VPN packets across an MPLS backbone, there must be an unbroken LSP tunnel between PE routers. This is because the second label in the stack is recognized only by the egress PE router that has originated it and will not be understood by any other router should it ever become exposed.

The following describes 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.
- 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.

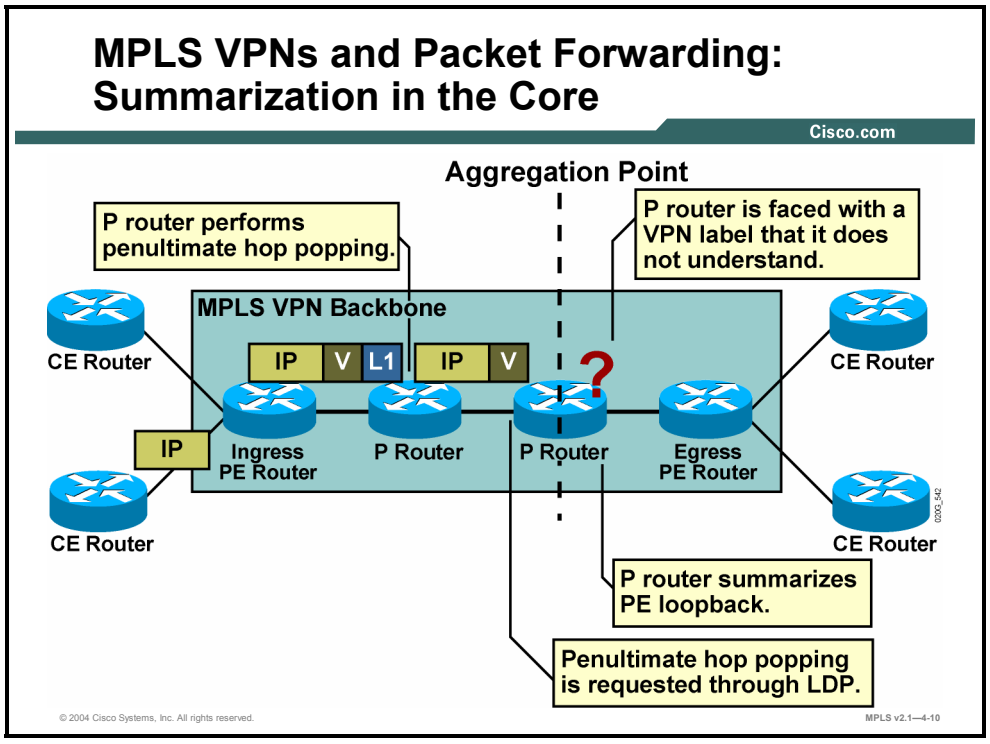

The LSP tunnel is broken at a summarization point, so the summarizing router needs to perform full IP lookup. In a frame-based MPLS network, the P router would request PHP for the summary route, and the upstream P router (or a PE router) would remove the LDP label, exposing the VPN label to the P router. 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 discussed in this lesson.

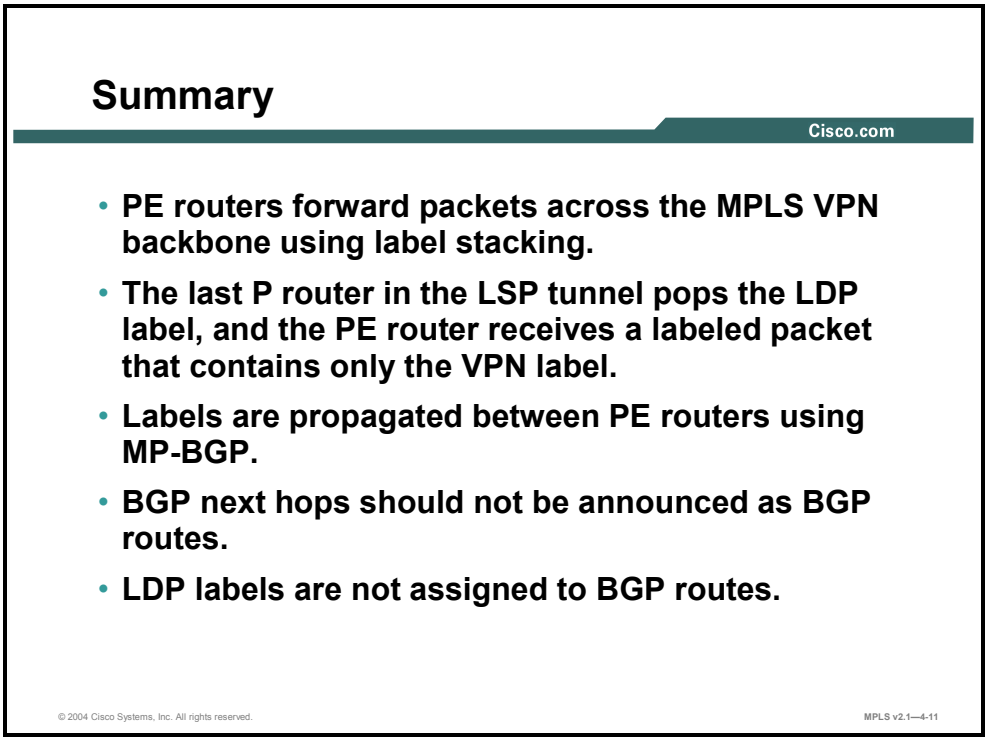

Module Summary

This topic summarizes the key points discussed in this module.

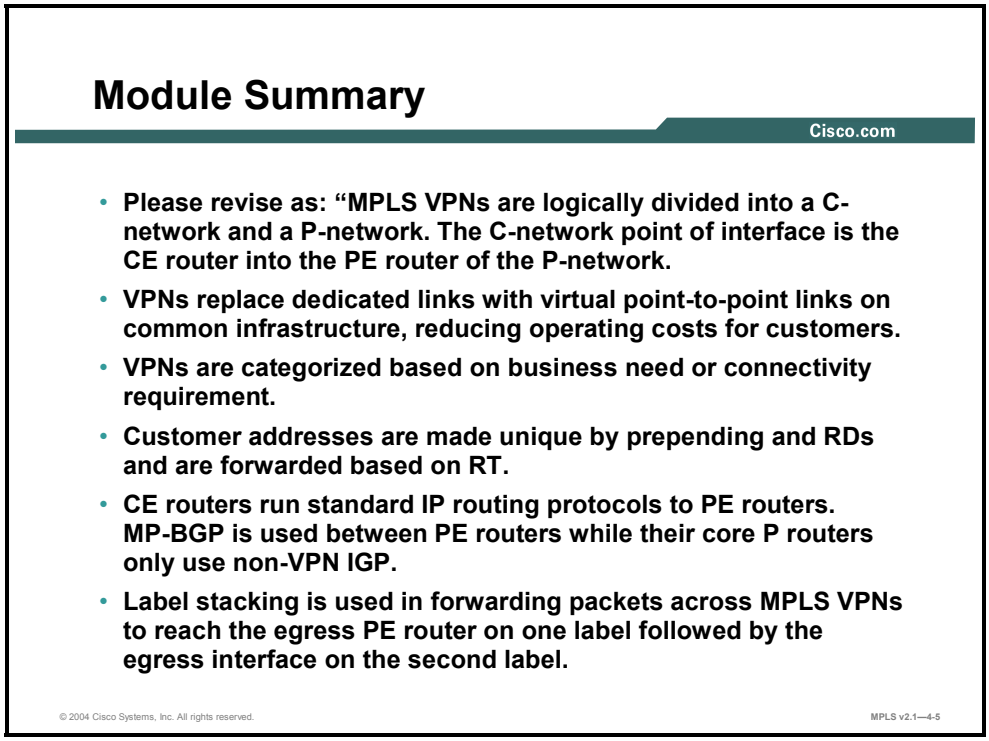

The two major VPN design options—overlay VPN and peer-to-peer VPN—have many benefits and drawbacks. The VPN topology categories and architectural components help determine the method for forwarding packets in an MPLS VPN environment.

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 were implemented using which type of links? (Source: Introducing Virtual Private Networks)

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 Virtual Private Networks)

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 Virtual Private Networks)

A) P device
B) CE device
C) PE device
D) CPE device

Q4) What are the two types of virtual circuits supported by switched WAN technologies? (Source: Introducing Virtual Private Networks)

_____

_____

Q5) Which of the following is a characteristic of an overlay VPN? (Source: Introducing Virtual Private Networks)

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.

Q6) In the traditional switched WAN model for Layer 2 VPN implementation, what are the service provider and customer responsible for? (Source: Introducing Overlay and Peer-to-Peer VPNs)

The service provider is responsible for _____.

The customer is responsible for _____.

The peer-to-peer VPN concept was introduced to help overcome what type of drawback?

_____

_____

_____

Q7) How is a peer-to-peer VPN implemented using packet filters? (Source: Introducing Overlay and Peer-to-Peer VPNs)

_____

_____

_____

Q8) How do you implement a peer-to-peer VPN based on controlled route distribution? (Source: Introducing Overlay and Peer-to-Peer VPNs)

_____

_____

_____

Q9) Which VPN type does NOT require the service provider to participate in customer routing? (Source: Introducing Overlay and Peer-to-Peer VPNs)

A) overlay
B) peer-to-peer

Q10) For which VPN type is it easier to provision an additional VPN? (Source: Introducing Overlay and Peer-to-Peer VPNs)

A) overlay
B) peer-to-peer

Q11) Which VPN type requires the PE router to carry all routes from all customers? (Source: Introducing Overlay and Peer-to-Peer VPNs)

A) overlay
B) peer-to-peer

Q12) Which VPN type requires the service provider to participate in customer routing? (Source: Introducing Overlay and Peer-to-Peer VPNs)

A) overlay
B) peer-to-peer

Q13) Describe the use of address space and packet routing in each of the following peer-to-peer implementations. (Source: Introducing Overlay and Peer-to-Peer 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 Overlay and Peer-to-Peer 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 Overlay and Peer-to-Peer 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 Overlay and Peer-to-Peer 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) Name the VPN topology that has many sites connecting to a central site. (Source: Categorizing VPNs)

Q18) When you are using a dynamic routing protocol such as RIP in a redundant hub-and-spoke topology, which of the following is true? (Source: Categorizing VPNs)

A) Static routing must be used to provide remote-site-to-remote-site connectivity.
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

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 VPN 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 of the following is true if you use the P-network IPG 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 nonunique IP address into a unique VPNv4 address

Q36) Which routing protocol does the CE router run? (Source: Introducing 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 routers exchange VPNv4 routes? (Source: Introducing 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 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 MPLS VPN Routing Model)

Q40) What BGP function do MPLS VPNs use to transport RTs? (Source: Introducing MPLS VPN Routing Model)

Q41) How does the PE router know in which VRF table to install received routes for a customer? (Source: Introducing MPLS VPN Routing Model)

Q42) What is the impact of an MPLS VPN on CE routers? (Source: Introducing 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 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 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 MPLS VPN Routing Model)

A) The route will not insert 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 P router at the point of summarization 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) switched virtual circuits, permanent virtual circuits

Q5) D

Q6) providing end-to-end connectivity, routing updates
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) hub-and-spoke

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 nonunique 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