MPLS

Implementing Cisco MPLS

Version 2.0

Student Guide

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 • 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

Printed in the USA

Table of Contents

Volume 1

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 are realizing that something else is needed. Not only must they be concerned with protecting their existing infrastructure, but they must also find ways to generate new services that are not currently supportable using existing technologies.

Multiprotocol Label Switching (MPLS) is a high-performance method for forwarding packets through a network. It 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, it allows service providers the ability to support value-added features critical for their networks.

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

Outline

The Course Introduction includes these topics:

- Course Objectives
- Cisco Certifications
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Flow Diagram
- Icons and Symbols
- Learner Introductions
- Course Evaluations

Course Objectives

This topic lists the course objectives.

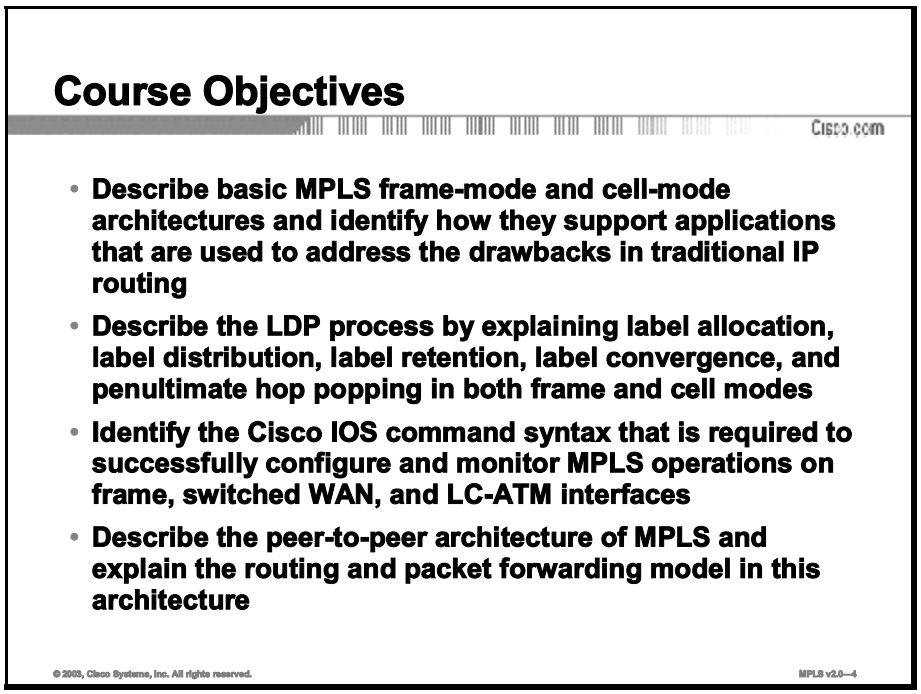

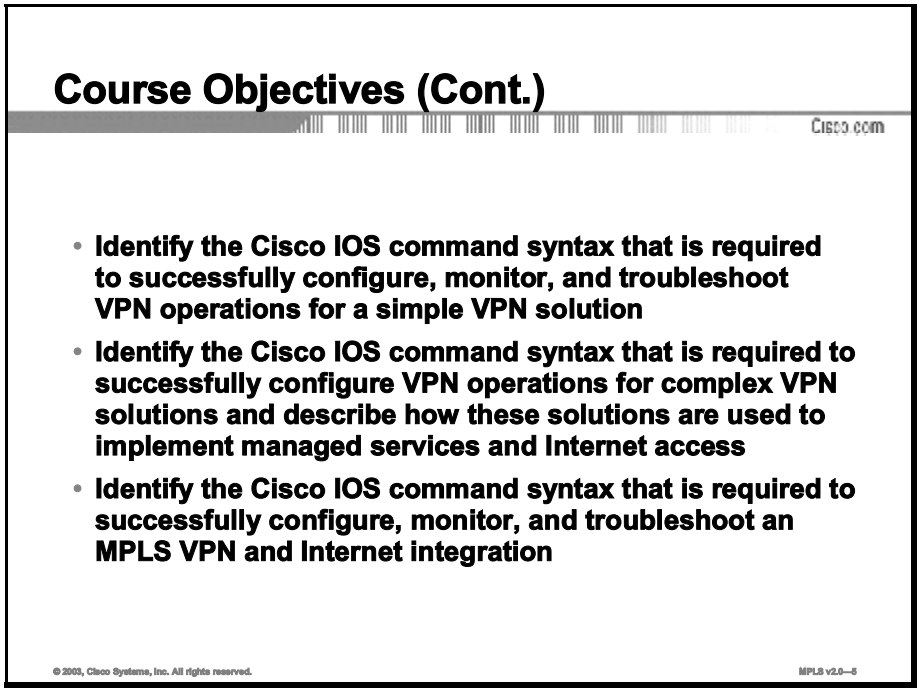

Cisco Certifications

This topic lists the certification requirements of this course.

Cisco provides three levels of general career certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused Cisco Qualified Specialist (CQS) certifications for designated areas such as cable communications, voice, and security.

There are many paths to Cisco certification, but only one requirement—passing one or more exams demonstrating knowledge and skill. For details, go to http://www.cisco.com/go/certifications.

Learner Skills and Knowledge

This topic lists the course prerequisites.

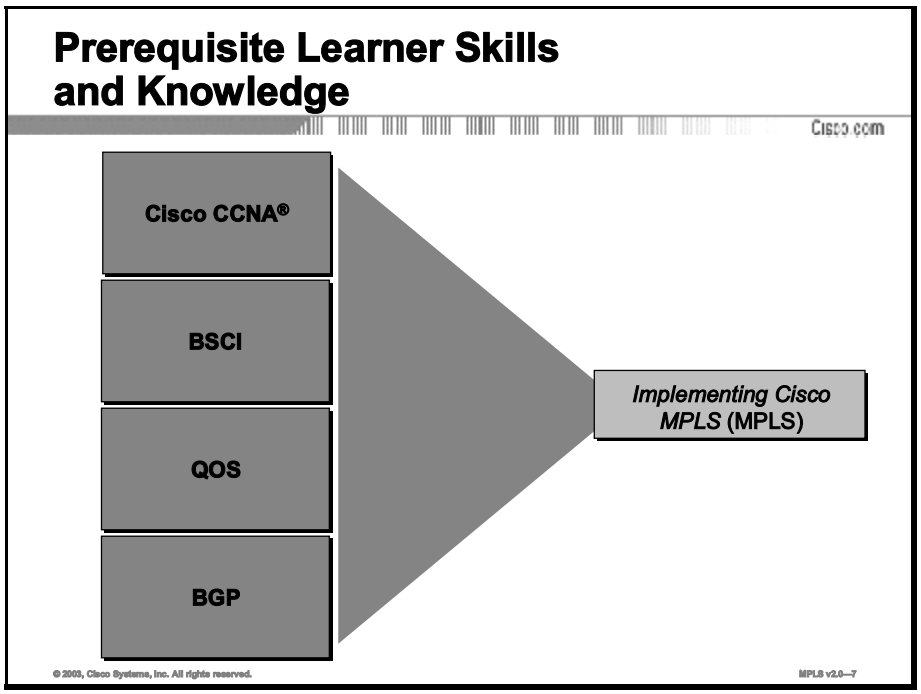

To benefit fully from this course, you must have these prerequisite skills and knowledge:

- Cisco CCNA® certification
- *Building Scalable Cisco Internetworks* (BSCI)
- *Implementing Cisco Quality of Service* (QOS)
- *Configuring BGP on Cisco Routers* (BGP)

Note	Practical experience with deploying and operating networks based on Cisco network devices and Cisco IOS software is strongly recommended.

Learner Responsibilities

This topic discusses the responsibilities of the learners.

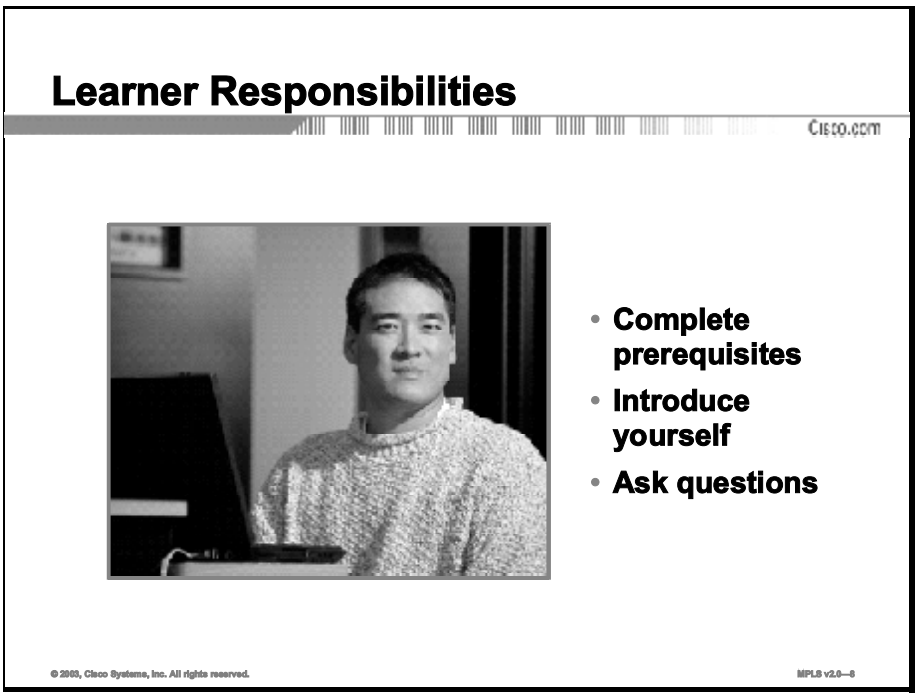

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This topic lists the administrative issues for the course.

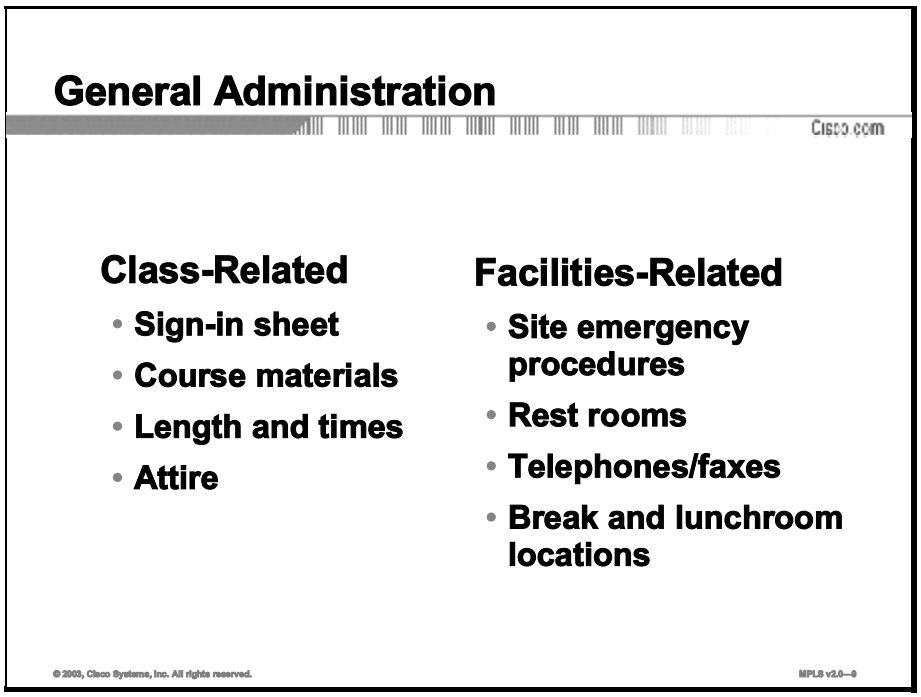

The instructor will discuss the administrative issues noted here so you know exactly what to expect from the class.

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Flow Diagram

This topic covers the suggested flow of the course materials.

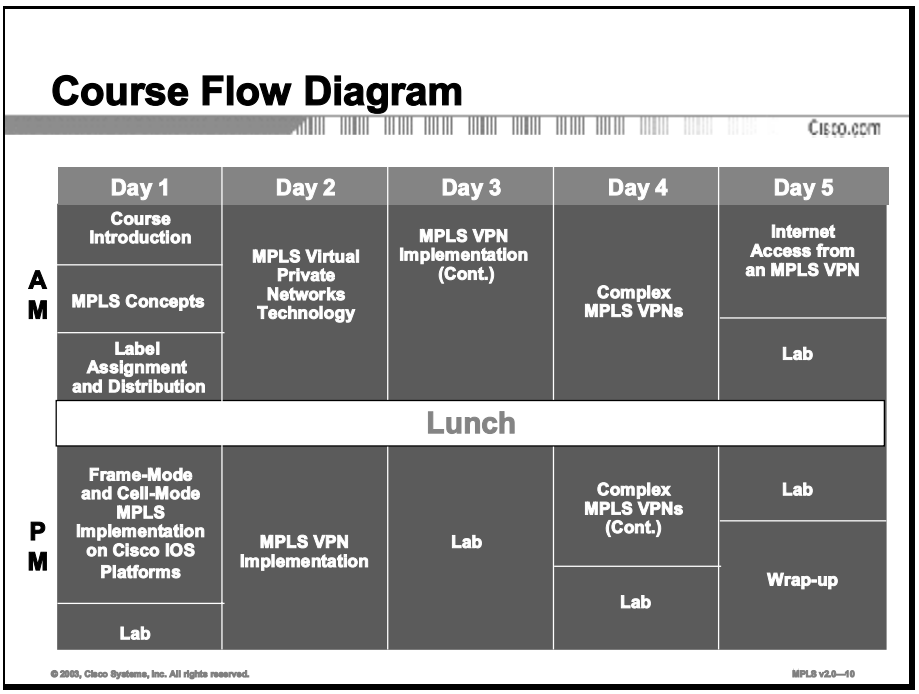

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

This topic shows the Cisco icons and symbols used in this course.

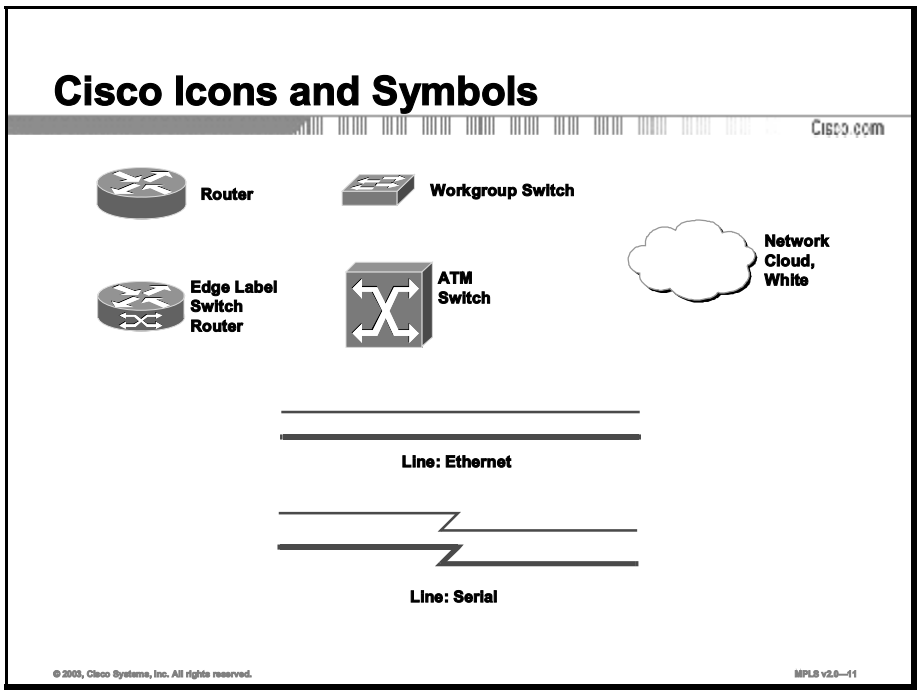

Learner Introductions

This is the point in the course where you introduce yourself.

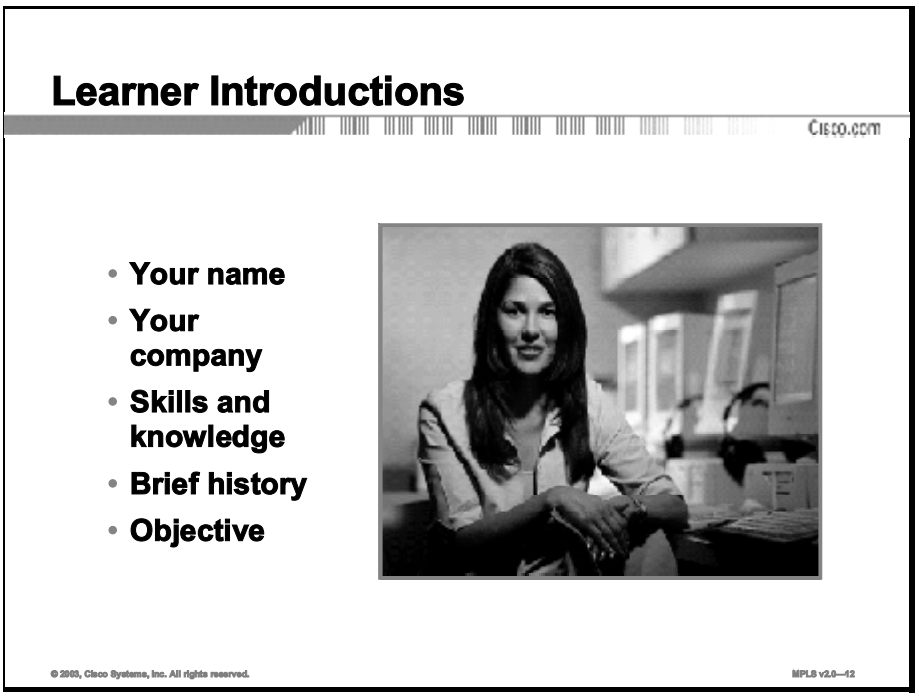

Prepare to share the following information:

- Your name
- Your company
- If you have most or all of the prerequisite skills
- A profile of your experience
- What you would like to learn from this course

Course Evaluations

Cisco relies on customer feedback to make improvements and guide business decisions. Your valuable input will help shape future Cisco learning products and program offerings.

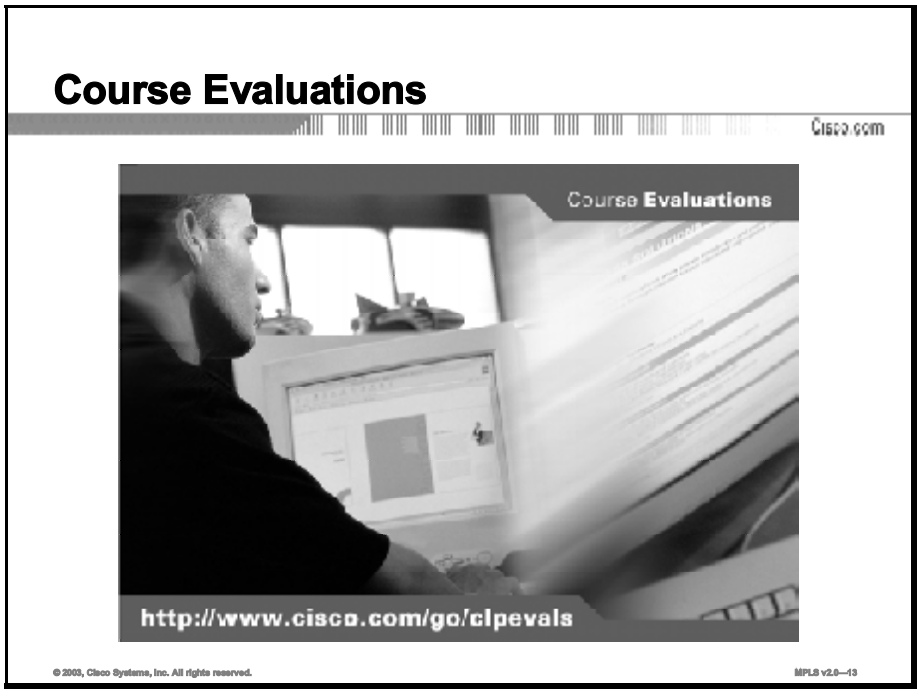

On the first and final days of class, your instructor will provide the following information needed to fill out the evaluation:

- Course acronym *(printed on student kit side label)* _____
- Course version number *(printed on student kit side label)* _____
- Cisco Learning Partner ID # _____
- Instructor ID # _____
- Course ID # *(for courses registered in Cisco Learning Locator)* _____

Please use this information to complete a brief (approximately 10 minutes) online evaluation concerning your instructor and the course materials in the student kit. To access the evaluation, go to http://www.cisco.com/go/clpevals.

After the completed survey has been submitted, you will be able to access links to a variety of Cisco resources, including information on the Cisco Career Certification programs and future Cisco Networkers events.

If you encounter any difficulties accessing the course evaluation URL or submitting your evaluation, please contact Cisco via email at clpevals_support@external.cisco.com.

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, as well as MPLS label format and label switch router (LSR) architecture and operations, are explained.

Module Objectives

Upon completing this module, you will be able to describe the basics of MPLS. This includes being able to do the following:

- Describe the basic MPLS concepts, including the drawbacks in traditional IP routing
- Describe MPLS labels and MPLS label stacks, including the format of the MPLS label and also when and why a label stack is created
- Describe the different MPLS applications, including MPLS VPNs and MPLS TE

Module Outline

The module contains these lessons:

- Basic MPLS Concepts
- MPLS Labels and Label Stack
- MPLS Applications

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. It lays the foundation for subsequent lessons that cover the key areas, such as Cisco MPLS Traffic Engineering (MPLS TE) and Virtual Private Networks (VPNs).

Relevance

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, as well as understand how information gets transferred across an MPLS domain.

Objectives

This lesson describes the basic MPLS concepts, including the drawbacks in traditional IP routing. Upon completing this lesson, you will be able to:

- Discuss the drawbacks of traditional IP routing
- Describe the basic MPLS concepts
- Describe MPLS versus IP over ATM
- List the advantages of using MPLS TE
- Describe the architecture of MPLS
- Describe the MPLS label format
- Describe the different types of label switch routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of IP routing
- Basic understanding of the various routing protocols, including RIP, IGRP, EIGRP, OSPF, IS-IS, and BGP

Outline

This lesson includes these topics:

- Overview
- Drawbacks of Traditional IP Routing
- Basic MPLS Concepts
- MPLS Versus IP over ATM
- Traffic Engineering with MPLS
- MPLS Architecture
- MPLS Labels
- Label Switch Routers
- Summary
- Quiz

Drawbacks of Traditional IP Routing

This topic details the drawbacks of traditional IP routing.

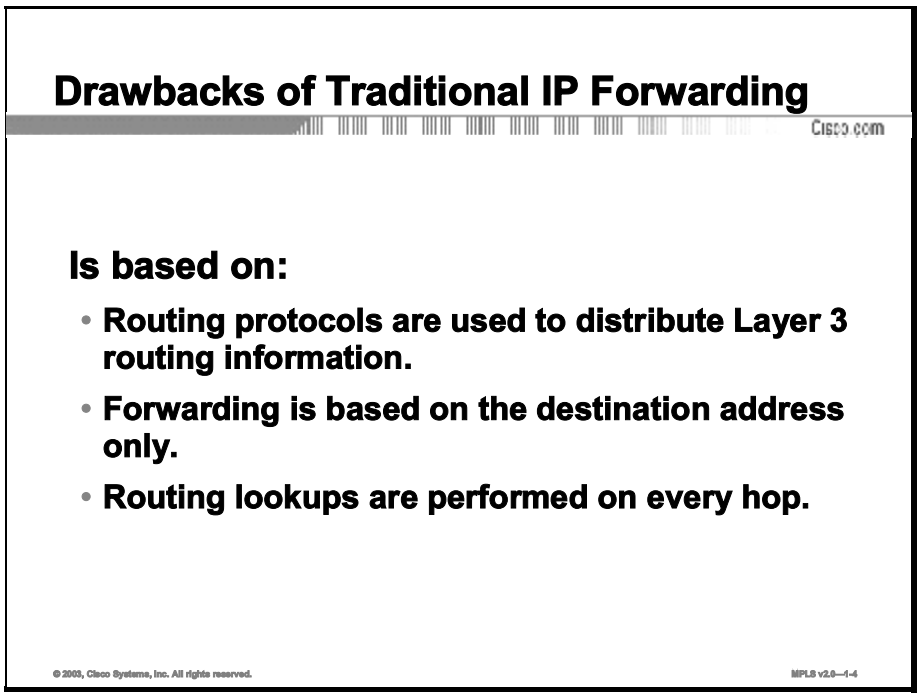

Before basic MPLS functionality is explained, three drawbacks of traditional IP forwarding need to be highlighted:

- 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.

Drawbacks of Traditional IP Forwarding (Cont.)
Traditional IP Forwarding

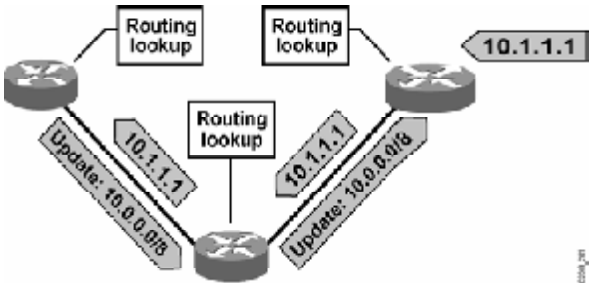

- **Every router may need full Internet routing information (more than 100,000 routes).**
- **Destination-based routing lookup is needed on every hop.**

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, or 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.

Drawbacks of Traditional IP Forwarding (Cont.)
IP over ATM

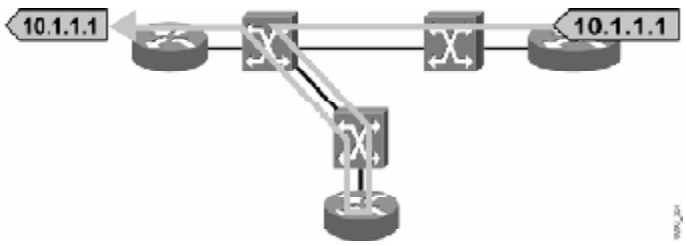

- **Layer 2 devices have no knowledge of Layer 3 routing information—virtual circuits must be manually established.**
- **Layer 2 topology may be different from Layer 3 topology, resulting in suboptimal paths and link use.**
- **Even if the two topologies overlap, the hub-and-spoke topology is usually used because of easier management.**

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 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.

Drawbacks of Traditional IP Forwarding (Cont.) Traffic Engineering

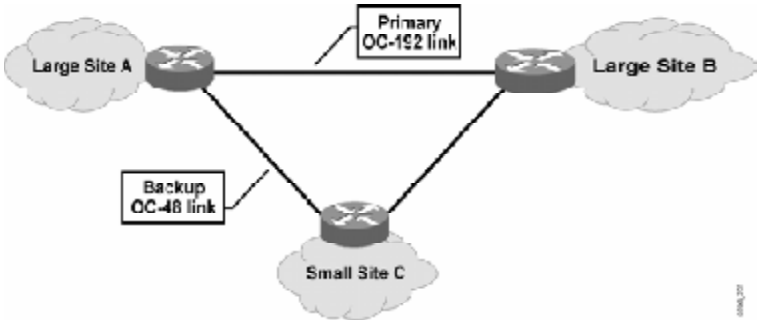

- **Most traffic goes between large sites A and B, and uses only the primary link.**
- **Destination-based routing does not provide any mechanism for load balancing across unequal paths.**
- **Policy-based routing can be used to forward packets based on other parameters, but this is not a scalable solution.**

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.

Basic MPLS Concepts

This topic details the basic concepts of MPLS.

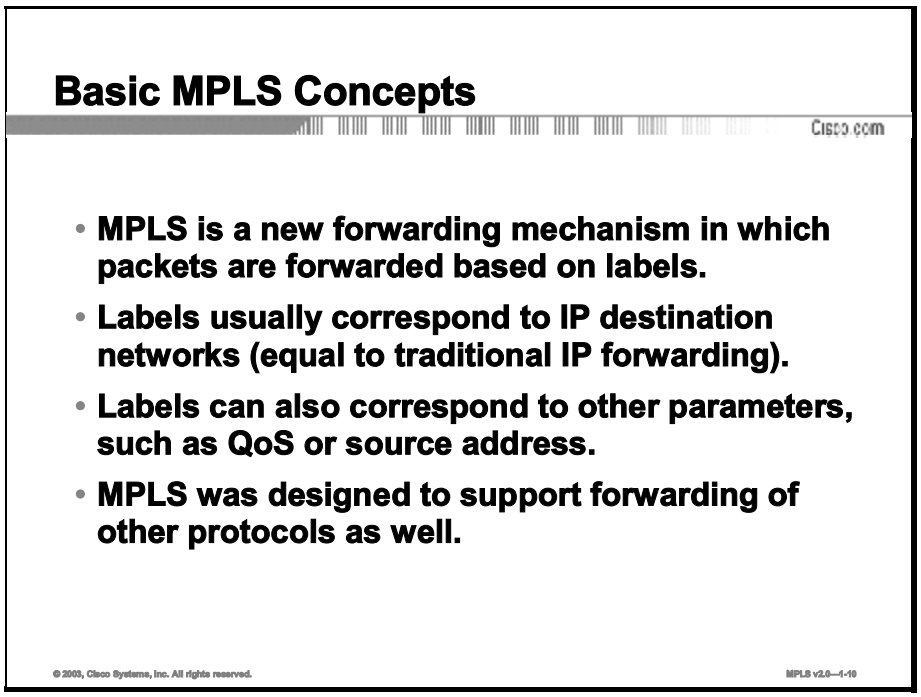

MPLS is a new switching mechanism that uses labels (numbers) to forward packets.

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.

Basic MPLS Concepts (Cont.) Example

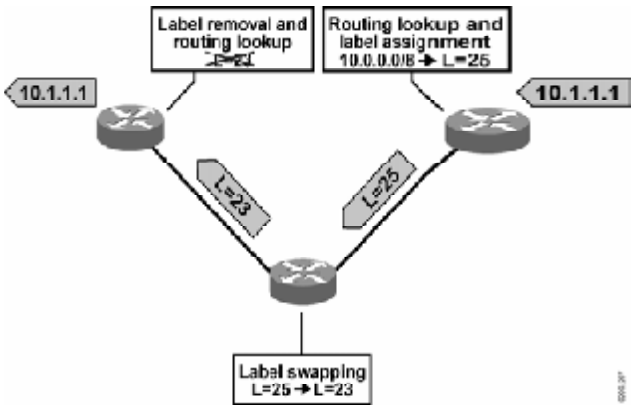

- **Only edge routers must perform a routing lookup.**
- **Core routers switch packets based on simple label lookups and swap labels.**

The figure illustrates a situation where 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.

MPLS versus IP over ATM

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

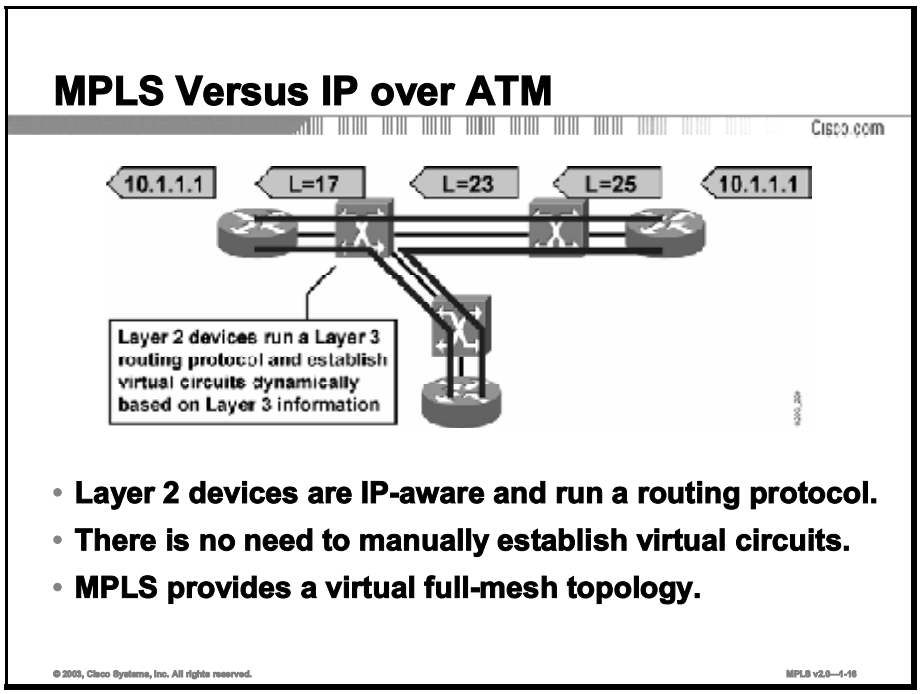

The figure shows how MPLS is used in ATM networks to provide optimal routing across Layer 2 ATM switches. In order for MPLS to work with ATM switches, the switches must be Layer 3–aware. In other words ATM switches must run a Layer 3 routing protocol.

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.

Traffic Engineering with MPLS

This topic introduces MPLS and traffic engineering interaction.

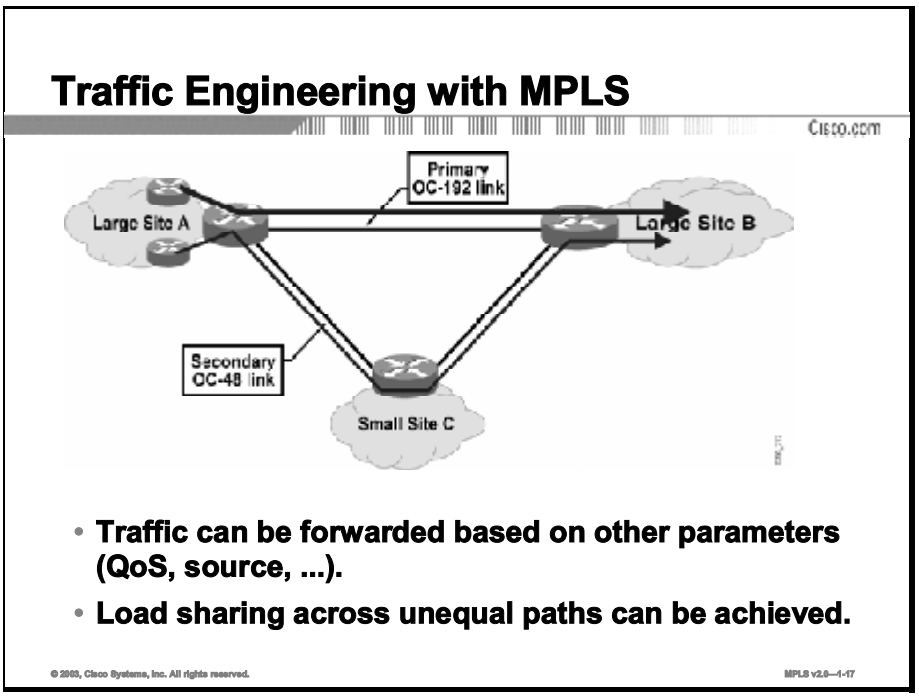

MPLS also supports traffic engineering. Traffic-engineered tunnels can be created based on traffic analysis to provide load balancing across unequal paths.

Multiple traffic engineering 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. Traffic engineering determines the path at the source based on additional parameters such as available resources and constraints in the network.

MPLS Architecture

This topic looks at the main components of the MPLS architecture.

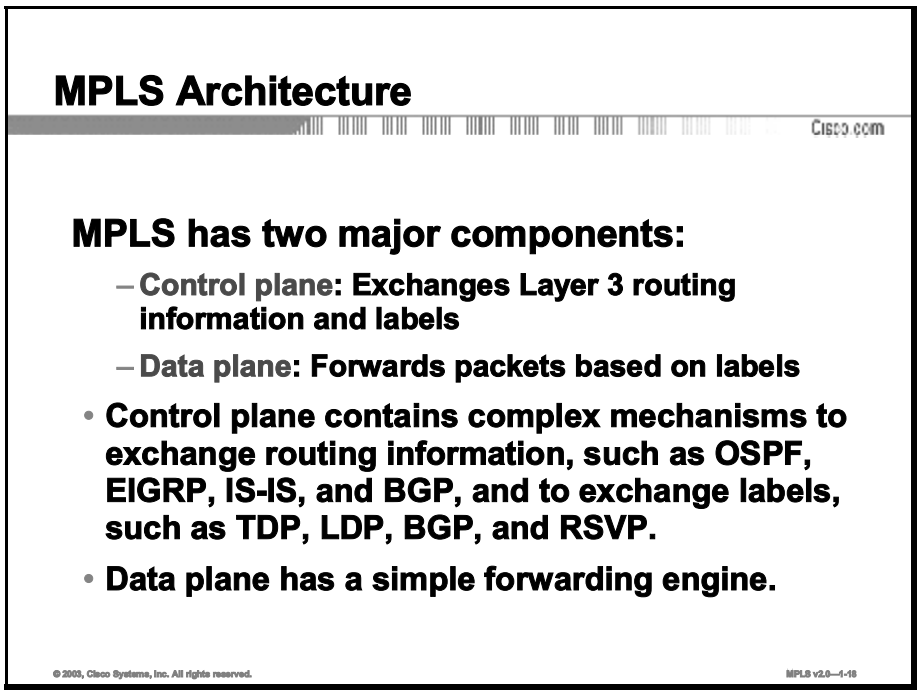

MPLS consists of two major components:

- **Control plane:** Takes care of the routing information exchange and the label exchange between adjacent devices
- **Data plane:** Takes care of forwarding based on either destination addresses or labels; 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 Architecture (Cont.)

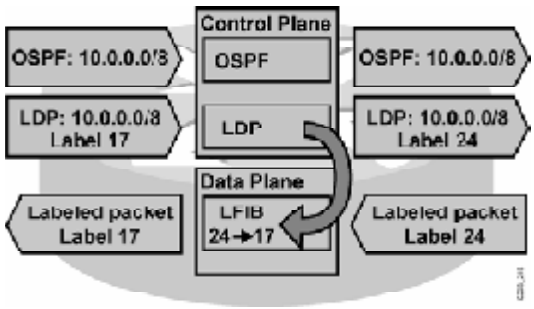

- **Router functionality is divided into two major parts: control plane and data plane**

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.

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 a destination address 10.x.x.x. A local label 24 is generated and sent to upstream neighbors so 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.

MPLS Labels

This topic introduces MPLS labels and their format.

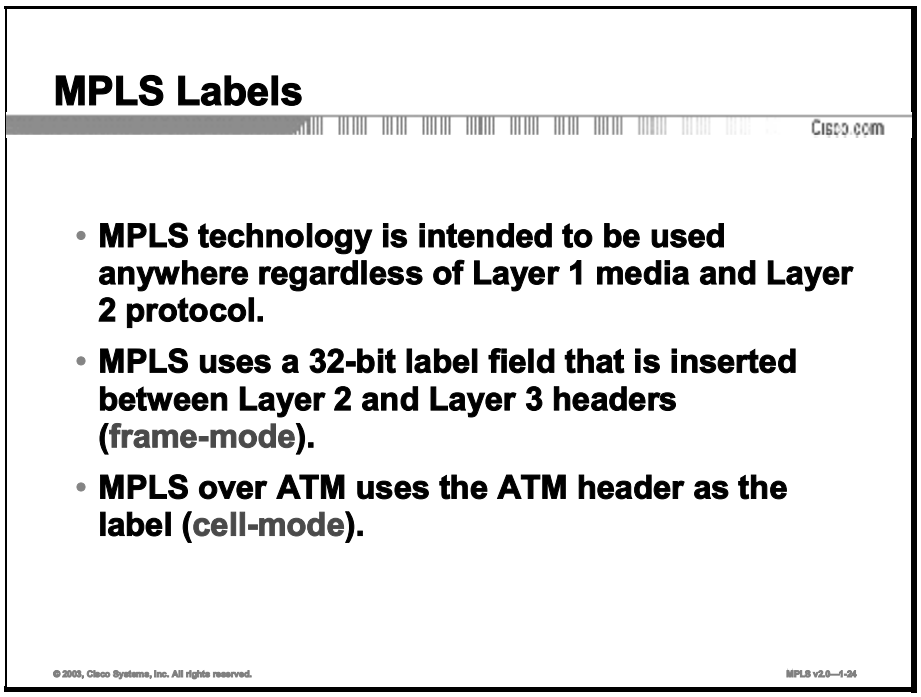

MPLS is designed for use on virtually any media and Layer 2 encapsulation. Most Layer 2 encapsulations are frame-based, and MPLS simply inserts a 32-bit label between the Layer 2 and Layer 3 headers (“frame-mode” MPLS).

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).

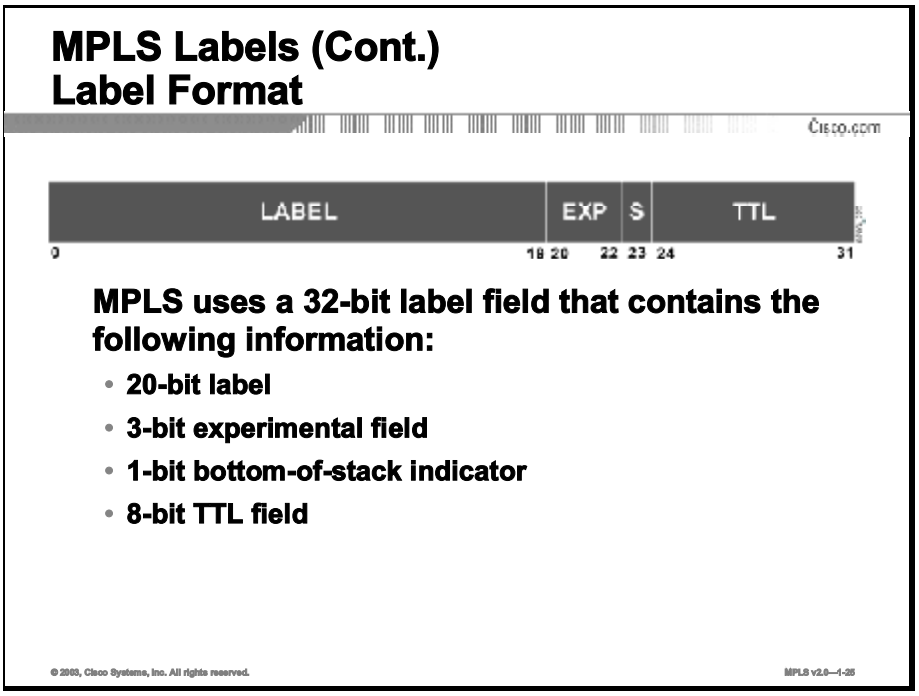

A 32-bit label contains these fields:

Label Fields

Field	Description
`20-bit label`	The actual label.
`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 time-to-live (TTL) field`	Has the same purpose as the TTL field in the IP header.

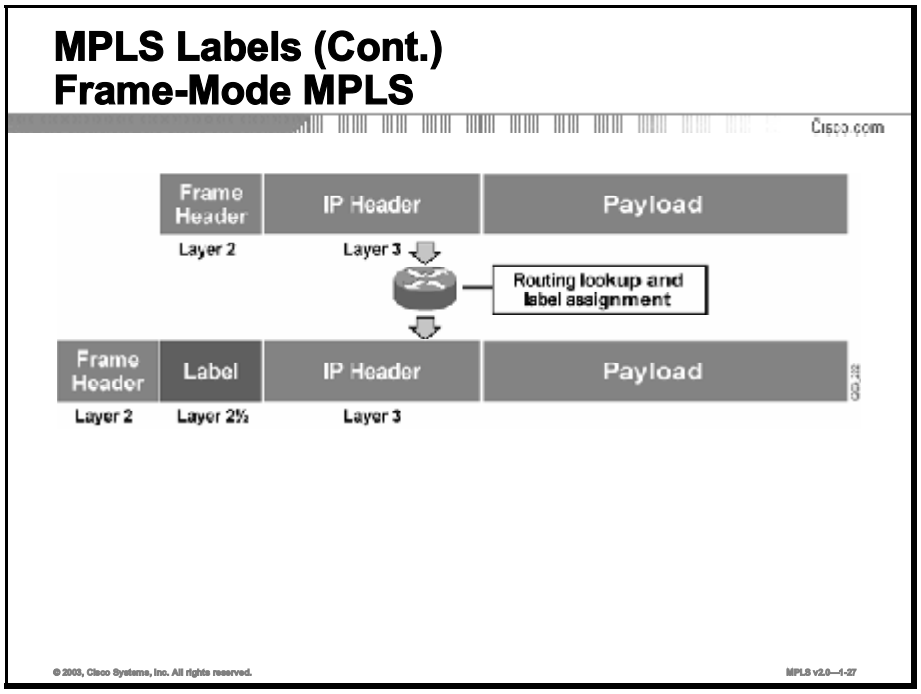

The figure shows an edge router that receives a normal IP packet. The router then does the following:

- Performs routing lookup to determine the outgoing interface
- 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
- Sends the labeled packet

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

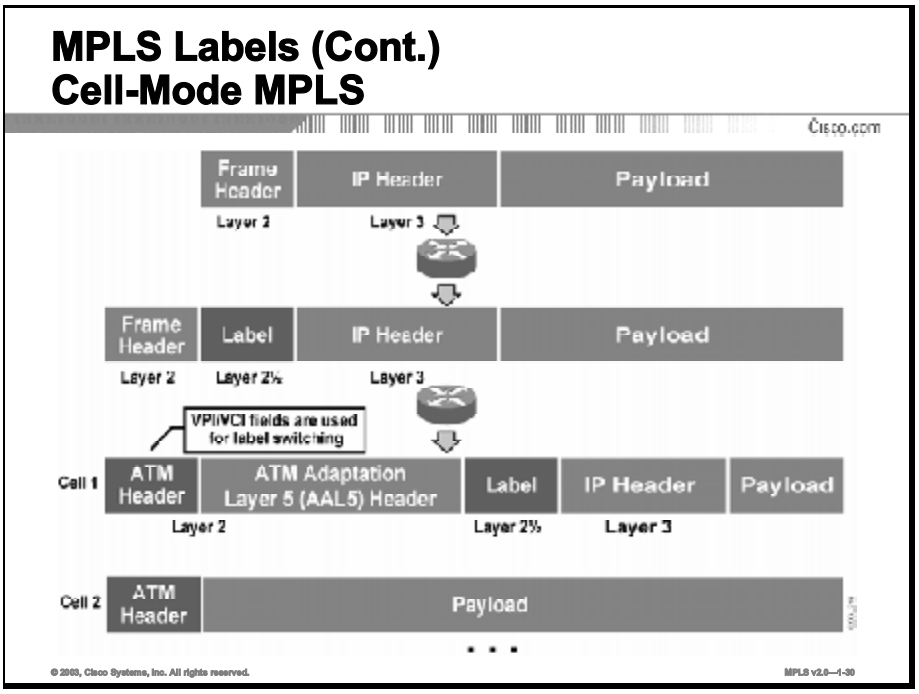

Cell-mode MPLS uses the ATM header VPI/VCI field for forwarding decisions. The 32-bit label is preserved in the frame but not used in the ATM network. The original label is present only in the first cell of a packet.

Label Switch Routers

This topic lists and describes label switch routers (LSRs).

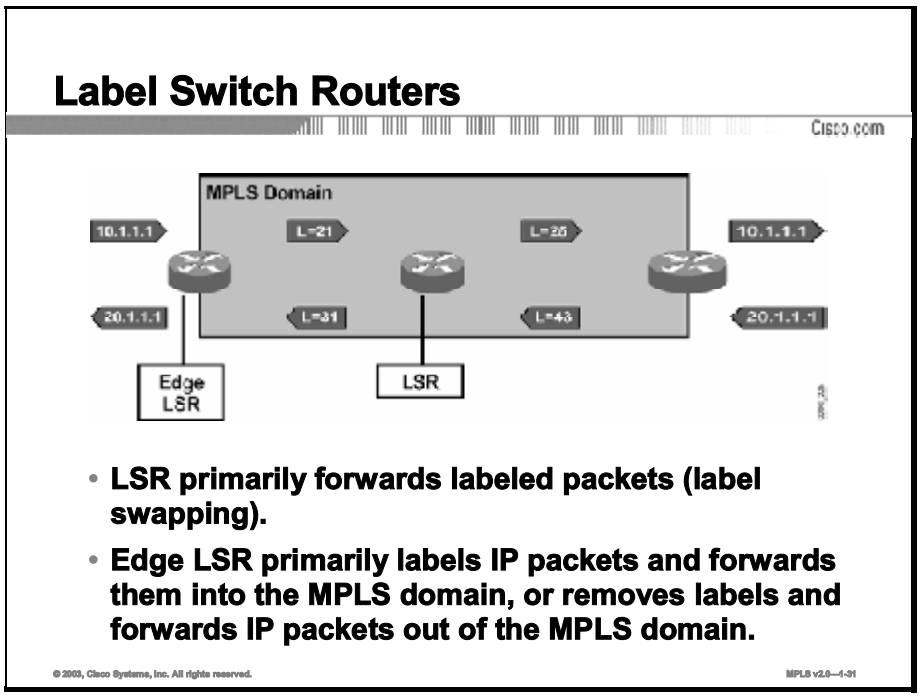

In preparation for a detailed description of MPLS, here is some of the terminology used in this course:

- **Label switch router (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, or AS). These routers also forward packets based on IP destination addresses and label them if the outgoing interface is enabled for MPLS.

Label Switch Routers (Cont.)
ATM Label Switch Router

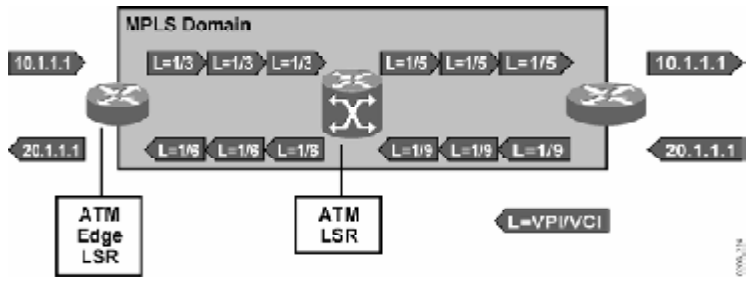

- **ATM LSR can forward only cells.**
- **ATM edge LSR segments packets into cells and forwards them into an MPLS ATM domain, or reassembles cells into packets and forwards them out of an MPLS ATM domain.**

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.

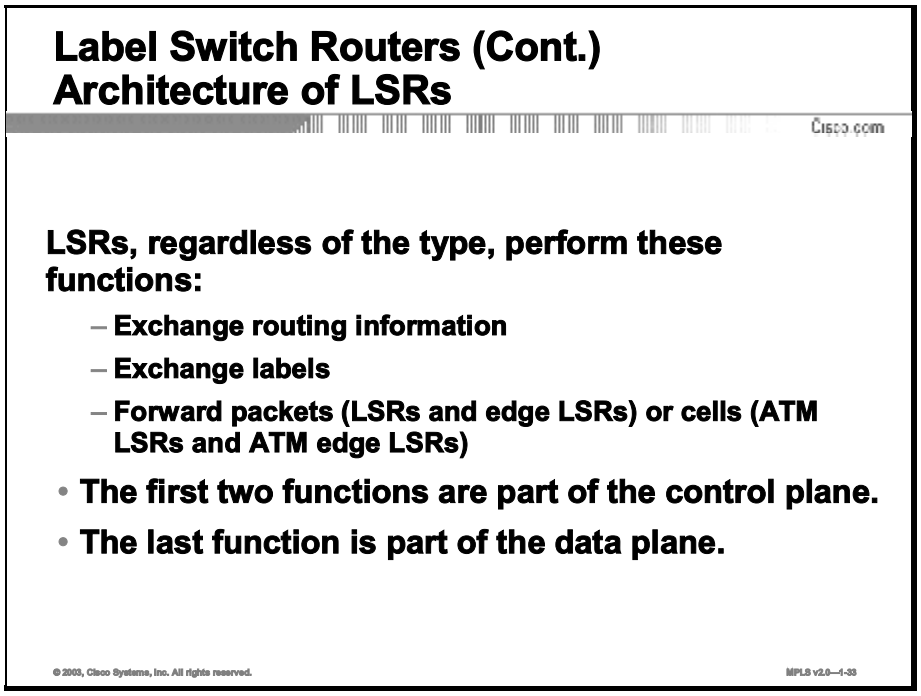

LSRs of all types must perform these functions:

- 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)

Label Switch Routers (Cont.)
Architecture of ATM LSRs (Cont.)

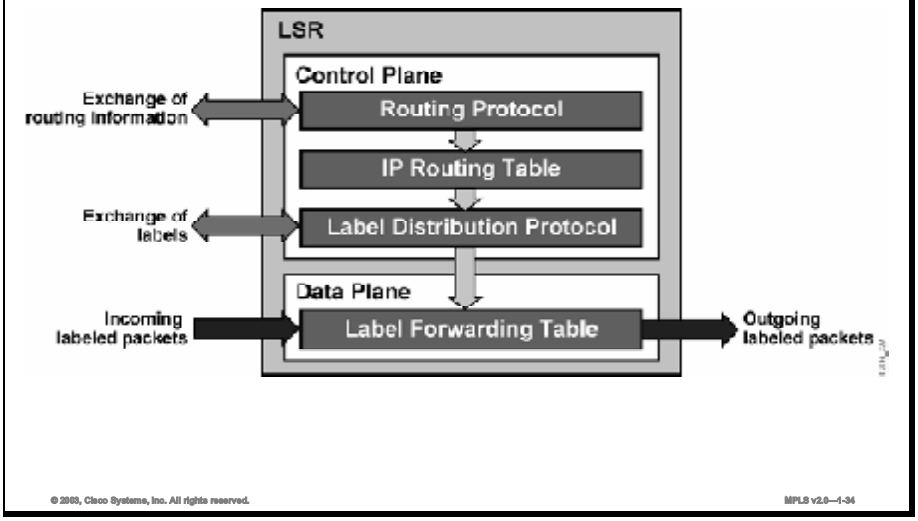

The primary function of an LSR is to forward labeled packets. Therefore, every LSR needs a Layer 3 routing protocol (for example, OSPF, EIGRP, IS-IS) and a label distribution protocol (for example, LDP, TDP).

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 the routing information.

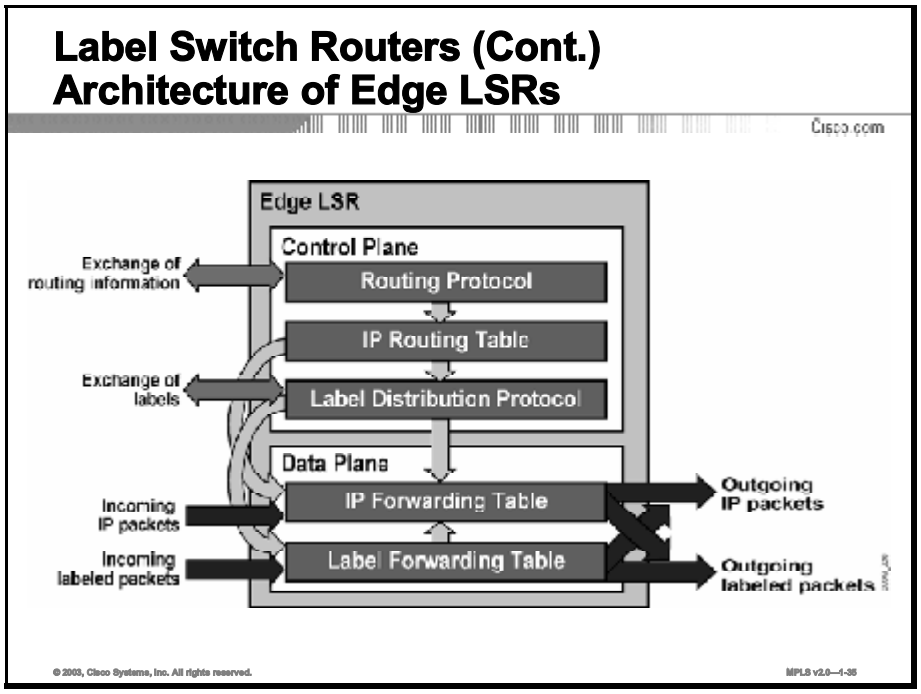

Edge LSRs also forward IP packets based on their IP destination addresses and optionally label them if a label exists.

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, or 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.

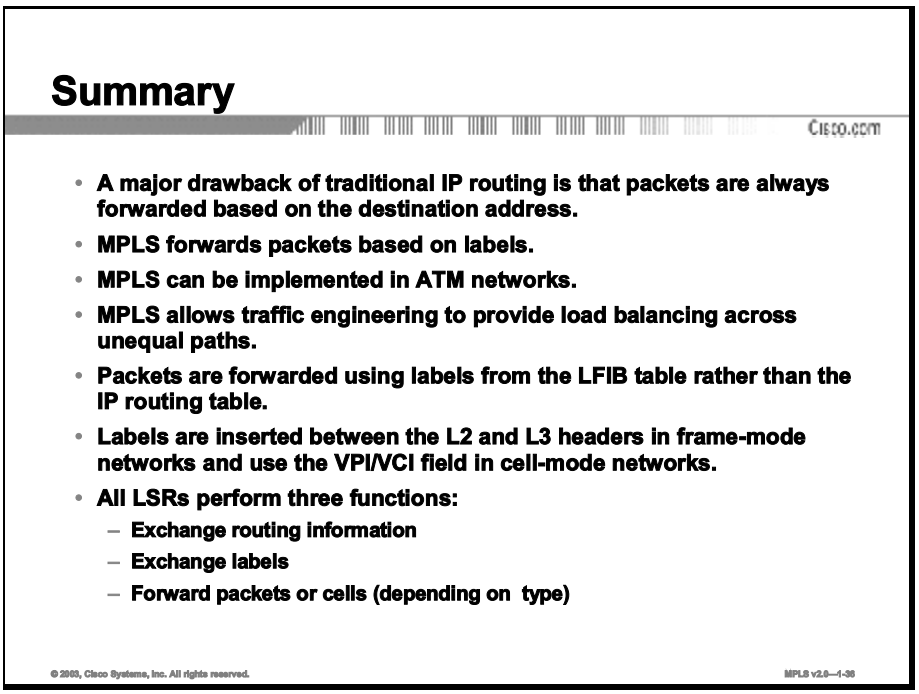

References

For additional information, refer to this resource:

- RFC 3031, “Multiprotocol Label Switching Architecture”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are three drawbacks of traditional IP routing? (Choose three.)

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 of the following is not true?

A) MPLS uses labels to forward packets.

B) MPLS works only in IP networks.

C) MPLS labels can correspond to L3 destination address, QoS, source address, or L2 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?

A) Layer 2 devices run a Layer 3 routing protocol

B) VCs (Virtual Circuits) still require that they be established manually

C) MPLS can not run in an ATM network

D) ATM switches needed to be made L3 AND L4 aware

Q4) In MPLS TE, which two of the following statements are true? (Choose two.)

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:

A) data plane

B) forwarding plane

C) system plane

D) control plane

Q6) The MPLS label field consists of how many bits?

A) 64 bits

B) 32 bits

C) 16 bits

D) 8 bits

Q7) Which two of the following are true? (Choose two.)

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.

Quiz Answer Key

Q1) A, B, C

Relates to: Drawbacks of Traditional IP Routing

Q2) B

Relates to: Basic MPLS Concepts

Q3) A

Relates to: MPLS versus IP over ATM

Q4) B, C

Relates to: Traffic Engineering with MPLS

Q5) D

Relates to: MPLS Architecture

Q6) B

Relates to: MPLS Labels

Q7) A, C

Relates to: Label Switch Routers

MPLS Labels and Label Stack

Overview

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

Relevance

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 MPLS label stacks, including the format of the MPLS label and also when and why a label stack is created.

Upon completing this lesson, you will be able to:

- Describe where MPLS labels are inserted in an IP packet
- List the format and fields of an MPLS label
- Describe an MPLS label stack and when a label stack is created
- Describe the steps that MPLS uses in forwarding packets

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Basic MPLS Concepts” lesson of this module

Outline

This lesson includes these topics:

- Overview
- MPLS Labels
- MPLS Label Format
- MPLS Label Stack
- MPLS Forwarding
- Summary
- Quiz

MPLS Labels

This topic provides an overview of MPLS labels.

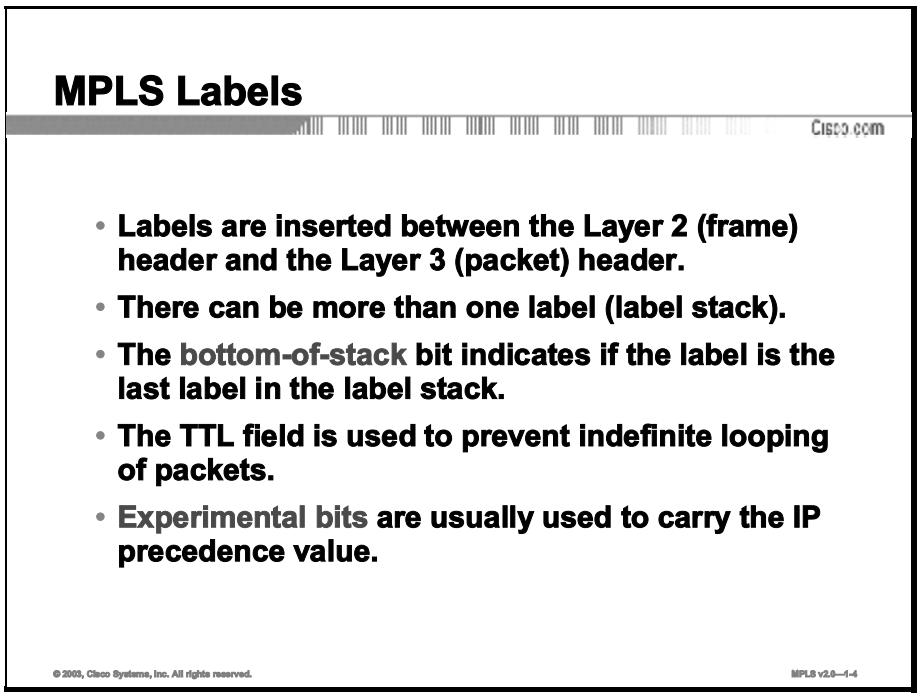

MPLS uses a 32-bit label that is inserted between the Layer 2 and Layer 3 headers. An MPLS label contains four fields:

- The actual label
- Experimental field
- Bottom-of-stack bit
- Time-to-live (TTL) field

These fields are explained in detail on the following pages.

MPLS Label Format

This topic describes the MPLS label format.

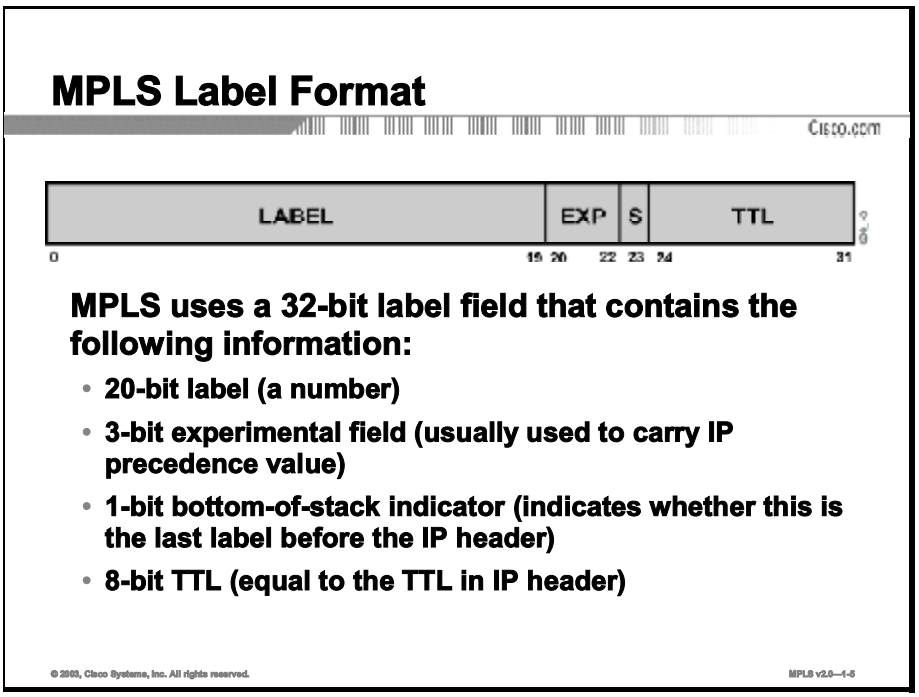

A label contains these fields:

Label Fields

Field	Description
`20-bit label`	The actual label.
`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 time-to-live (TTL) field`	Has the same purpose as the TTL field in the IP header.

MPLS Label Stack

This topic describes the MPLS label stack.

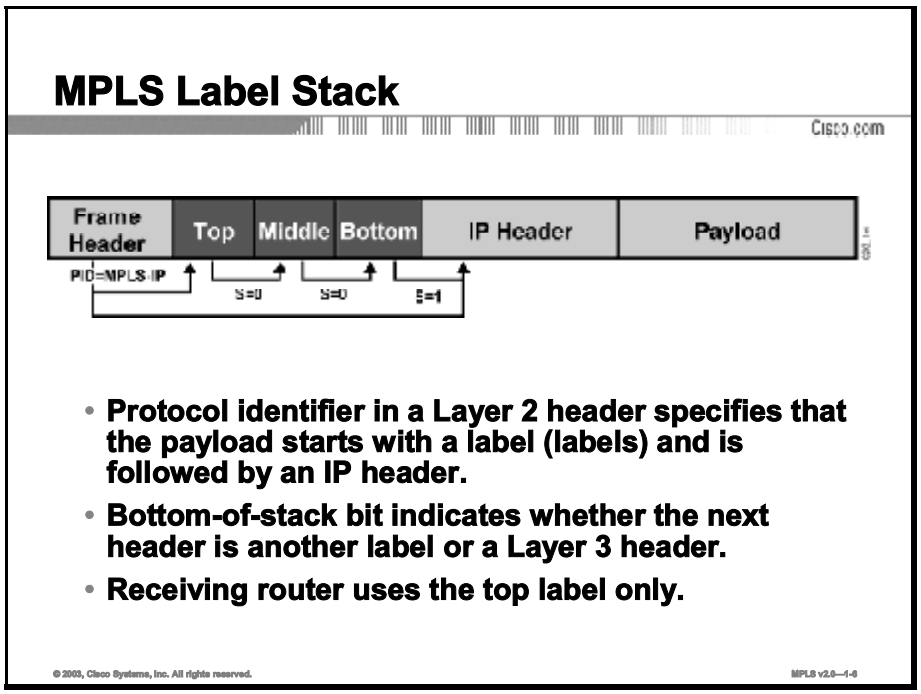

A label does not contain any information about the Layer 3 protocol being carried in a packet. A new protocol identifier is used for every MPLS-enabled Layer 3 protocol.

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.

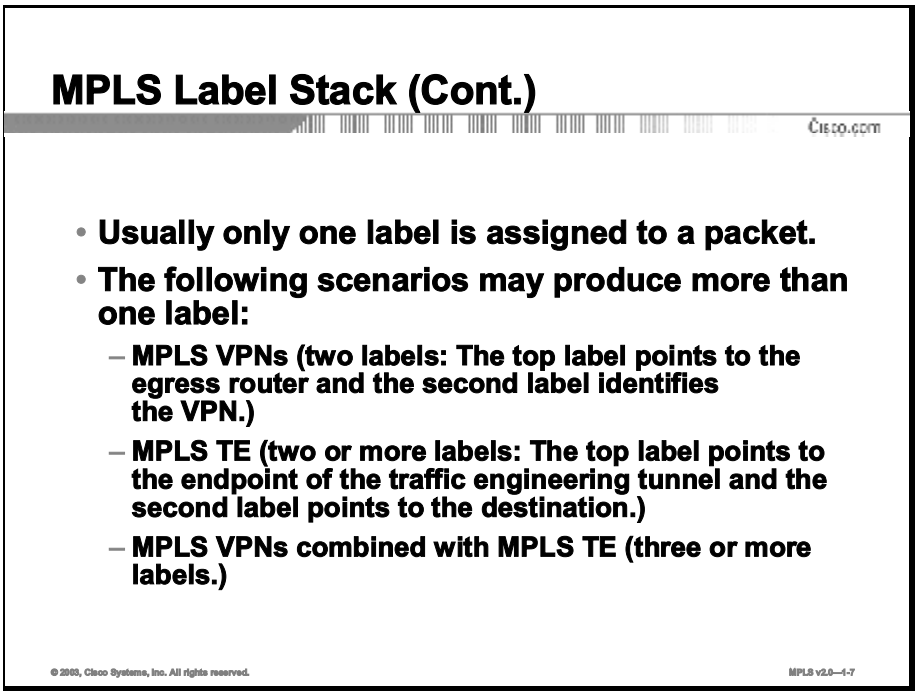

As previously noted, MPLS supports multiple labels in one packet. Simple MPLS uses just one label in each packet. The following applications may add labels to packets:

- **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.

MPLS Forwarding

This topic lists the steps involved with MPLS forwarding.

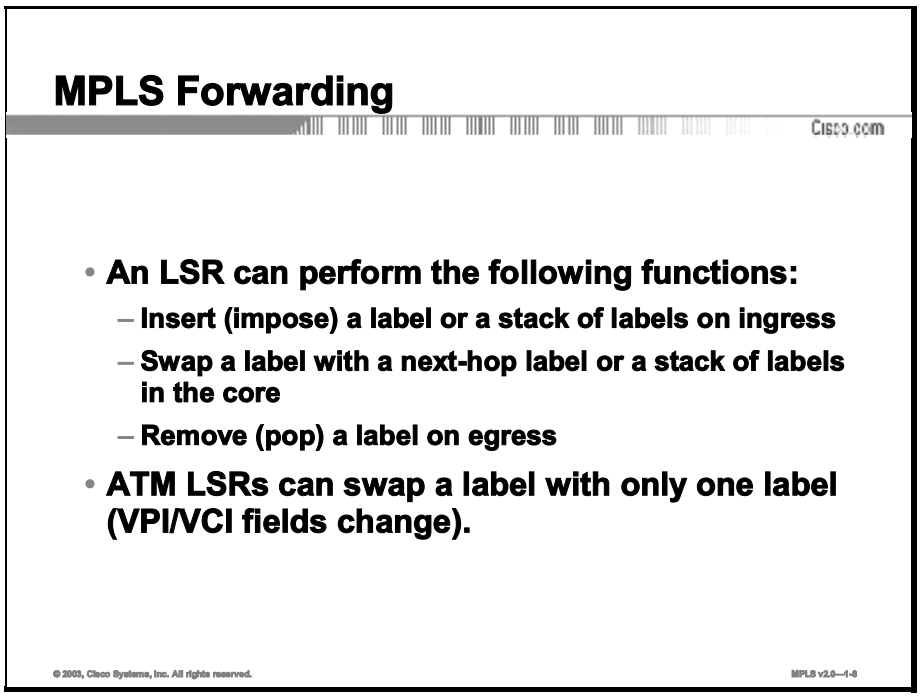

An IP packet going through an MPLS domain experiences the following:

- 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).

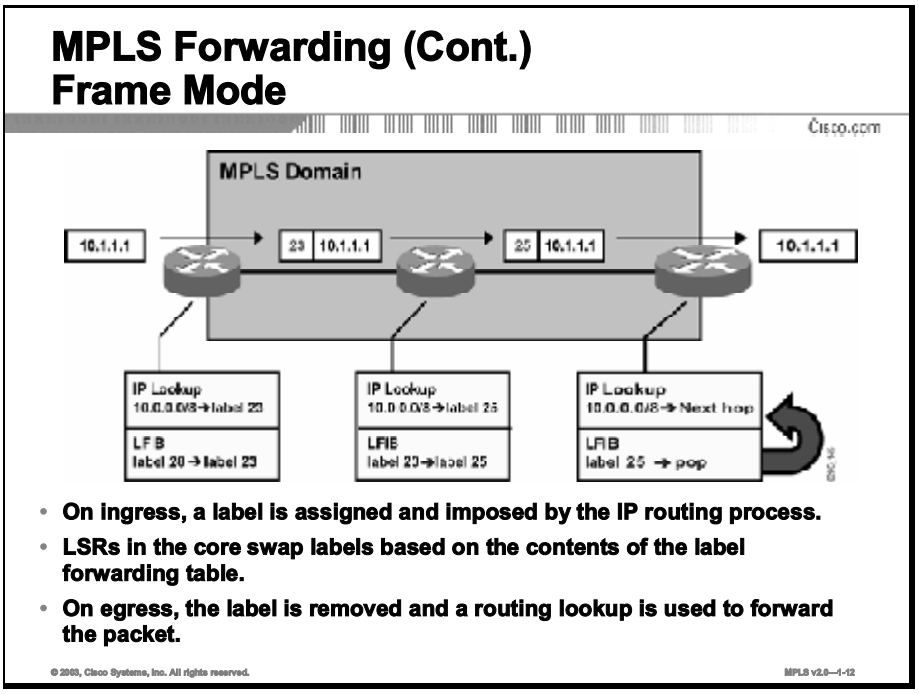

The figure shows an MPLS network using frame-mode MPLS.

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.

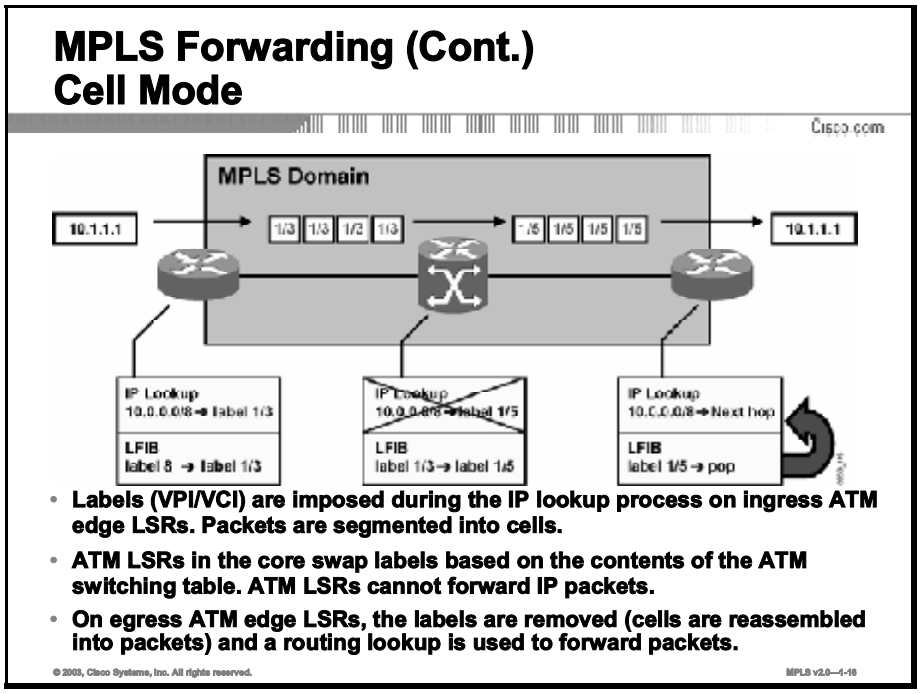

Cell-mode MPLS is similar to frame-mode MPLS. The difference is that ATM LSRs (ATM switches) cannot forward IP packets.

Summary

This topic summarizes the key points discussed in this lesson.

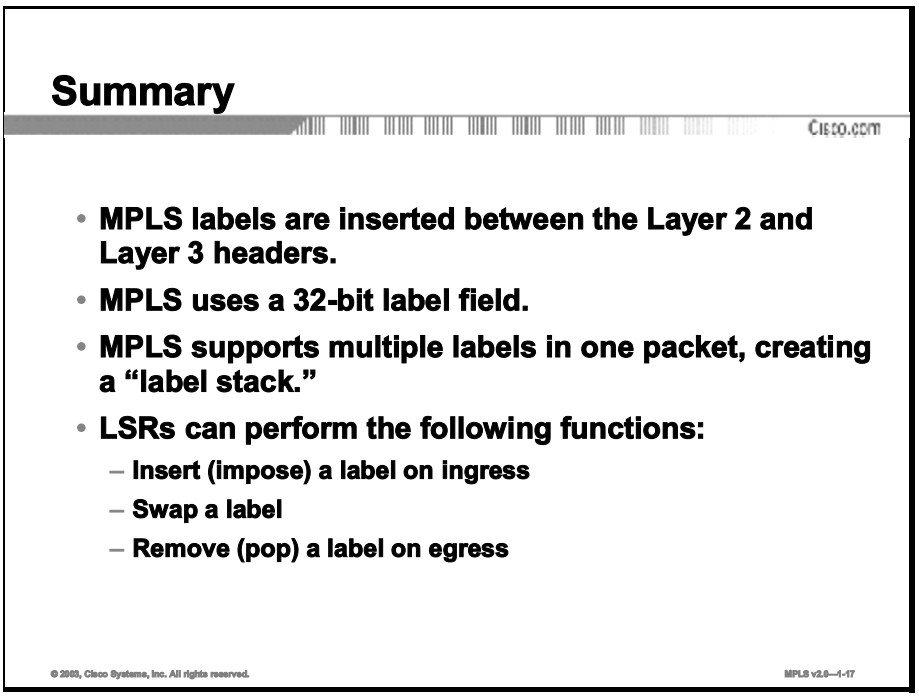

References

For additional information, refer to these resources:

- RFC 3031, "Multiprotocol Label Switching Architecture"
- RFC 3032, "Label Stack Encoding"

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two of the following are true? (Choose two.)

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.

Q2) The actual MPLS label, contained in the MPLS label field, is how long?

A) 32 bits long

B) 8 bits long

C) 16 bits long

D) 20 bits long

Q3) Which two of the following are true? (Choose two.)

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.

Q4) Which two of the following are normal functions of an LSR? (Choose two.)

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

Quiz Answer Key

Q1) A, D

Relates to: MPLS Labels

Q2) D

Relates to: MPLS Label Format

Q3) A, C

Relates to: MPLS Label Stack

Q4) A, D

Relates to: MPLS Forwarding

MPLS Applications

Overview

This lesson looks at some of the different types of applications where you can use MPLS. These applications are discussed at a high level. Each will be discussed in much greater detail in later modules. Interaction among multiple applications is also discussed because there are various methods for exchanging labels. No matter the differences in the control plane, all of the applications use a single label-forwarding engine in the data plane.

Relevance

It helps to be able to apply MPLS principles with applications that are in use today. Doing so helps clarify the picture for the learner.

Objectives

This lesson describes the different MPLS applications, including MPLS VPNs and MPLS TE.

Upon completing this lesson, you will be able to:

- Describe the various applications that are used with MPLS
- Describe MPLS use in unicast IP routing
- Describe MPLS use in multicast IP routing
- Describe MPLS use in traffic engineering environments
- Describe MPLS use in QoS environments
- Describe MPLS use in VPNs
- Identify the interactions that occur between various MPLS applications

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of routing principles and MPLS concepts

Outline

This lesson includes these topics:

- Overview
- MPLS Applications
- Unicast IP Routing
- Multicast IP Routing
- MPLS Traffic Engineering
- Quality of Service
- Virtual Private Networks
- Interactions Between MPLS Applications
- Summary
- Quiz

MPLS Applications

This topic describes various applications used in IP routing where MPLS may also be used.

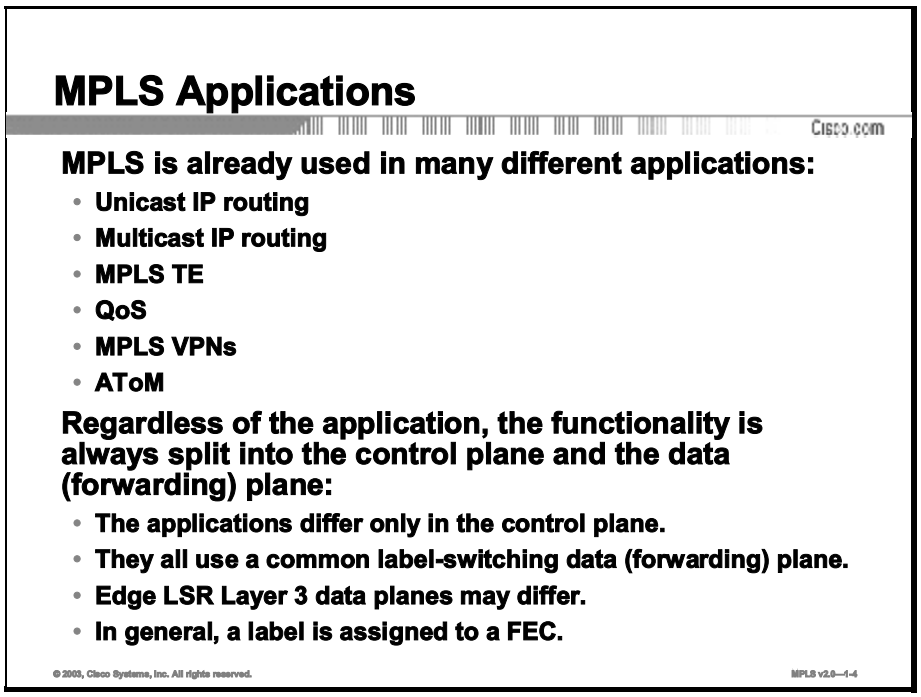

MPLS can be used in different applications:

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

Unicast IP Routing

This topic provides an overview of unicast IP routing.

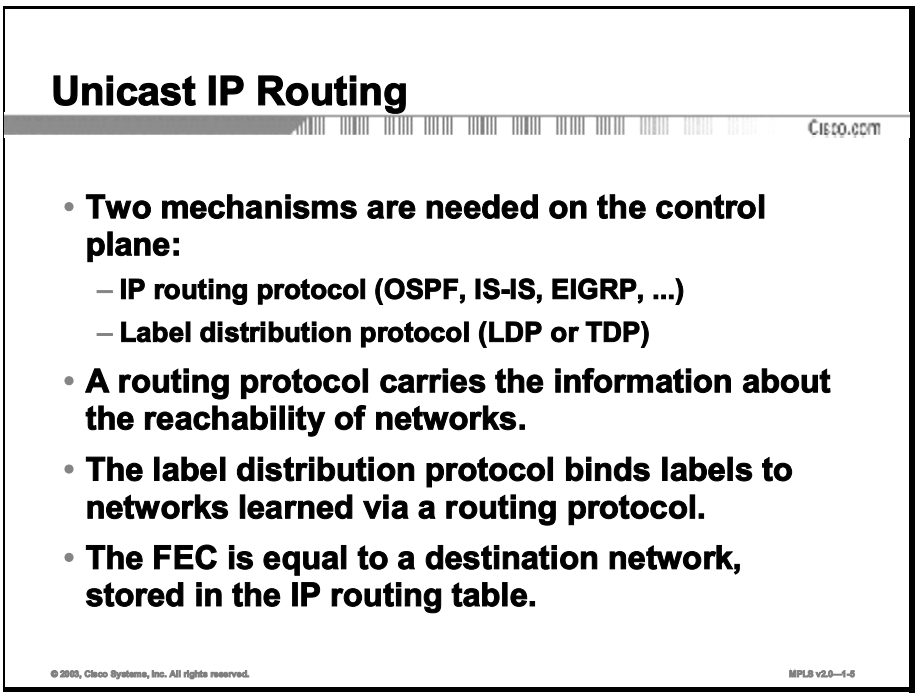

A unicast IP routing setup usually requires two components:

- 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.

Multicast IP Routing

This topic provides an overview of multicast IP routing.

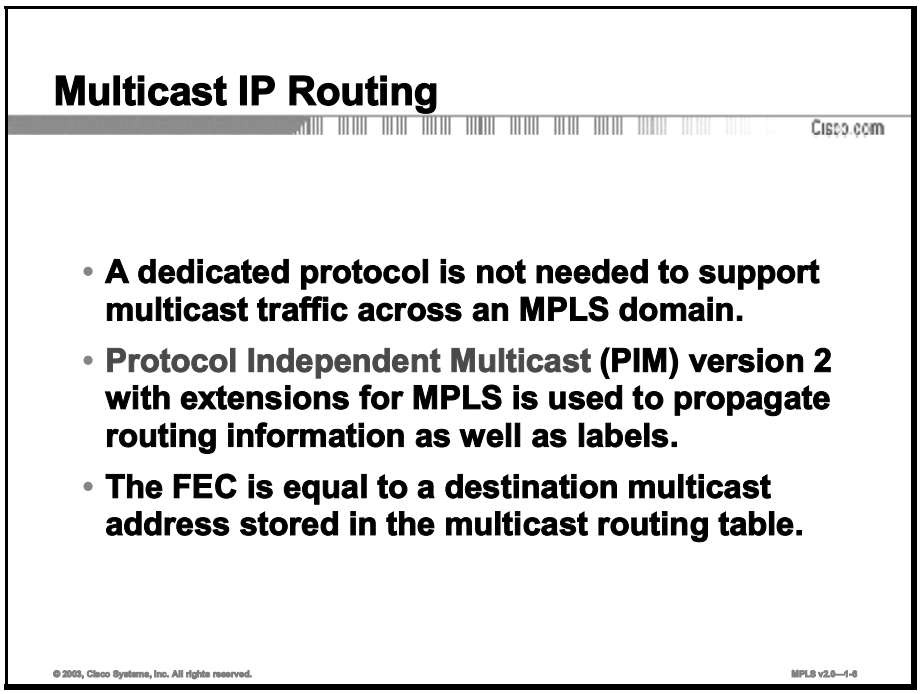

Multicast IP routing can also use MPLS. Cisco Protocol Independent Multicast (PIM) version 2 with extensions for MPLS is used to propagate routing information and labels.

The FEC is equal to a destination multicast address.

MPLS Traffic Engineering

This topic provides an overview of MPLS TE.

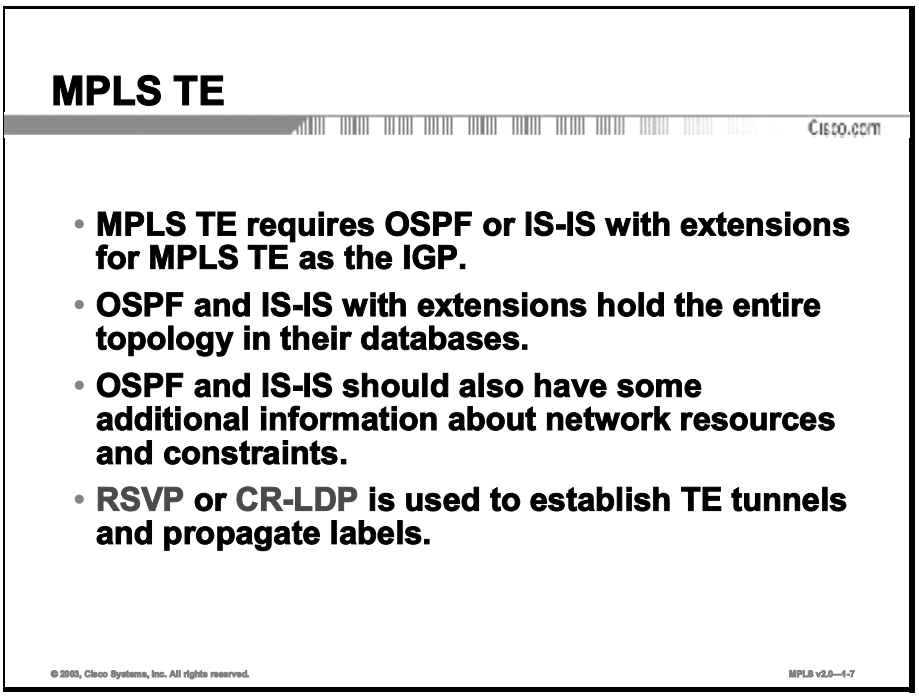

MPLS TE has special requirements:

- 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 CR-LDP (Constraint Route-LDP) is used to establish traffic engineering (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.

Quality of Service

This topic provides an overview of QoS.

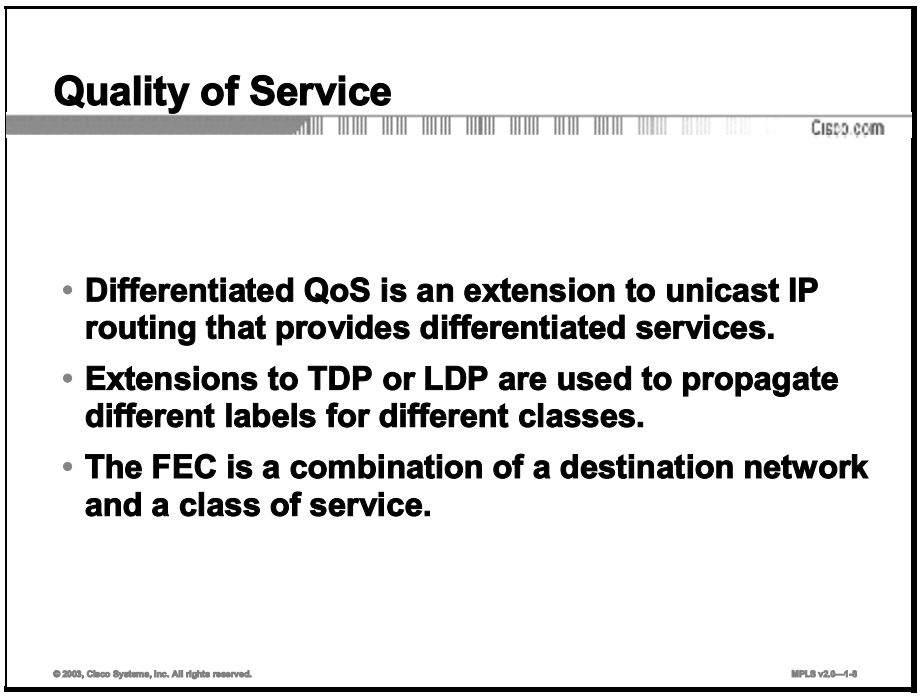

Differentiated QoS is achieved by using MPLS experimental bits or by creating separate LSP tunnels for different classes. Extensions to TDP or LDP are used to create multiple LSP tunnels for the same destination (one for each class).

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

Virtual Private Networks

This topic provides an overview of VPNs.

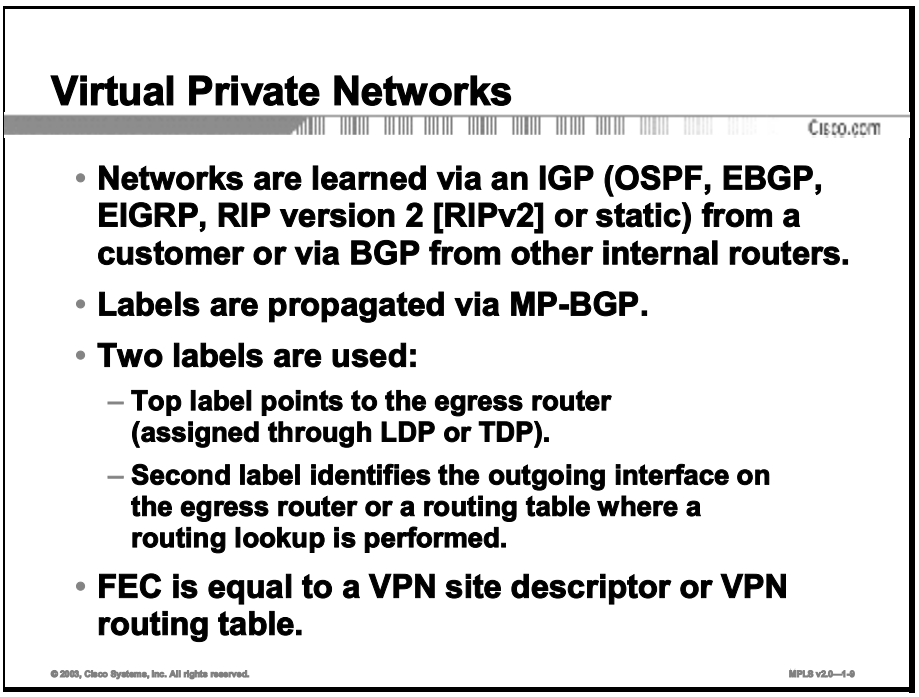

MPLS VPNs use an additional label to determine the VPN and the corresponding VPN destination network. MP-BGP is used to propagate VPN routing information and labels across the MPLS domain. TDP or LDP is needed to link edge LSRs with a single LSP tunnel.

The FEC is equal to a VPN destination network.

Interactions Between MPLS Applications

This topic briefly describes the interactions that occur between MPLS applications.

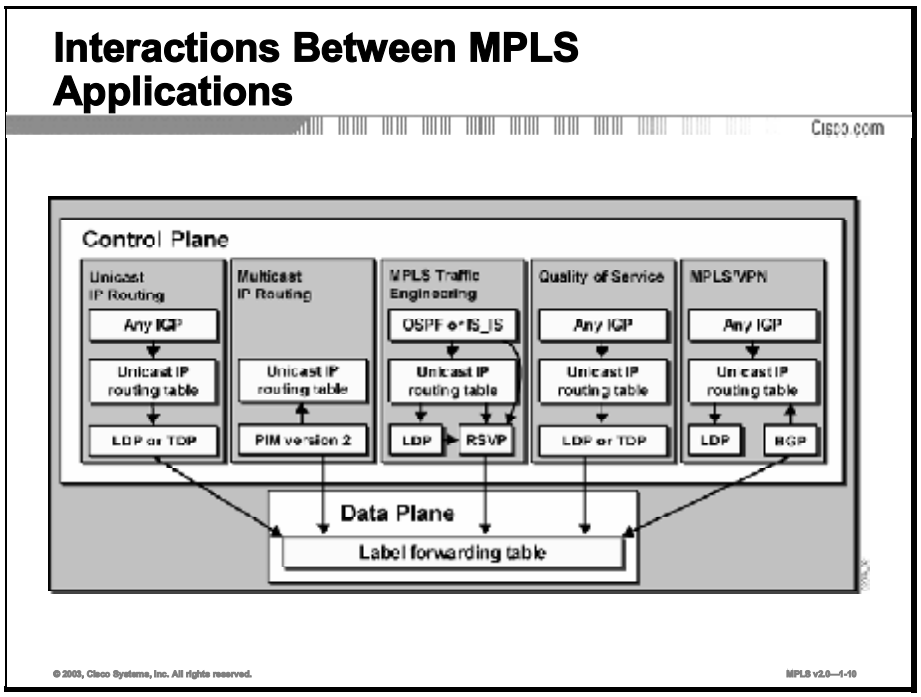

The figure shows the complete architecture when all applications are used. Each application may use a different routing protocol and a different label exchange protocol, but all applications use a single label-forwarding engine.

Summary

This topic summarizes the key points discussed in this lesson.

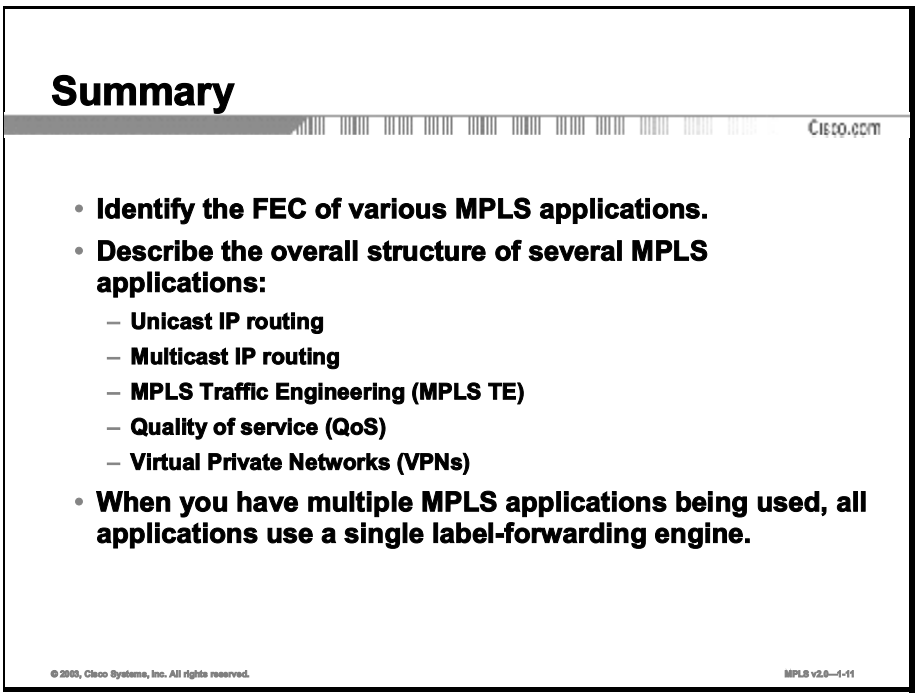

References

For additional information, refer to this resource:

- RFC 3031, “Multiprotocol Label Switching Architecture”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The label distribution protocol is found on what plane?

A) forwarding plane

B) data plane

C) control plane

D) ground plane

Q2) Which two of the following are true regarding RSVP? (Choose two.)

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.

Q3) When MPLS is used for QoS, which of the following is true?

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.

Q4) In MPLS VPN networks which of the following is true?

A) Labels are propagated via LDP or TDP.

B) Labels are not used in an MPLS VPN network, just next-hop addresses.

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.

Q5) Which two statements are true regarding interactions between MPLS applications? (Choose two.)

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.

Quiz Answer Key

Q1) C

Relates to: Unicast IP Routing

Q2) A, B

Relates to: MPLS Traffic Engineering

Q3) C

Relates to: Quality of Service

Q4) C

Relates to: Virtual Private Networks

Q5) A, D

Relates to: Interactions Between MPLS Applications

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: MPLS Concepts

Complete the quiz to assess what you have learned in this module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Describe the basic MPLS concepts, including the drawbacks in traditional IP routing
- Describe MPLS labels and MPLS label stacks, including the format of the MPLS label and also when and why a label stack is created
- Describe the different MPLS applications, including MPLS VPNs and MPLS TE

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key at the end of this quiz.

Step 3 Review the related lesson for each question that you answered incorrectly.

Quiz

Answer these questions:

Q1) MPLS labels can correspond to which of the following?

A) Layer 2 source addresses

B) Layer 3 source addresses

C) Layer 2 destination addresses

D) Layer 3 destination addresses

Q2) Which of the following best describes the term "a simple label-based forwarding engine?"

A) control plane

B) ground plane

C) data plane

D) routing plane

Q3) Which two of the following statements are true? (Choose two.)

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.

Q4) Cisco routers automatically assign the IP precedence value to which field in the MPLS label?

A) TTL field

B) experimental field

C) top-of-stack field

D) The IP precedence value is not copied to the MPLS field. It remains in the IP packet.

Q5) Which of the following is NOT a valid Ethertype used to identify Layer 3 protocols with most Layer 2 encapsulations?

A) unlabeled IP unicast (PID = 0x0800)

B) labeled IP unicast (PID = 0x0847)

C) unlabeled IP multicast (PID = 0x8846)

D) labeled IP multicast (PID = 0x8848)

Q6) In MPLS VPNs what does the FEC refer to?

A) IP destination network

B) MPLS ingress router

C) core of the MPLS network

D) VPN destination network

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) D

Relates to: Basic MPLS Concepts

Q2) C

Relates to: Basic MPLS Concepts

Q3) A, C

Relates to: Basic MPLS Concepts

Q4) B

Relates to: MPLS Labels and Label Stack

Q5) C

Relates to: MPLS Labels and Label Stack

Q6) D

Relates to: MPLS Applications

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. The 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 includes being able to do the following:

- 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

Module Outline

The module contains these lessons:

- Typical Label Distribution in Frame-Mode MPLS
- Convergence in Frame-Mode MPLS
- Typical Label Distribution over LC-ATM Interfaces and VC Merge
- MPLS Label Allocation, Distribution, and Retention Modes
- LDP Neighbor Discovery

Typical Label Distribution in Frame-Mode MPLS

Overview

This lesson talks about how label allocation and distribution function in a frame-mode network. Also covered are penultimate hop popping (PHP) and how the MPLS data structures are built.

Relevance

This lesson is key in understanding the basic fundamentals of how information gets distributed and placed into the appropriate tables for both label and unlabeled packet usage.

Objectives

This lesson describes how the LIB, FIB, and LFIB tables are populated with label information.

Upon completing this lesson, you will be able to:

- Describe MPLS unicast IP routing and architecture
- Describe label switch paths
- Describe how label allocation is done in a frame-mode MPLS network
- Describe how labels are distributed in a frame-node network
- Describe how the LFIB table is populated
- Describe packet propagation across an MPLS network
- Describe frame-mode loop detection
- Describe penultimate hop popping
- Describe per-platform label allocation

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Concepts” module of this course

Outline

This lesson includes these topics:

- Overview
- MPLS Unicast IP Routing Architecture
- Label Switched Paths
- Label Allocation in a Frame-Mode MPLS Network
- Label Distribution and Advertisement
- Populating LFIB
- Packet Propagation Across an MPLS Network
- Frame-Mode Loop Detection
- Penultimate Hop Popping
- Per-Platform Label Allocation
- Summary
- Quiz

MPLS Unicast IP Routing Architecture

This topic introduces how labels are propagated across a network.

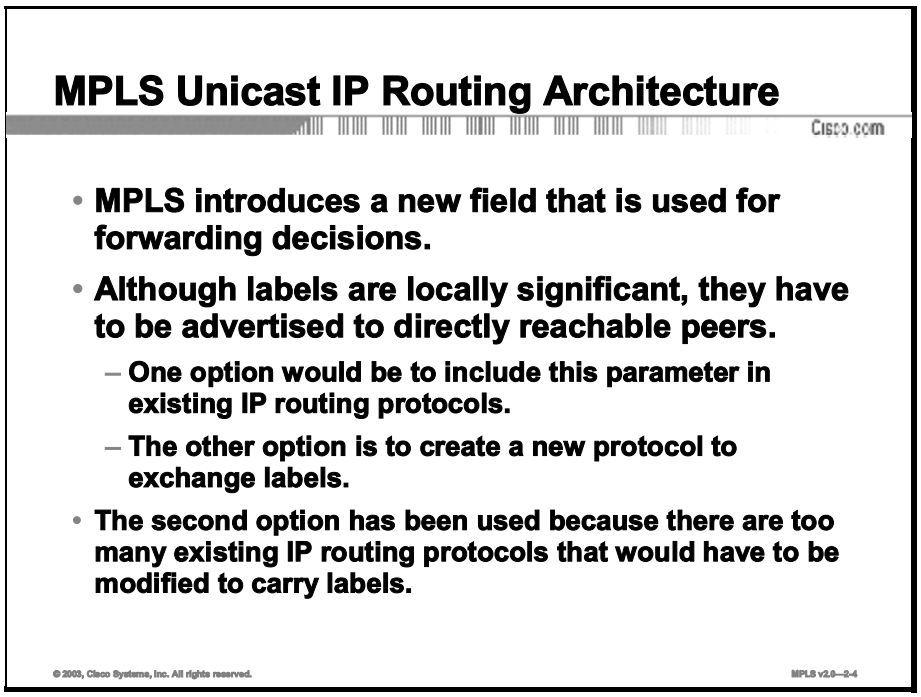

One application of MPLS is unicast IP routing. A label is assigned to destination IP networks and is later used to label packets sent toward those destinations.

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.

There are 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, or OSPF; Intermediate System-to-Intermediate System, or IS-IS; Enhanced Interior Gateway Routing Protocol, or EIGRP; Interior Gateway Routing Protocol, or IGRP; Routing Information Protocol, or RIP; and so on). The 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 option.

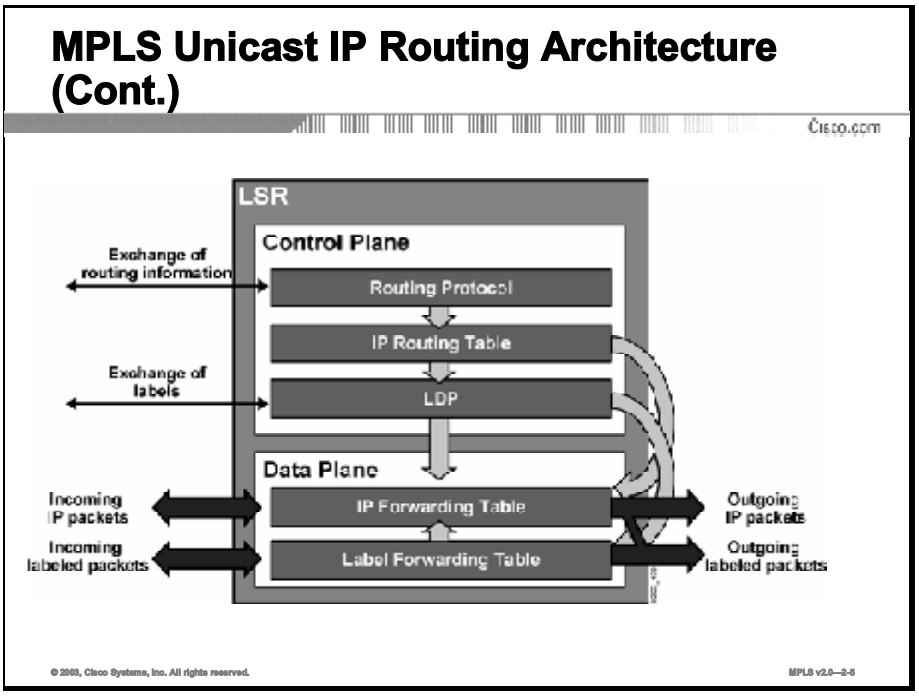

The figure shows the building blocks used by routers to perform traditional IP forwarding.

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 LDP (Label Distribution Protocol) (or the Cisco proprietary protocol TDP, or Tag Distribution Protocol) in the control plane exchanges labels and stores them in the label information base (LIB). This information is then used in the data plane to provide MPLS functionality:

- A label is added to the IP forwarding table (Forwarding Information Base, or FIB) to map an IP prefix to a next-hop label.
- A locally generated label is added to the label forwarding information base (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.

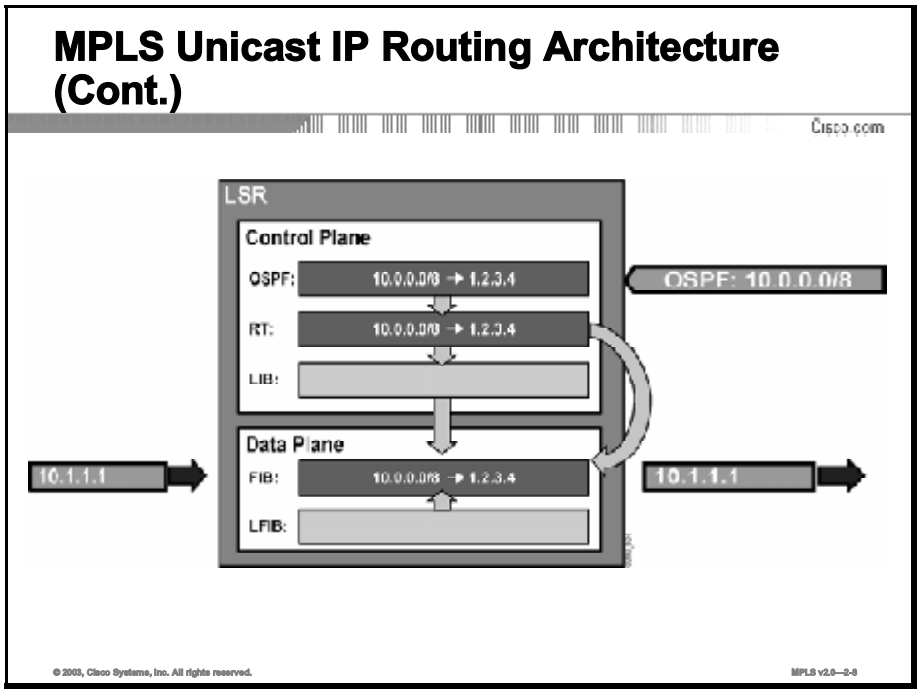

The figure shows a scenario where IP packets are successfully forwarded by using the FIB table.

Labeled packets, on the other hand, are not forwarded due to 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 it.

The example simply 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.

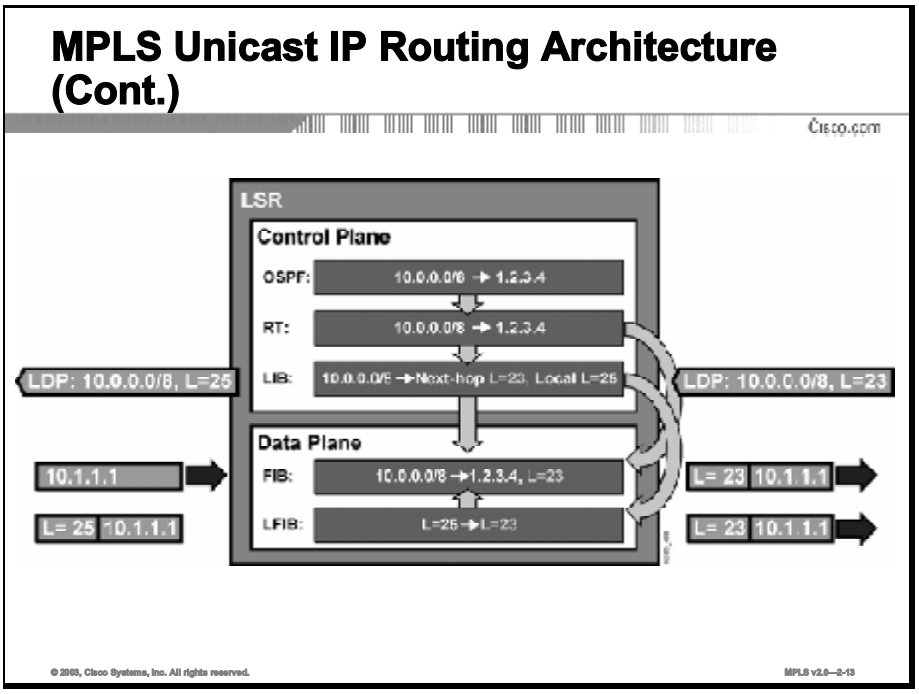

The figure shows a router where OSPF is used to exchange IP routing information and LDP is used to exchange labels.

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.

Label Switched Paths

This topic describes label switched paths (LSPs) and how they are built.

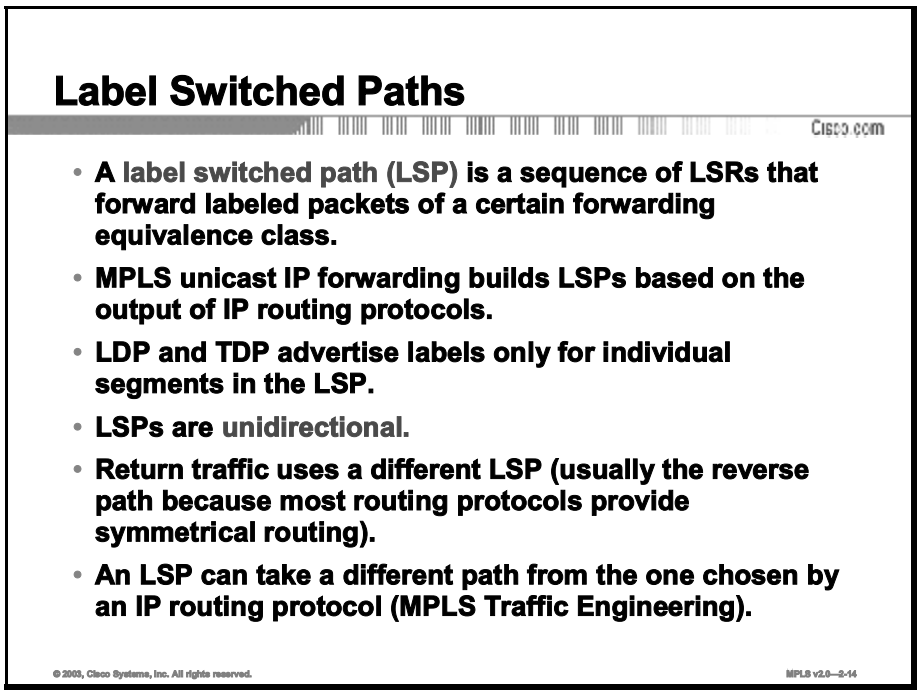

An LSP is a sequence of LSRs that forward labeled packets for a particular FEC. Each LSR swaps the top label in a packet traversing the LSP. An LSP is similar to Frame Relay or ATM virtual circuits. In cell-mode MPLS, an LSP *is* a virtual circuit.

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, or 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 they form the LSP in the opposite order.

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

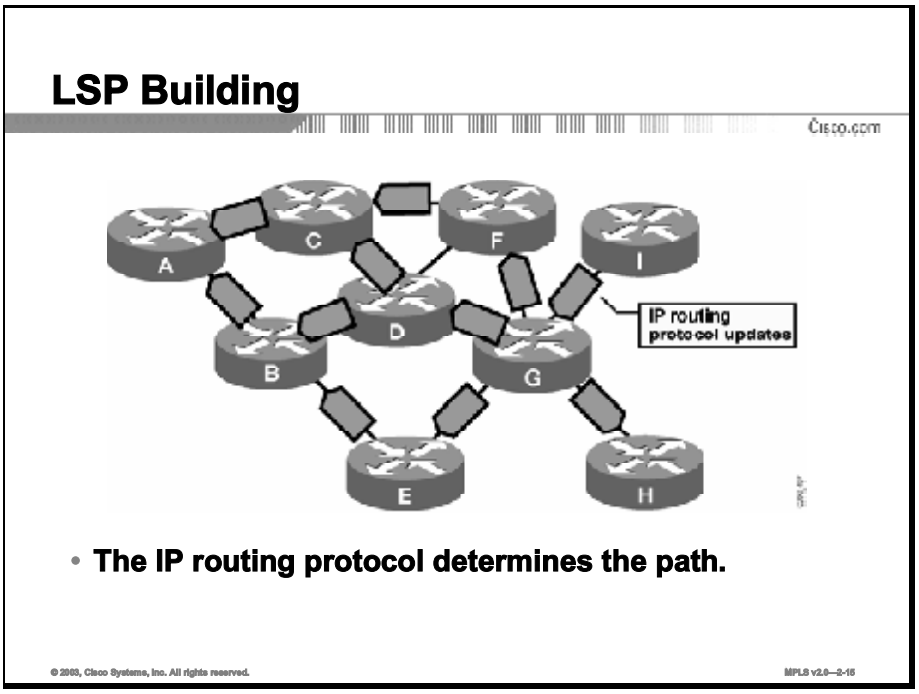

The example here illustrates how an IGP such as OSPF, IS-IS, or EIGRP propagates routing information to all routers in an MPLS domain. Each router determines its own shortest path.

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.

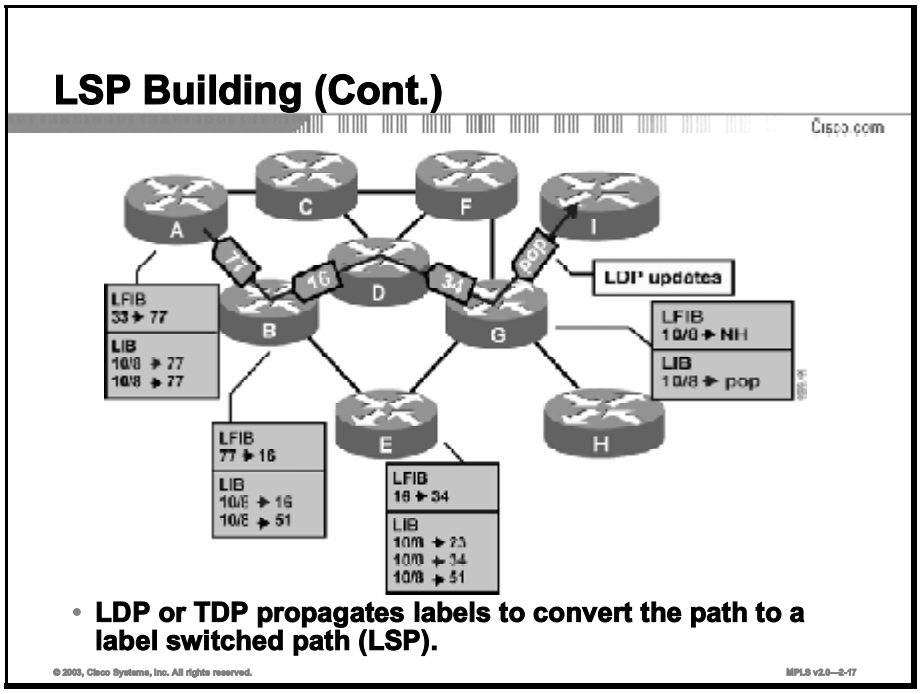

The example here shows the contents of LFIB and LIB tables. Frame-mode MPLS uses a liberal retention mode, which is evident from the contents of the LIB tables. Only those labels that come from the next-hop router are inserted into the LFIB table.

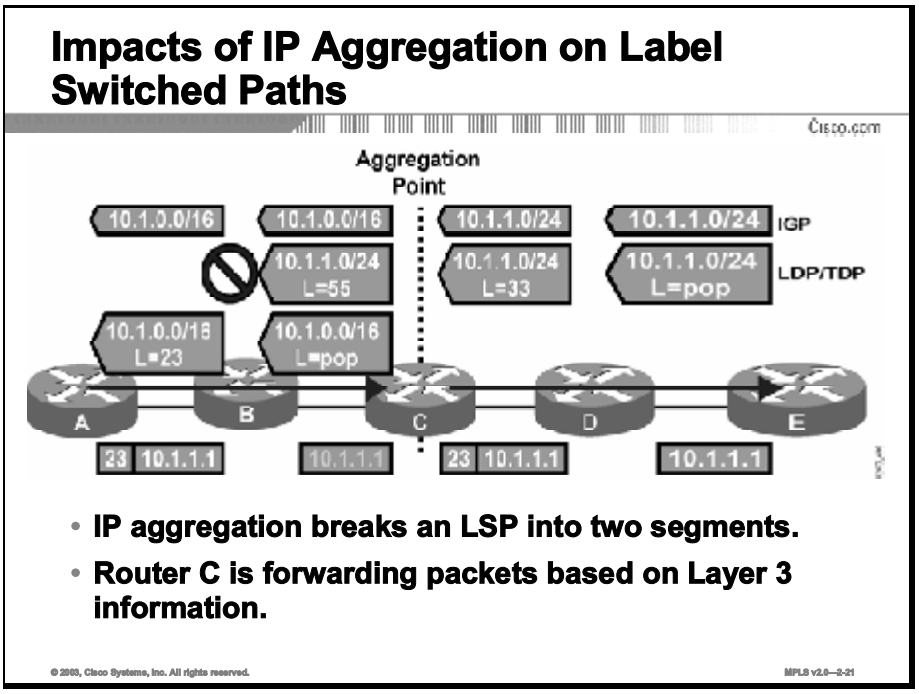

The example here illustrates a potential problem in an MPLS domain.

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 subnets 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:

- 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. LDP also sends a pop label for the new summary network, because it 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, due to the fact that 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.

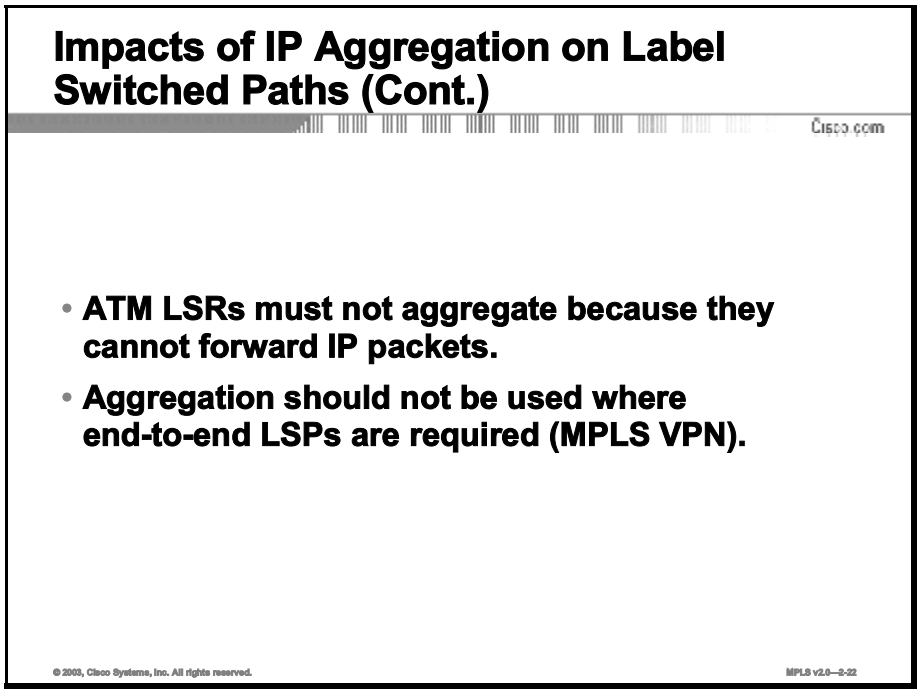

When cell-mode MPLS is used, ATM switches are IP-aware; they run an IP routing protocol, and LDP or TDP, and are generally seen as IP routers. In reality, however, ATM switches are capable of forwarding only cells, not IP packets.

Aggregation (or summarization) should not be used on ATM LSRs because it 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

Label Allocation in a Frame-Mode MPLS Network

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

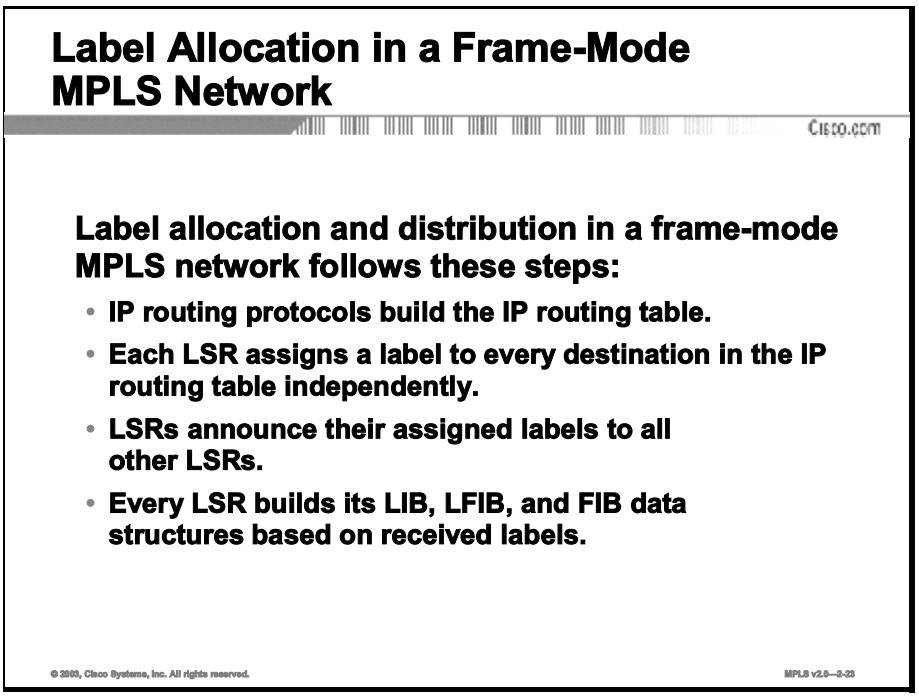

Unicast IP routing and MPLS functionality can be divided into the following steps:

- 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.

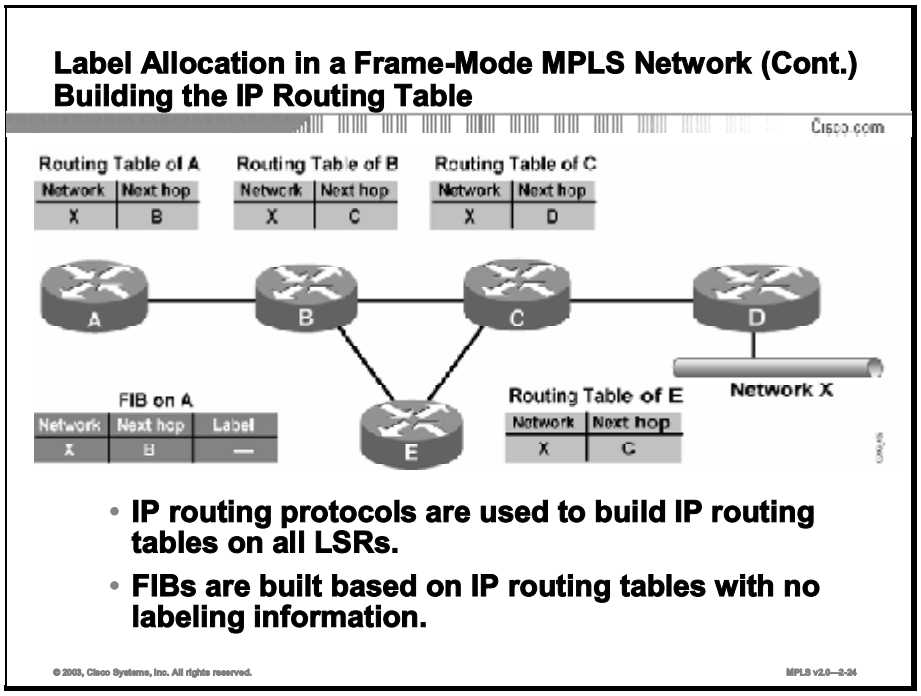

The figure illustrates how all routers learn about network X via an IGP such as OSPF, IS-IS, or EIGRP. The FIB table on router A contains the entry for network X that is mapped to the IP next-hop address B. At this time a next-hop label is not available, which means that all packets are forwarded in a traditional fashion (as unlabeled packets).

Label Allocation in a Frame-Mode MPLS Network (Cont.)
Allocating Labels

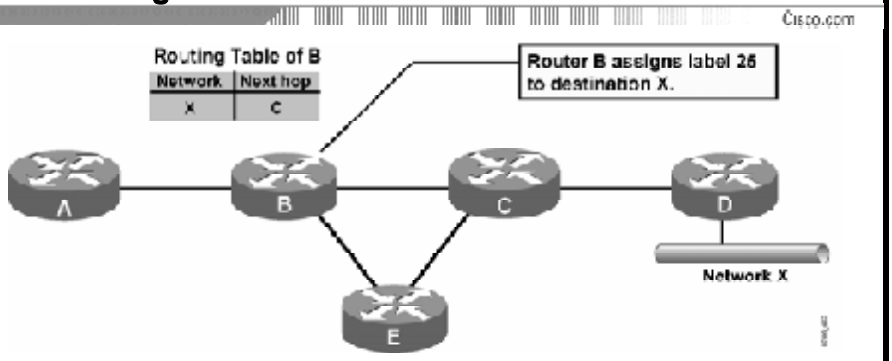

- **Every LSR allocates a label for every destination in the IP routing table.**
- **Labels have local significance.**
- **Label allocations are asynchronous.**

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 – 15 are reserved.

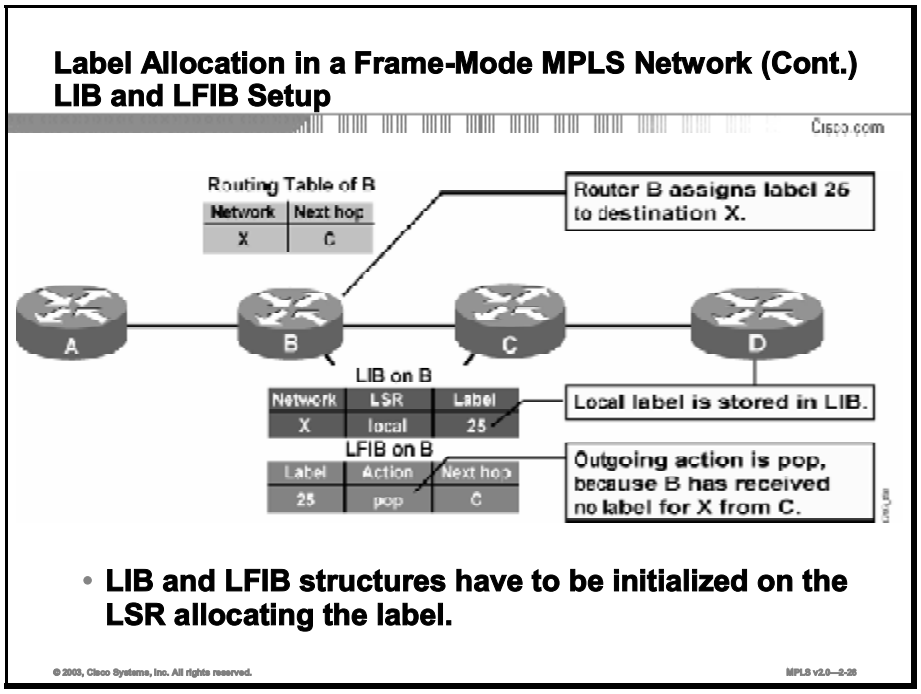

When a label is assigned to an IP prefix, it is stored in two tables:

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

Note The pop action results in labels being removed rather than swapped. Therefore, if a labeled packet is received, the label is removed and the remainder of the packet (which might be labeled if the incoming packet contained a label stack) is forwarded to the appropriate IP next hop. A similar action is the untagged action, which is the equivalent to the pop action with one exception: the resulting packet *must be unlabeled*; otherwise, it will be discarded.

Label Distribution and Advertisement

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

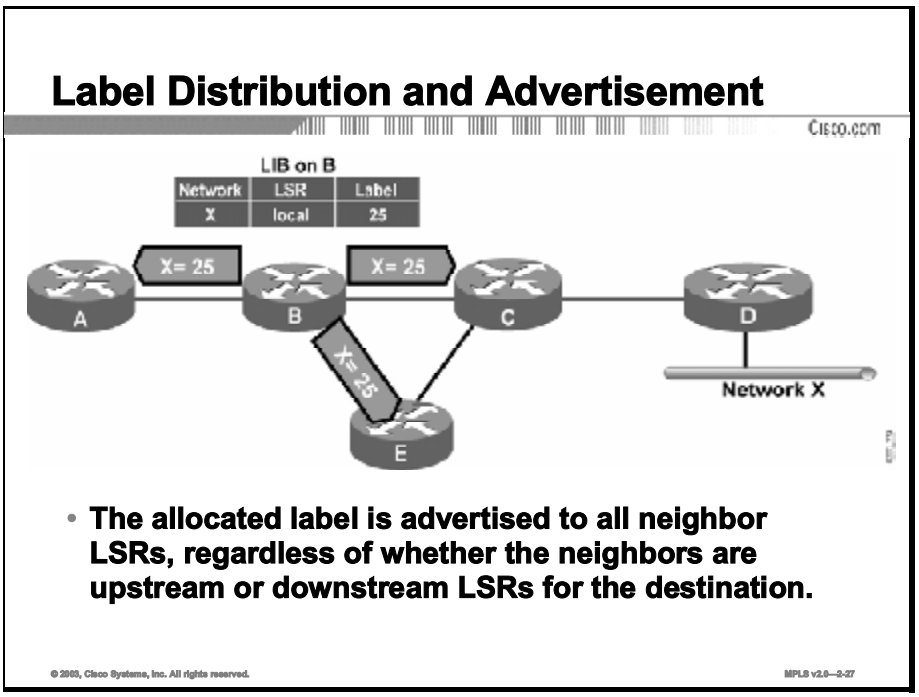

The figure illustrates the next step after a local label has been generated. Router B propagates this label, 25, to all adjacent neighbors where this label can be used as a next-hop label.

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

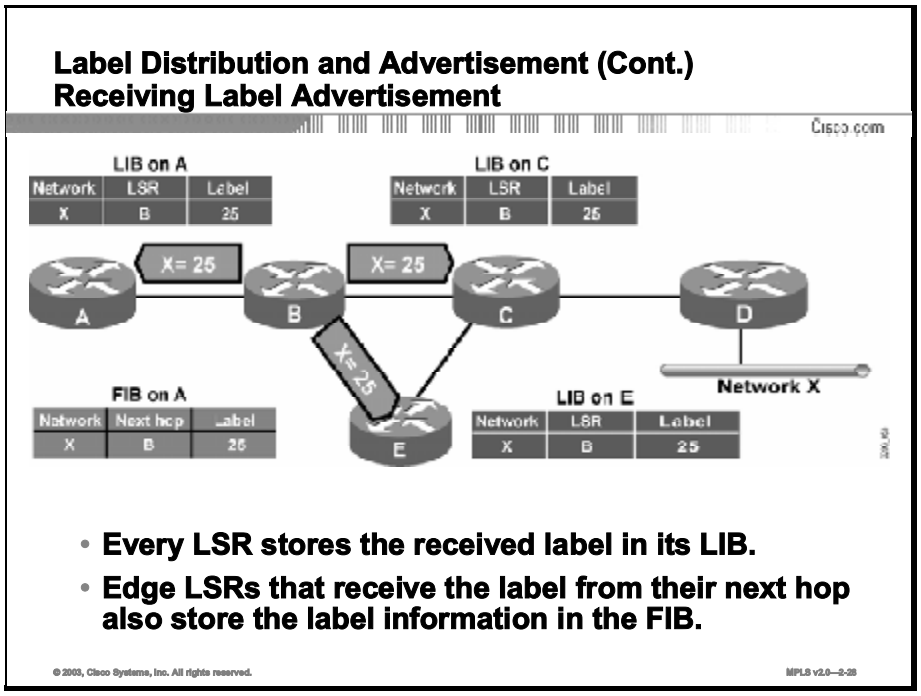

Upon receiving an LDP update, router A can fill in the missing piece in its LIB, LFIB, and FIB tables:

- 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).

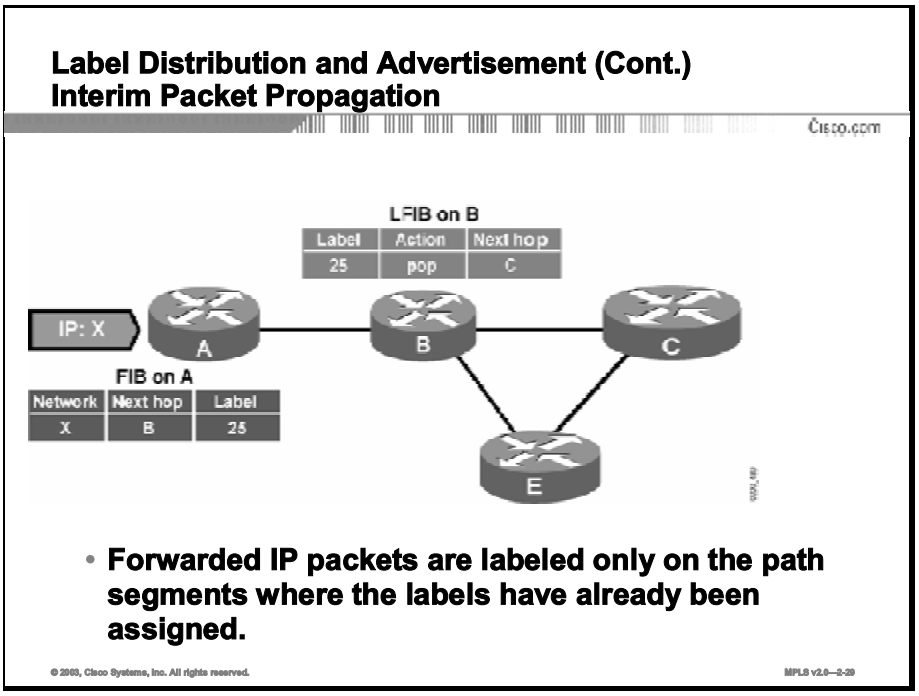

The figure shows how an unlabeled IP packet is forwarded based on the information found in the FIB table on router A. Label 25 found in the FIB table is used to label the packet.

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.

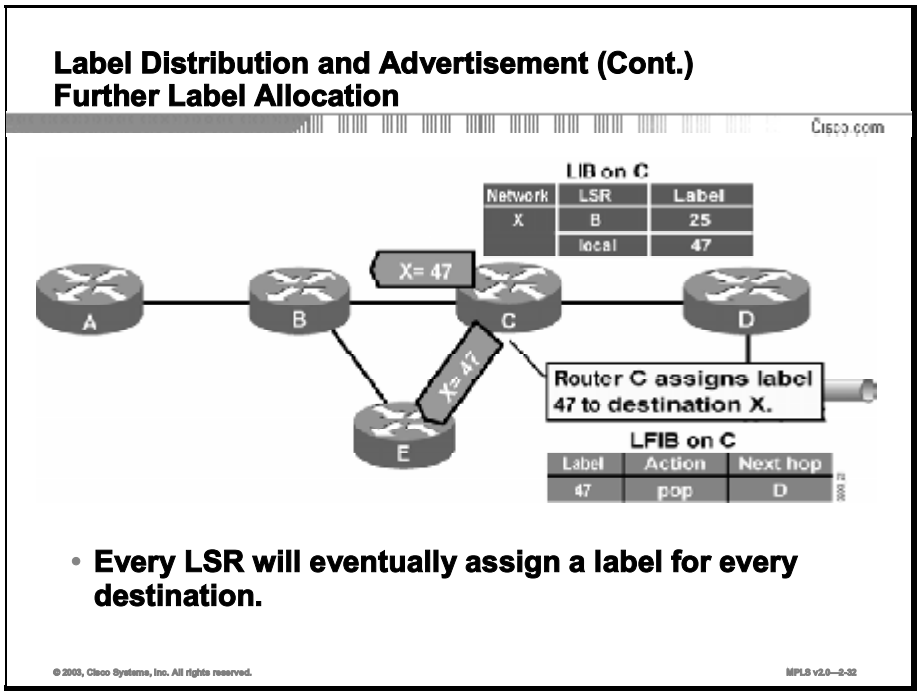

Because all routers in an MPLS domain asynchronously do the same as routers A and B, an LSP tunnel is generated, spanning from router A to router D.

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.

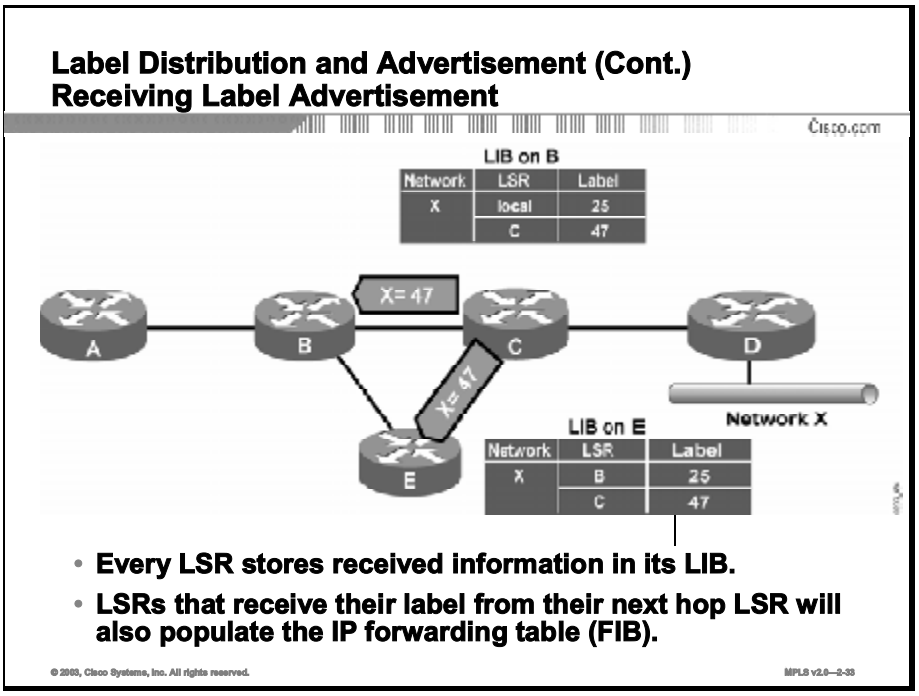

Router B can now map the entry for network X in its FIB, and the local label 25 in its LFIB, to the next-hop label 47 received from the downstream neighbor router C.

Populating LFIB

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

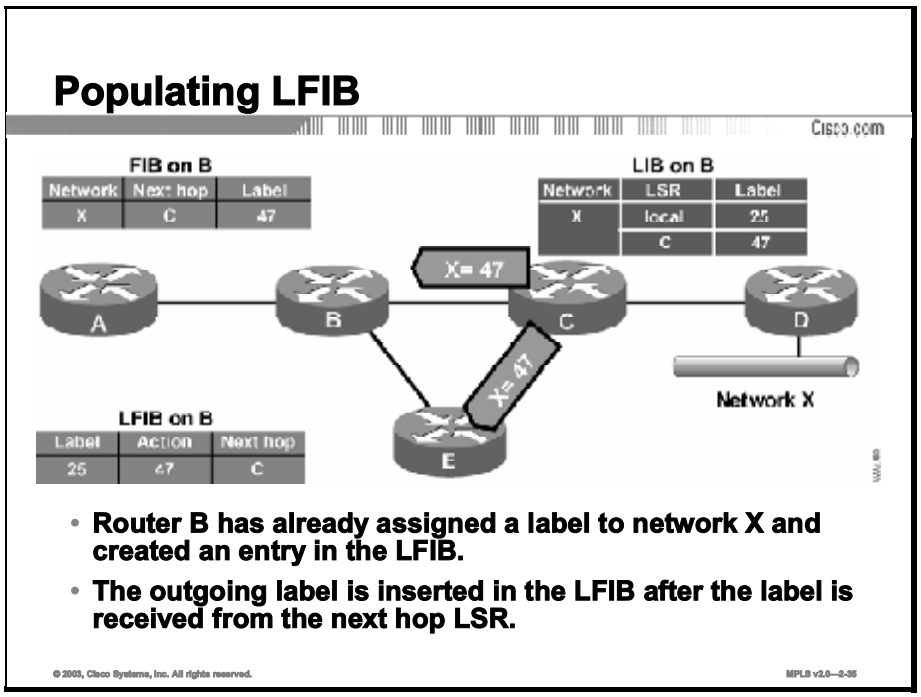

After router C advertises label 47 to adjacent routers, the LSP tunnel for network X has two hops:

- 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 still mapped to the pop action.

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

Packet Propagation Across an MPLS Network

This topic describes how IP packets cross an MPLS network.

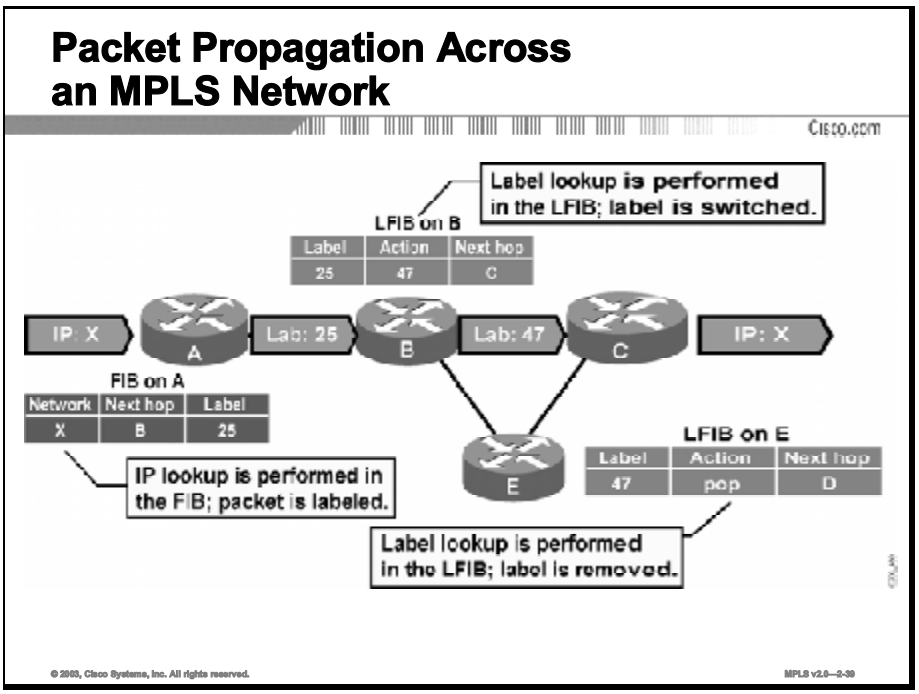

The figure illustrates how IP packets are propagated across an MPLS domain. The steps are as follows:

1. Router A labels a packet destined for network X by using the next-hop label 25 (CEF switching by using the FIB table).
2. Router B swaps label 25 with label 47 and forwards the packet to router C (label switching by using the LFIB table).
3. Router C removes the label and forwards the packet to router D (label switching by using the LFIB table).

Frame-Mode Loop Detection

This topic describes how loop detection can be handled within a frame-mode network.

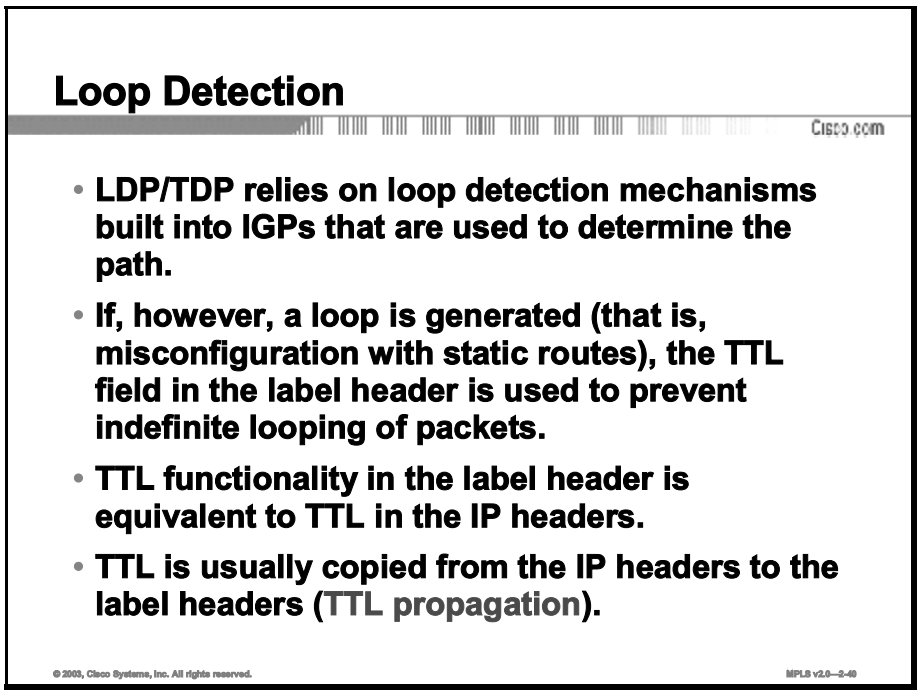

Loop detection in an MPLS-enabled network relies on more than one mechanism.

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.”

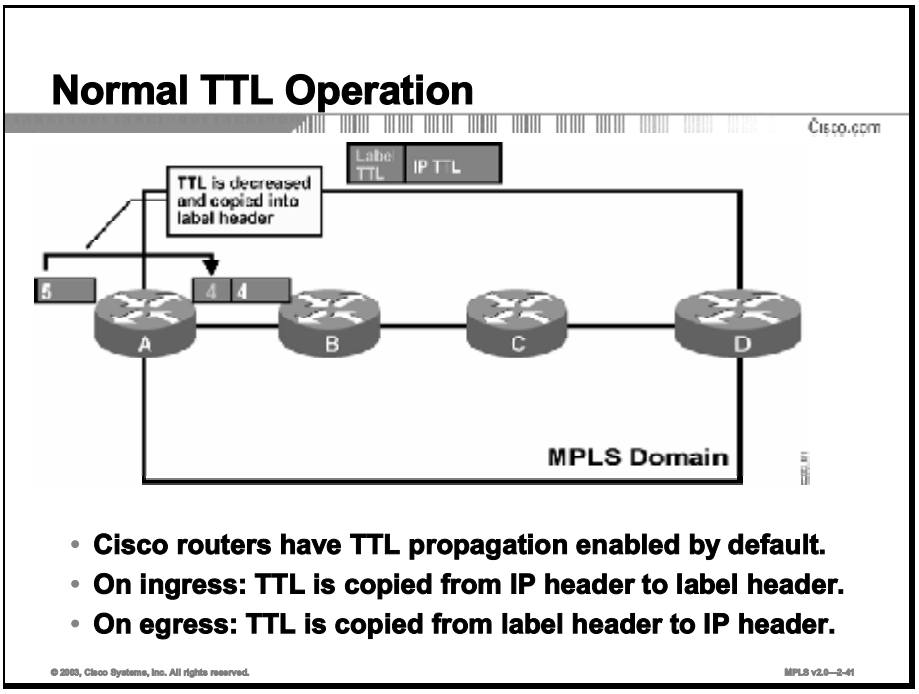

The figure illustrates how the TTL value 5 in the IP header is decreased and copied into the TTL field of the label when a packet enters an MPLS domain.

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.

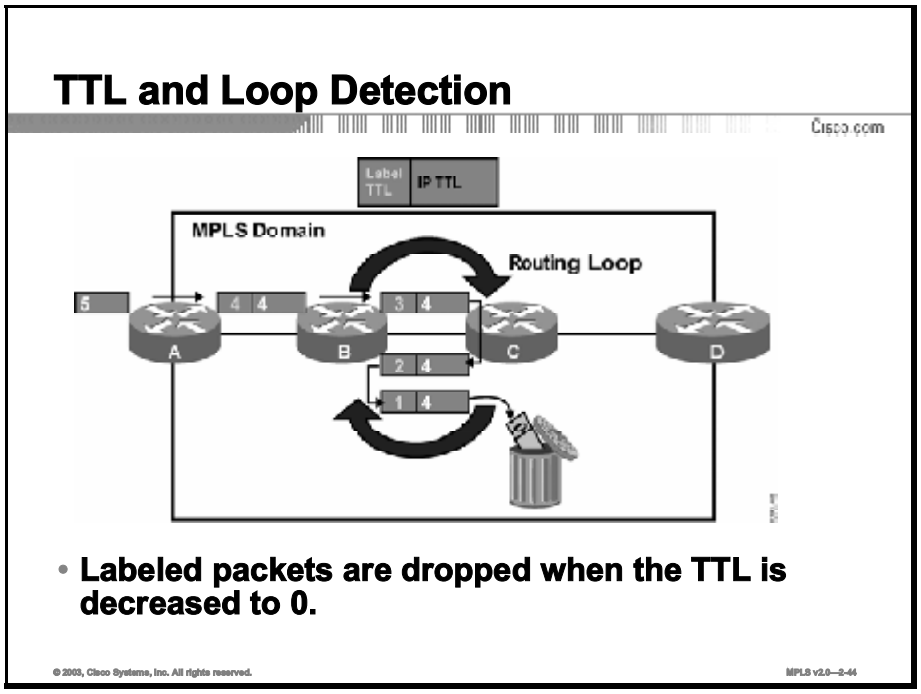

The figure illustrates a routing loop between routers B and C. The packet looping between these two routers is eventually dropped because the value of its TTL field reaches 0.

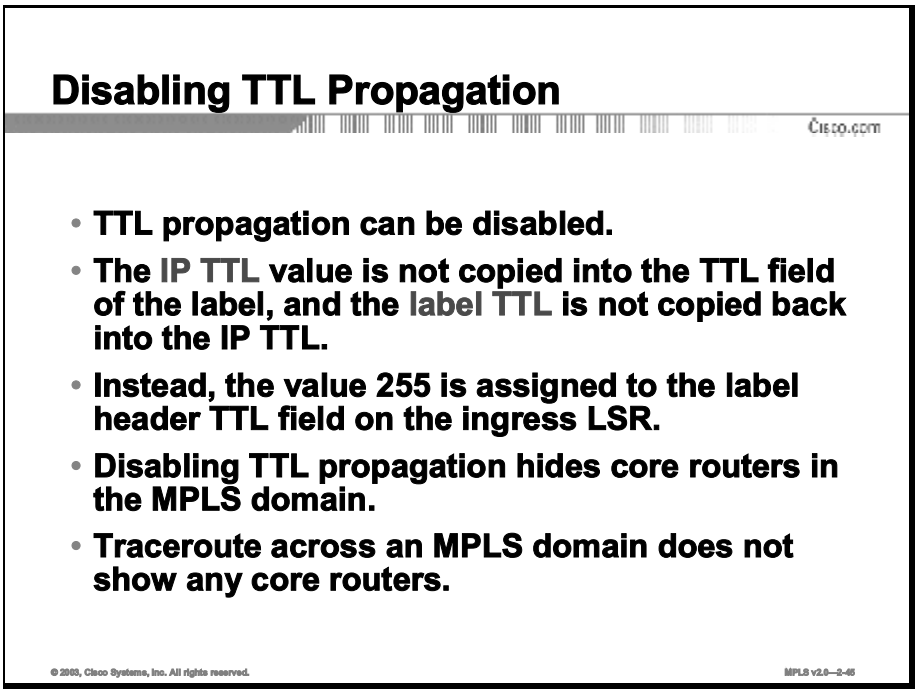

TTL propagation can be disabled to hide the core routers from the end users. Disabling TTL propagation causes routers to set the value 255 into the TTL field of the label when an IP packet is labeled.

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.

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

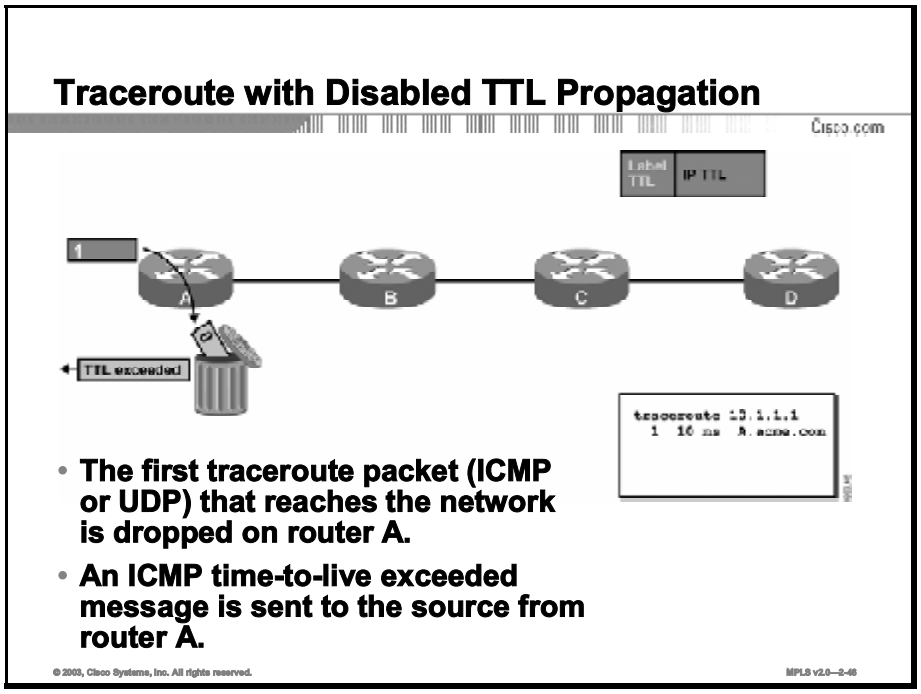

The first traceroute packet (ICMP or User Datagram Protocol [UDP]) that reaches the MPLS network is dropped on the first router (A), and an ICMP reply is sent to the source. This action results in an identification of router A by the traceroute application.

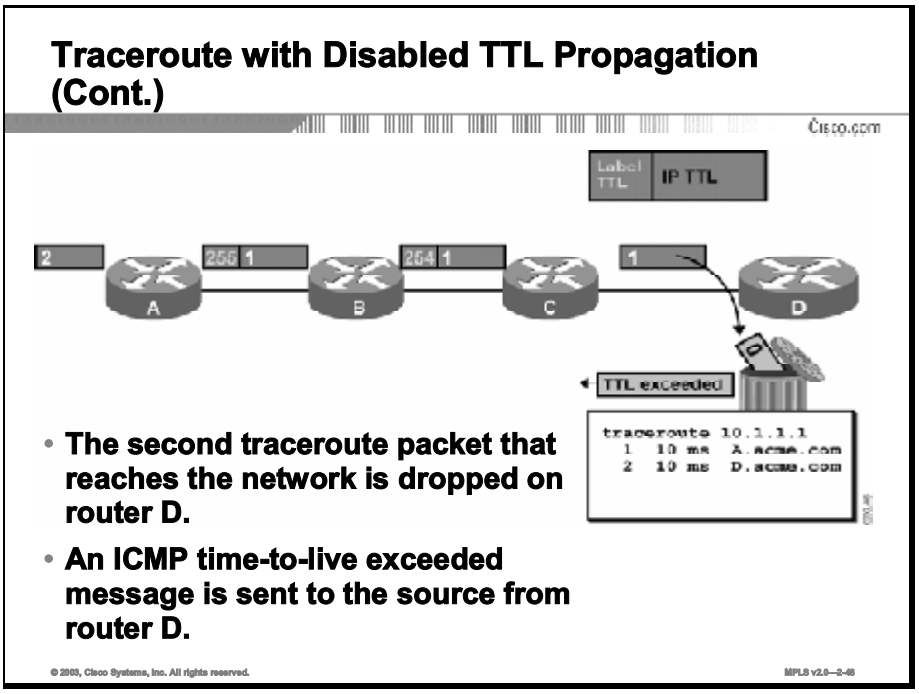

The traceroute application increases the initial TTL for every packet that it sends. The second packet, therefore, would be able to reach one hop farther (router B in the example). However, the TTL value is not copied into the TTL field of the label. Instead, router A sets the TTL field of the label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

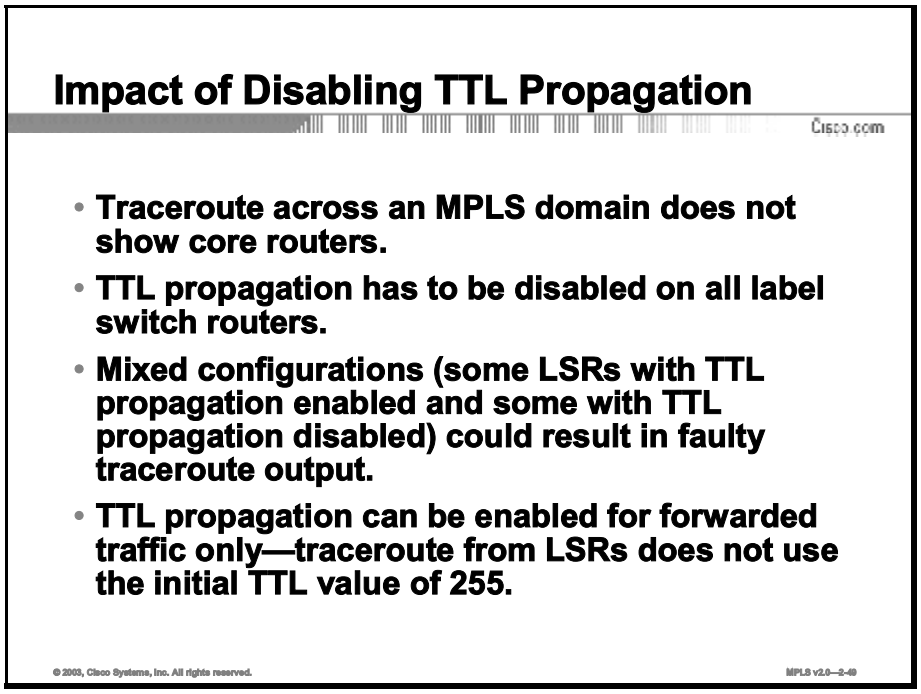

Cisco routers have TTL propagation enabled by default.

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.

Penultimate Hop Popping

This topic describes PHP.

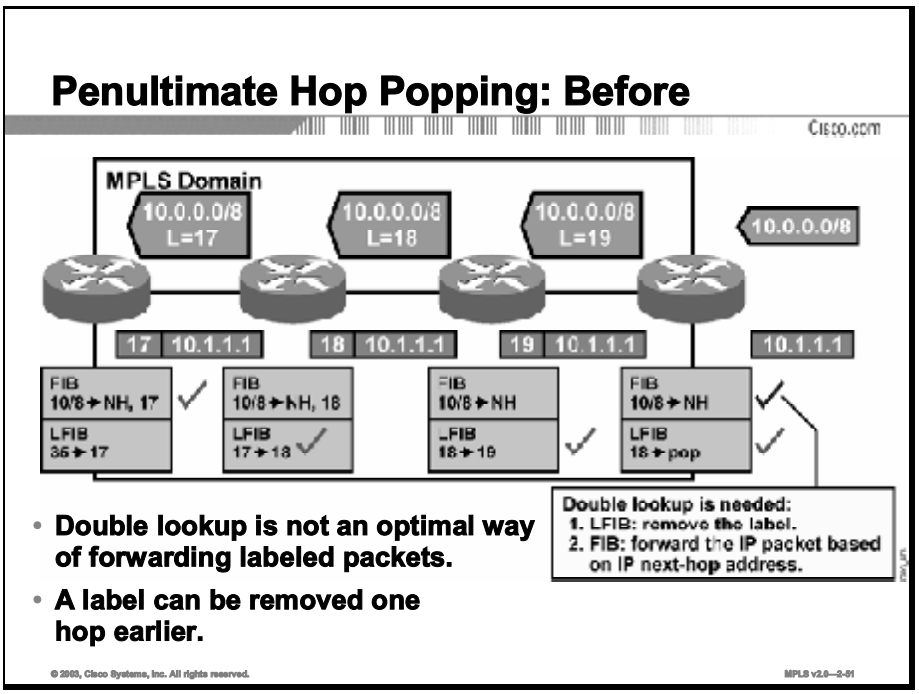

The figure illustrates how labels are propagated and used in a typical frame-mode MPLS network. The check marks show which tables are used on individual routers. The egress router in this example must do a lookup in the LFIB table to determine whether the label must be removed and if a further lookup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

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

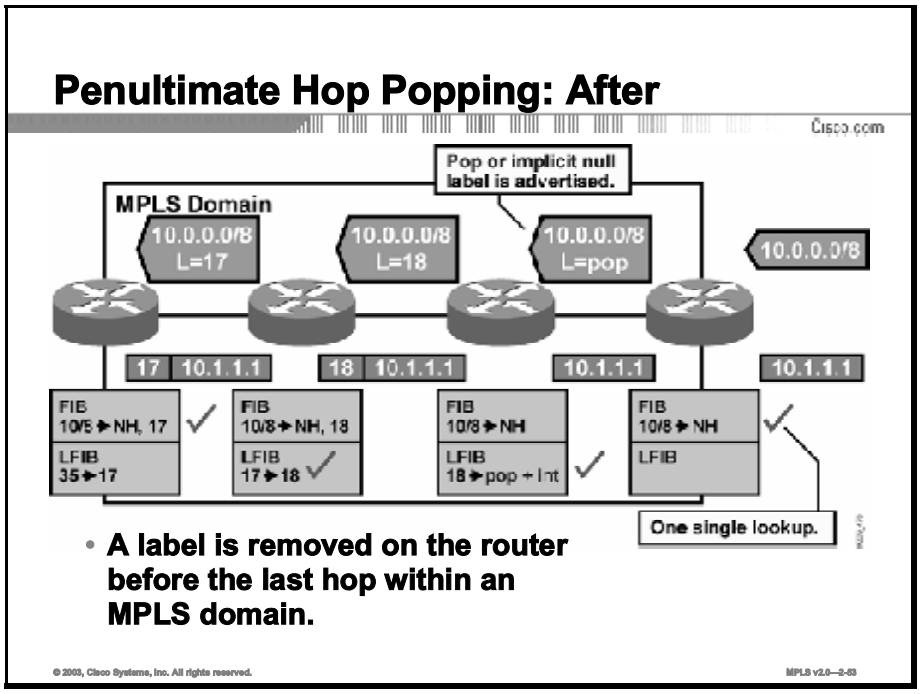

The figure illustrates how a predefined label pop, which corresponds to the pop action in the LFIB, is propagated on the first hop or the last hop, depending on the perspective. The term "pop" means to remove the top label in the MPLS label stack instead of swapping it with the next-hop label. The last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

PHP slightly optimizes MPLS performance by eliminating one LFIB lookup.

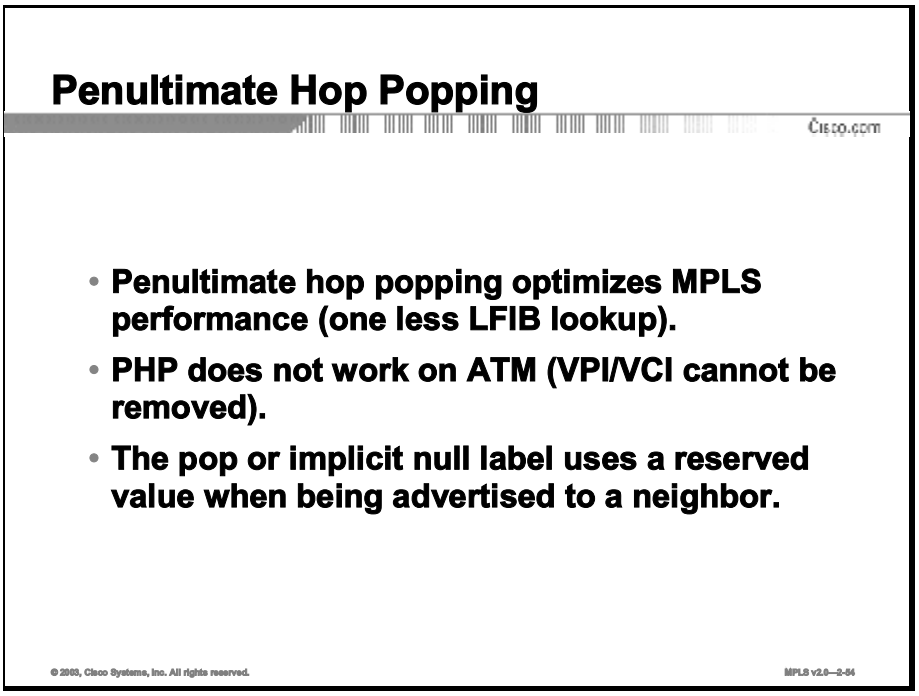

PHP optimizes MPLS performance by reducing the number of table lookups on the egress router.

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 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.

Per-Platform Label Allocation

This topic describes per-platform label allocation.

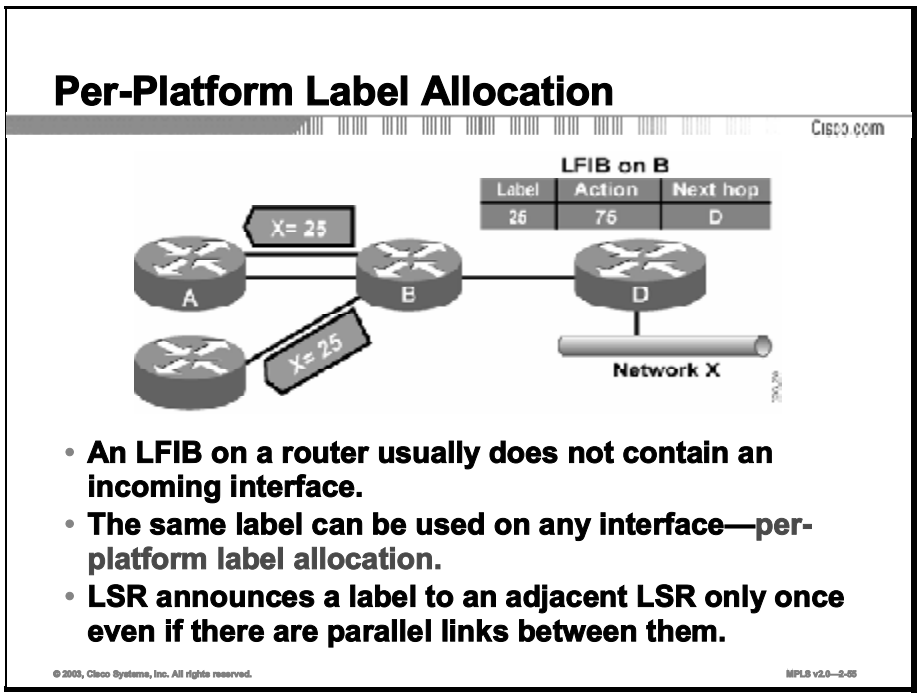

There are two possible approaches for assigning labels to networks:

- **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.

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.

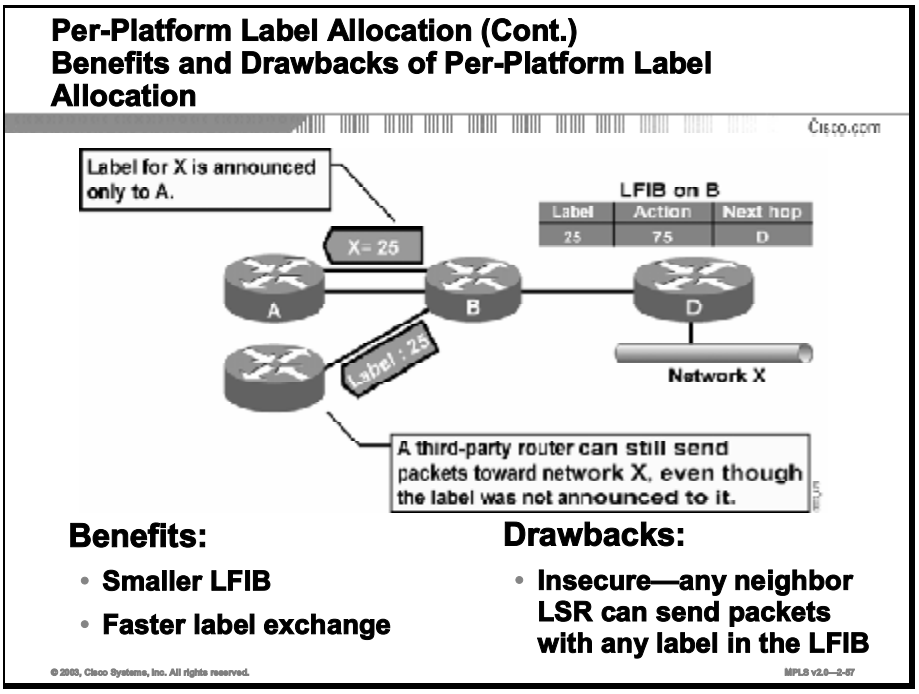

A potential drawback of per-platform label allocation is illustrated in the figure, which shows how an adjacent router can send a labeled packet with a label that has not been previously advertised to this router (label spoofing). If label switching has not been enabled on that interface, the packet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

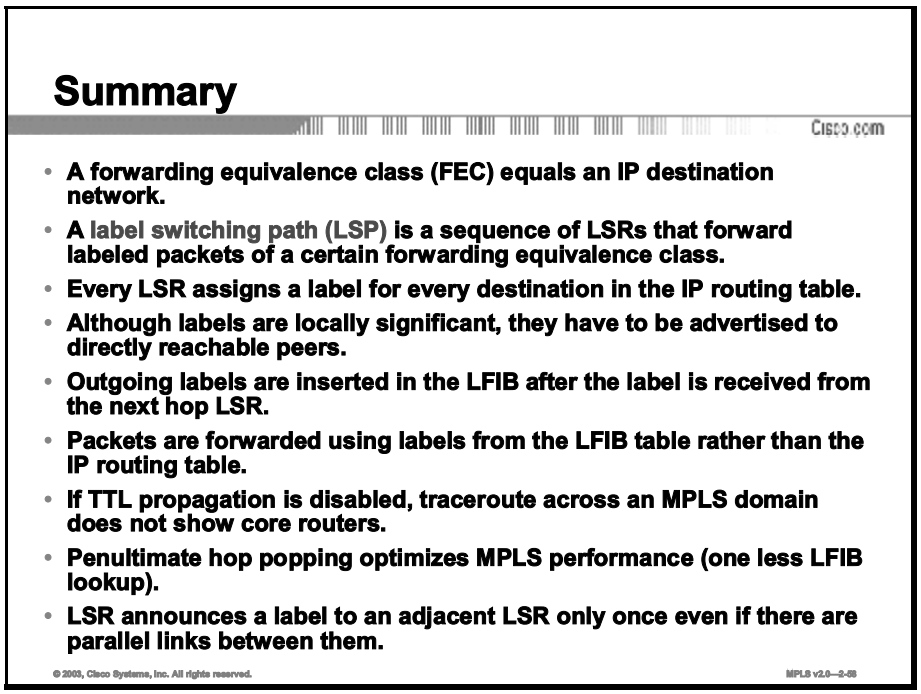

References

For additional information, refer to these resources:

- RFC 3031, “Multiprotocol Label Switching Architecture”
- RFC 3036, “LDP Specification”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following best describes PHP?

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) Per-platform label allocation is which of the following?

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.)

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 of the following tables will be used to perform the label switching?

A) LIB

B) FIB

C) FLIB

D) LFIB

Q5) Which of the following is correct?

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 of the following tables contain label information? (Choose two.)

A) LIB

B) main IP routing label

C) FLIB

D) LFIB

Q7) Which of the following generates a label update?

A) UDP

B) OSPF

C) EIGRP

D) LDP

Q8) Which two of the following are correct? (Choose two.)

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 of the following is correct regarding TTL propagation being disabled?

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.

Quiz Answer Key

Q1) C

Relates to: Penultimate Hop Popping

Q2) A

Relates to: Per-Platform Label Allocation

Q3) A,, B, D

Relates to: Populating LFIB

Q4) D

Relates to: Packet Propagation Across an MPLS Network

Q5) B

Relates to: MPLS Unicast IP Routing Architecture

Q6) A, D

Relates to: Label Allocation in a Frame-Mode MPLS Network

Q7) D

Relates to: Label Distribution and Advertisement

Q8) B, D

Relates to: Label Switched Paths

Q9) C

Relates to: Frame-Mode Loop Detection

Convergence in Frame-Mode MPLS

Overview

This lesson presents LDP convergence issues, and describes how routing protocols and MPLS convergence interact. The lesson concludes with a look at link failure, convergence after a link failure, and link recovery.

Relevance

It is important to understand the convergence times for LDP, and also 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

This lesson describes how convergence occurs in a frame-mode MPLS network

Upon completing this lesson, you will be able to:

- Identify the MPLS steady-state environment
- Describe link failure actions
- Describe routing protocol convergence after a link failure and recovery
- Describe frame-mode MPLS convergence after link failure and link recovery
- Describe link recovery actions

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Typical Label Distribution in Frame-Mode MPLS” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Steady State
- Link Failure Actions
- Routing Protocol Convergence
- MPLS Convergence
- Link Recovery Actions
- Summary
- Quiz

Steady State

This topic describes an MPLS network steady state operation.

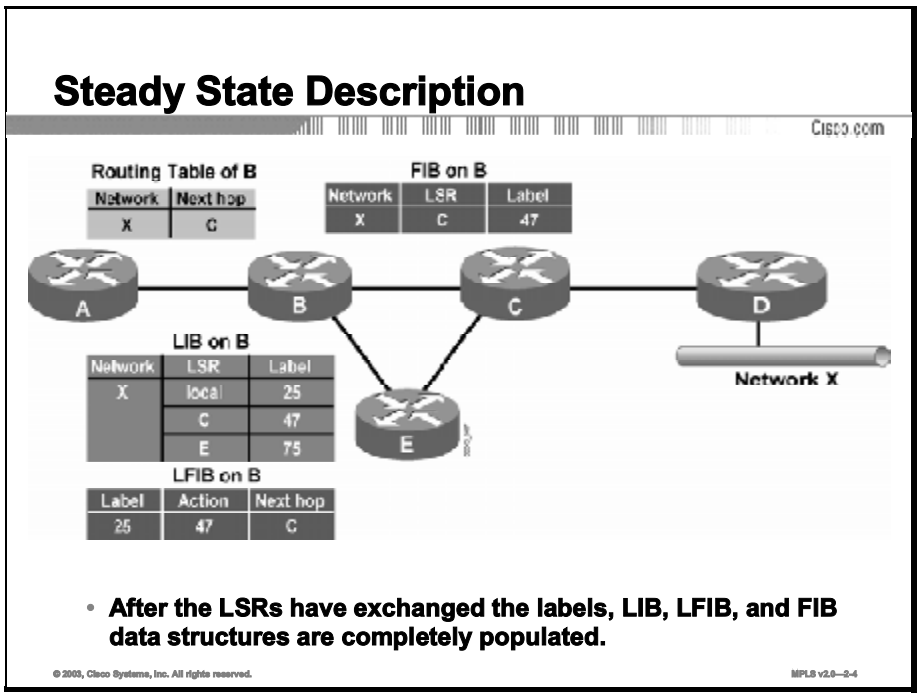

MPLS is fully functional when the IGP and LDP (or TDP) have populated all the tables:

- 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.

Note	The term LDP (Label Distribution Protocol) also applies to TDP (Tag Distribution Protocol).

Link Failure Actions

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

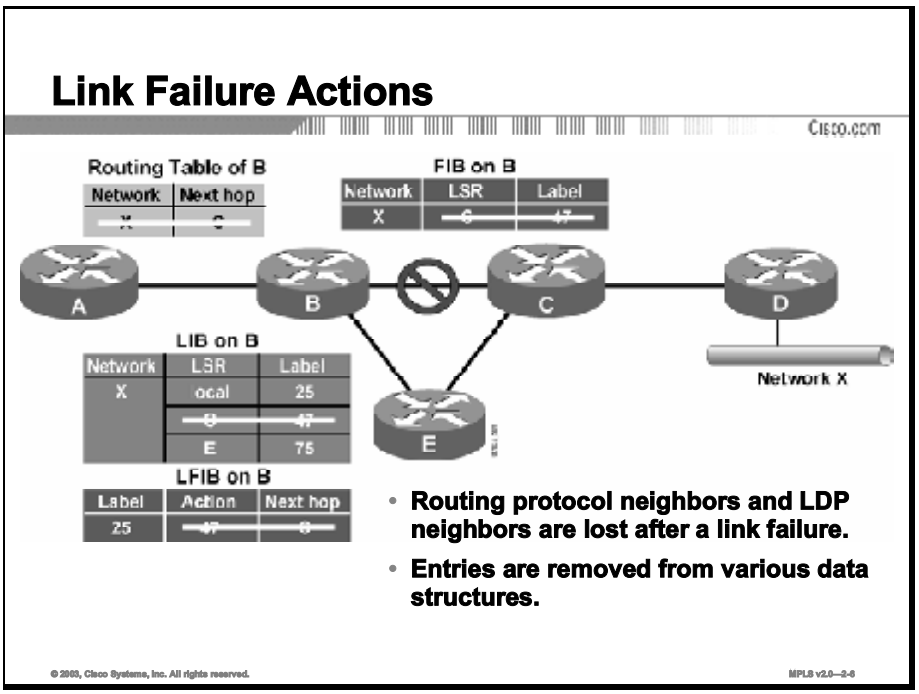

The figure illustrates how a link failure is handled in an MPLS domain:

- 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 they 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:

- Label 47 was learned from router C and is currently unavailable; therefore, because of the failure it 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.

Routing Protocol Convergence

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

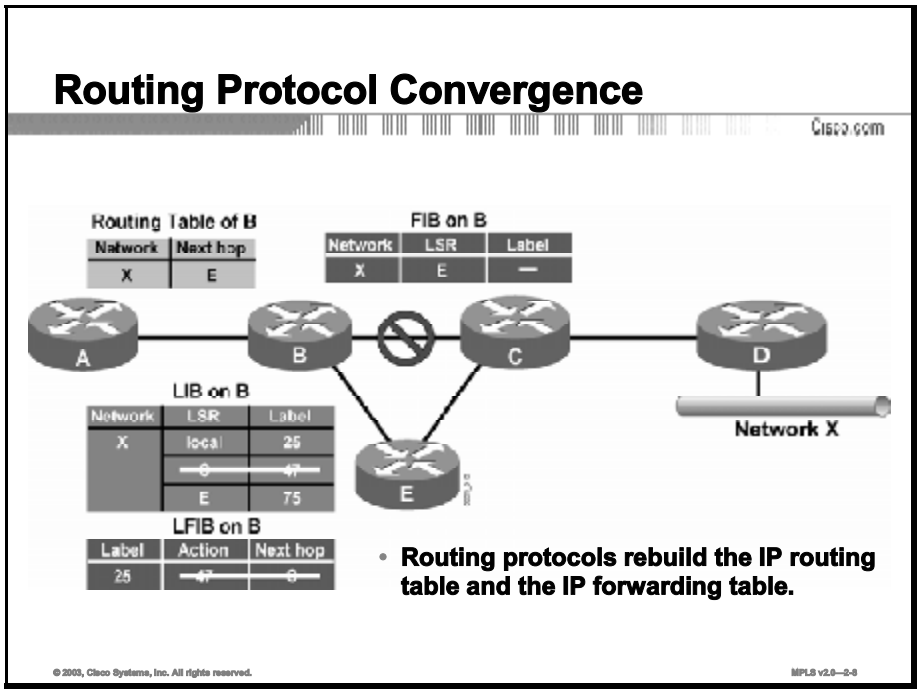

The figure illustrates how two entries are removed, one from the LIB table and one from the LFIB table, when the link between routers B and C fails:

- 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.

MPLS Convergence

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

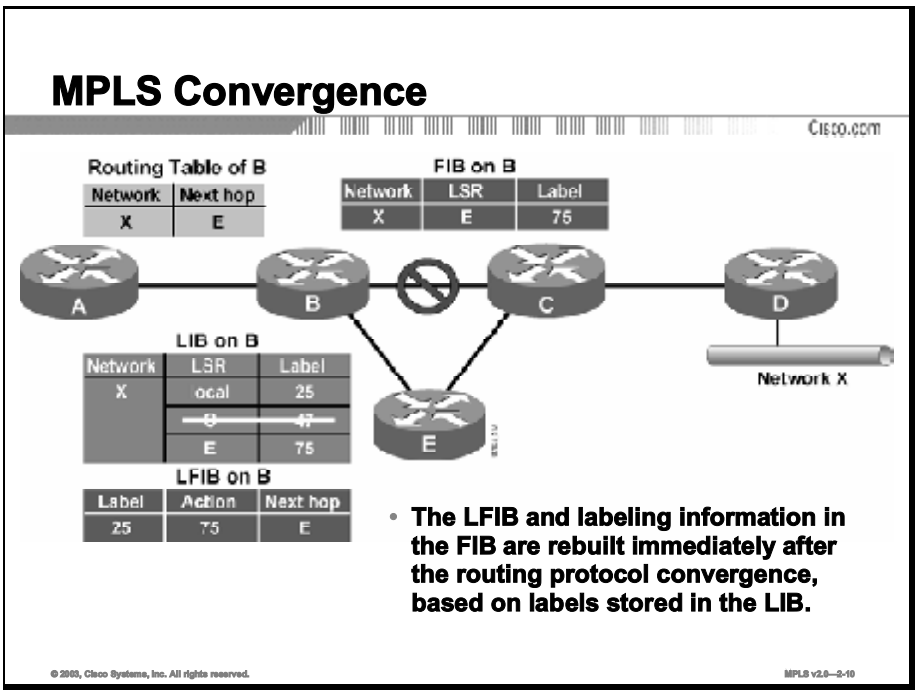

After the IGP determines that there is another path available, a new entry is created in the FIB table.

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.

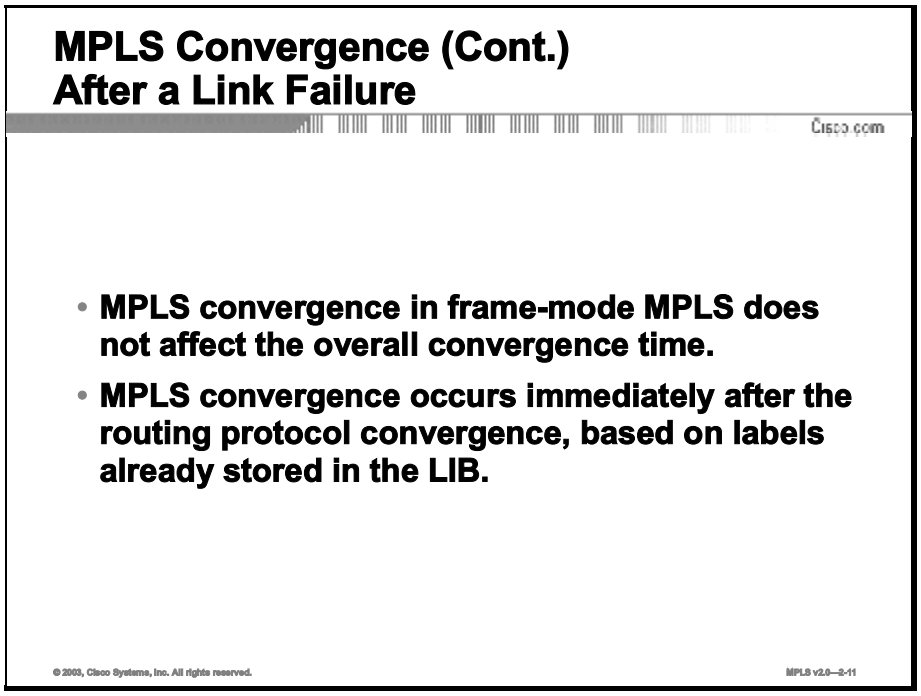

The overall convergence in an MPLS network is not affected by LDP convergence when there is a link failure.

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

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

Link Recovery Actions

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

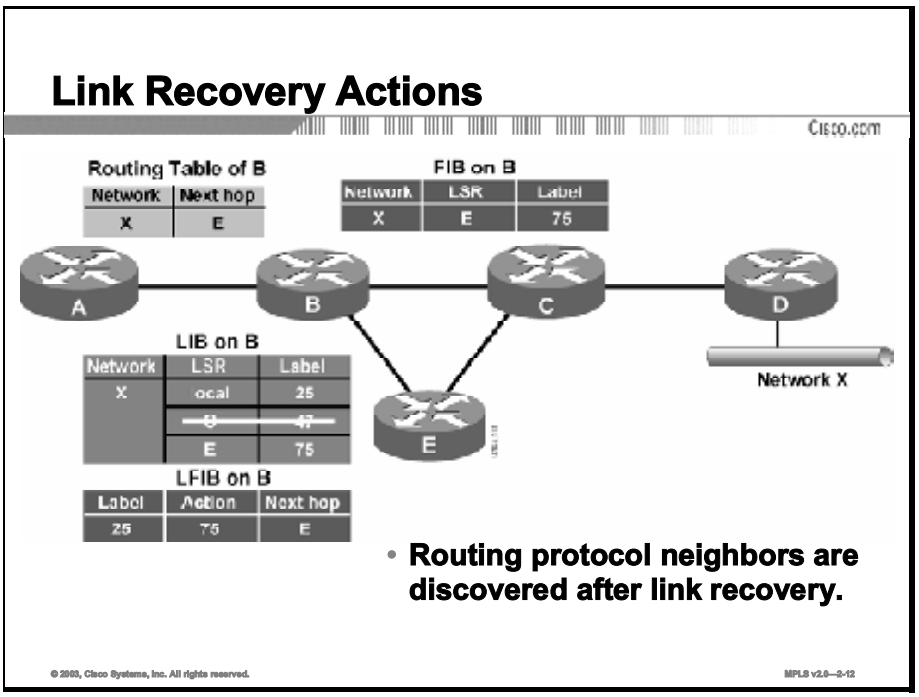

The figure illustrates the state of the routing tables at the time the link between routers B and C becomes available again.

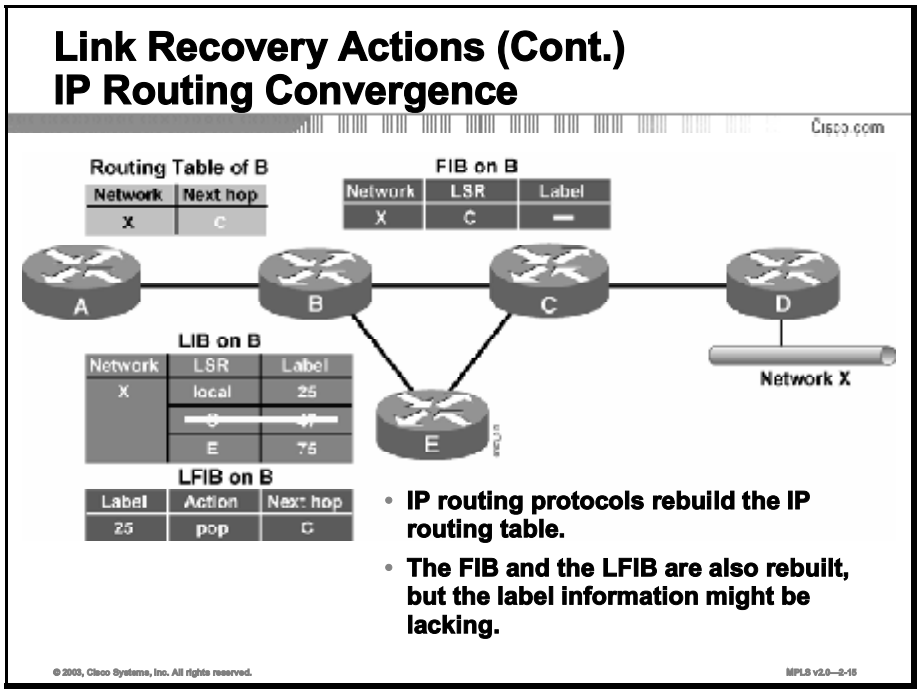

The IGP determines that the link is available again and changes the next-hop address for network X to point to router C. However, when router B also tries to set the next-hop label for network X, it has to wait for the LDP session between routers B and C to be re-established.

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 because the VPN traffic cannot be forwarded before the LDP session is fully operational.

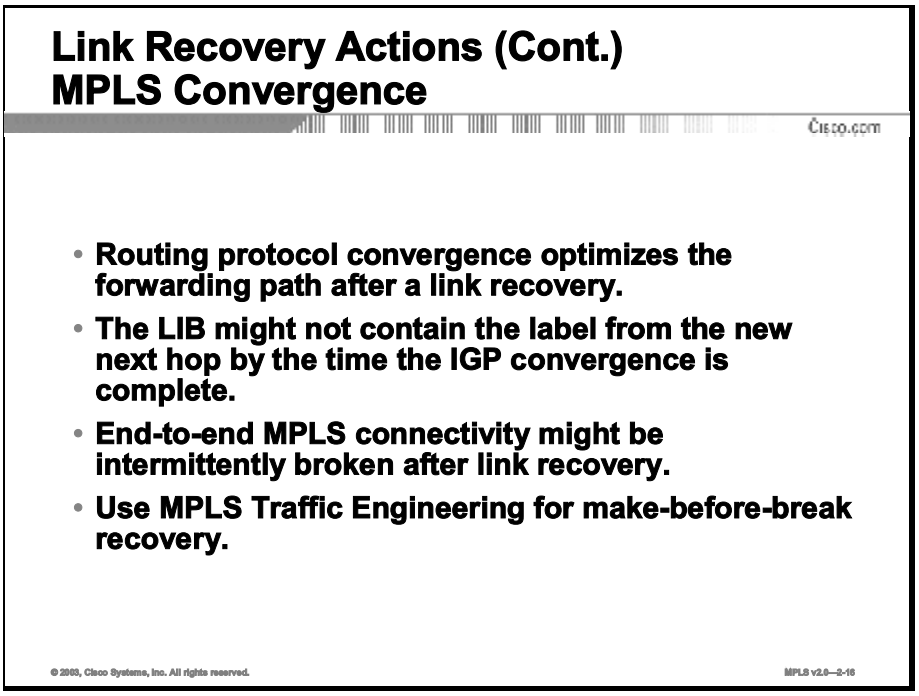

Link recovery requires that an LDP session be established (re-established), which adds to the convergence time of LDP.

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.

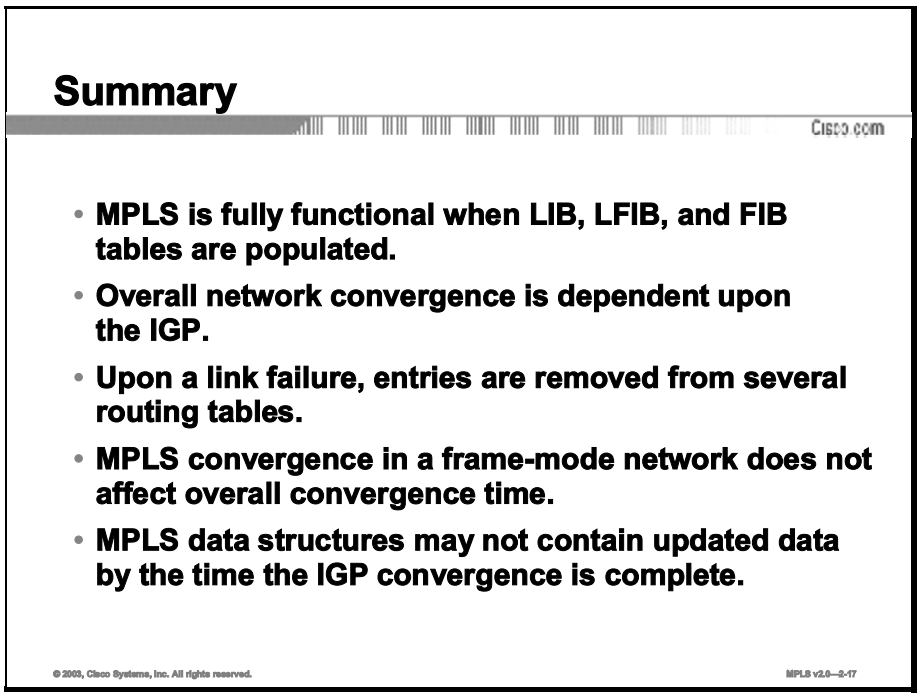

References

For additional information, refer to these resources:

- RFC 3031, “Multiprotocol Label Switching Architecture”
- RFC 3036, “LDP Specification”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following enables routers in a frame-mode MPLS network to store all received labels, even if they are not being used?

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. It keeps only the labels that it will use.

Q2) Which table is NOT used to determine if MPLS is fully functional?

A) LIB

B) LFIB

C) FIB

D) FLIB

Q3) Upon a link failure what three tables are updated to reflect the failed link? (Choose three.)

A) LIB

B) LFIB

C) FIB

D) FLIB

Q4) Which statement best describes how a link failure is handled in an MPLS network?

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.

Q5) Upon a link recovery what three tables are updated to reflect the failed link? (Choose three.)

A) LFIB

B) FLIB

C) FIB

D) LIB

Quiz Answer Key

Q1) C

Relates to: MPLS Convergence

Q2) D

Relates to: Steady State

Q3) A, B, C

Relates to: Routing Protocol Convergence

Q4) B

Relates to: Link Failure Actions

Q5) A, C, D

Relates to: Link Recovery Actions

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).

Relevance

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

This lesson describes typical label distribution over LC-ATM (label controlled-ATM) interfaces and VC merge.

Upon completing this lesson, you will be able to:

- Identify issues within cell-mode MPLS networks
- Describe how the IP routing table is built in a cell-mode MPLS network
- Describe how the IP forwarding table is built 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
- Describe cell-mode loop detection in MPLS networks
- Identify cell interleave issues in cell-mode MPLS networks
- Describe the benefits and drawbacks of VC merge
- Describe the benefits and drawbacks of per-interface label allocation

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Convergence in Frame-Mode MPLS” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Cell-Mode MPLS Network Issues
- Building the IP Routing Table
- Building the IP Forwarding Table
- Requesting a Label
- Allocating a Label
- Cell Interleave Issues
- VC Merge
- Loop Detection in Cell-Mode MPLS Networks
- Per-Interface Label Allocation
- Summary
- Quiz

Cell-Mode MPLS Network Issues

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

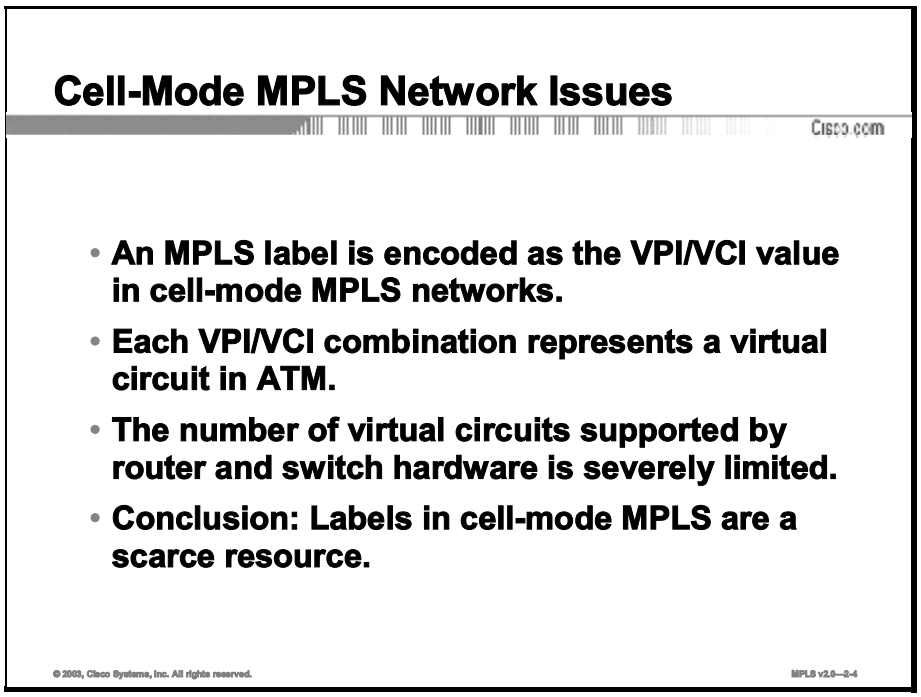

Cell-mode MPLS is significantly different from frame-mode MPLS because of some ATM-specific requirements:

- ATM uses cells and 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.

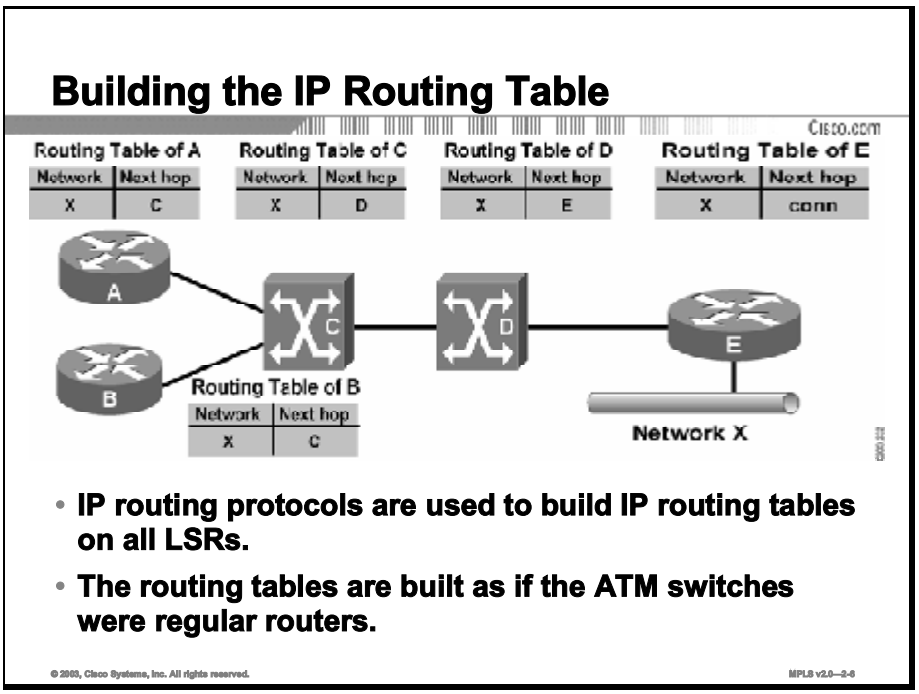

The figure shows how IP- and MPLS-aware ATM switches exchange IP routing information with routers.

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.

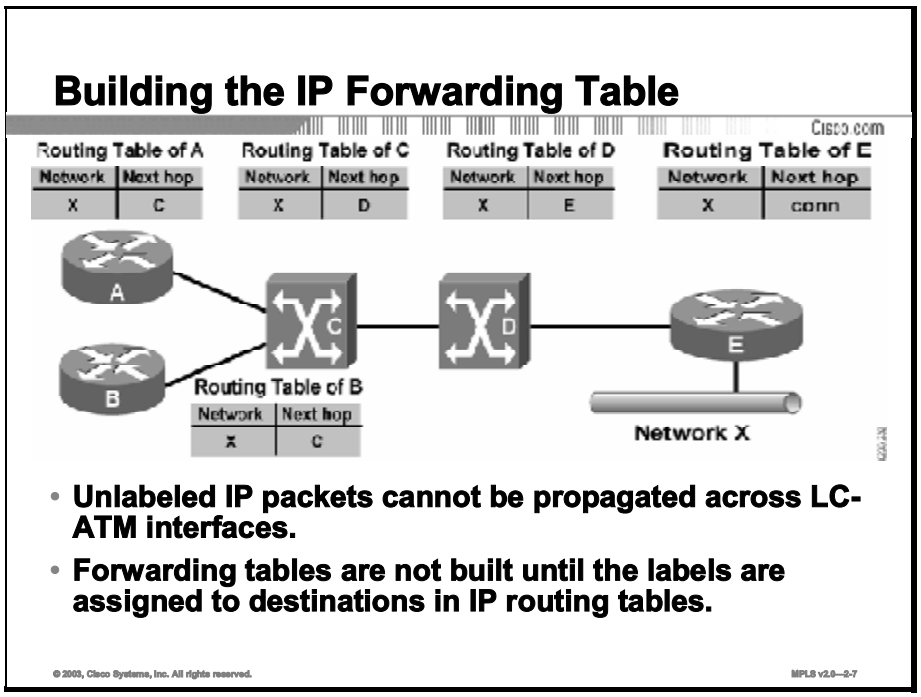

Because ATM switches cannot forward IP packets, labels cannot be asynchronously assigned and distributed.

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.

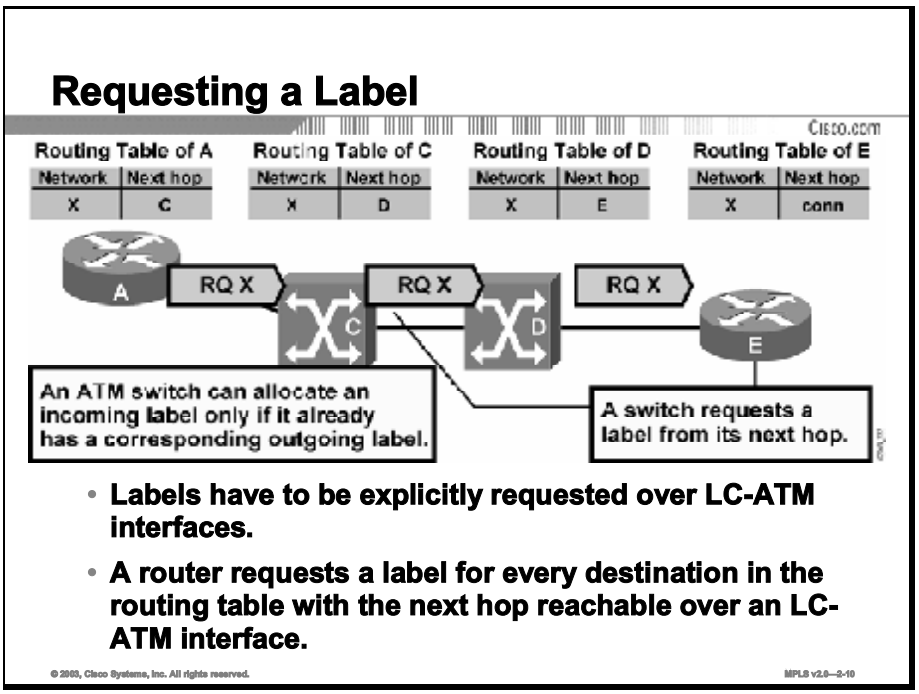

In the example here, a request is sent from router A to the ATM switch C. Because the ATM switch cannot perform IP lookups, the switch is not allowed to reply with the local label unless it already has the next-hop label. If switch C does not have the next-hop label, it must forward the request to the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

If switch D does not have the next-hop label, it 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.

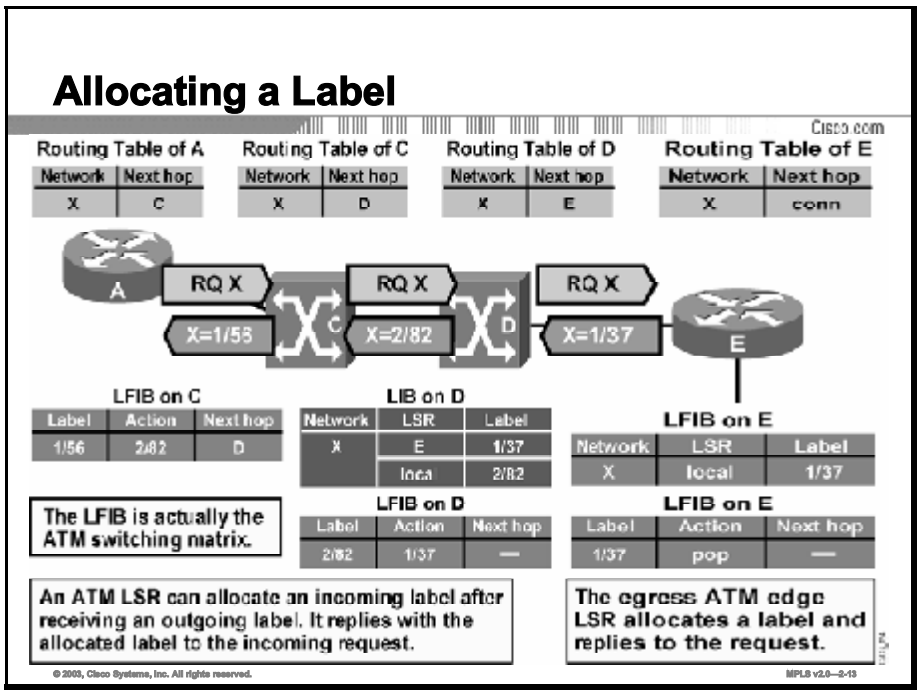

In this figure router E replies with its local label 1/37. The ATM switch D can now generate and use its local label 2/82. Switch C receives the next-hop label from switch D and forwards its own local label 1/56 to router A.

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.

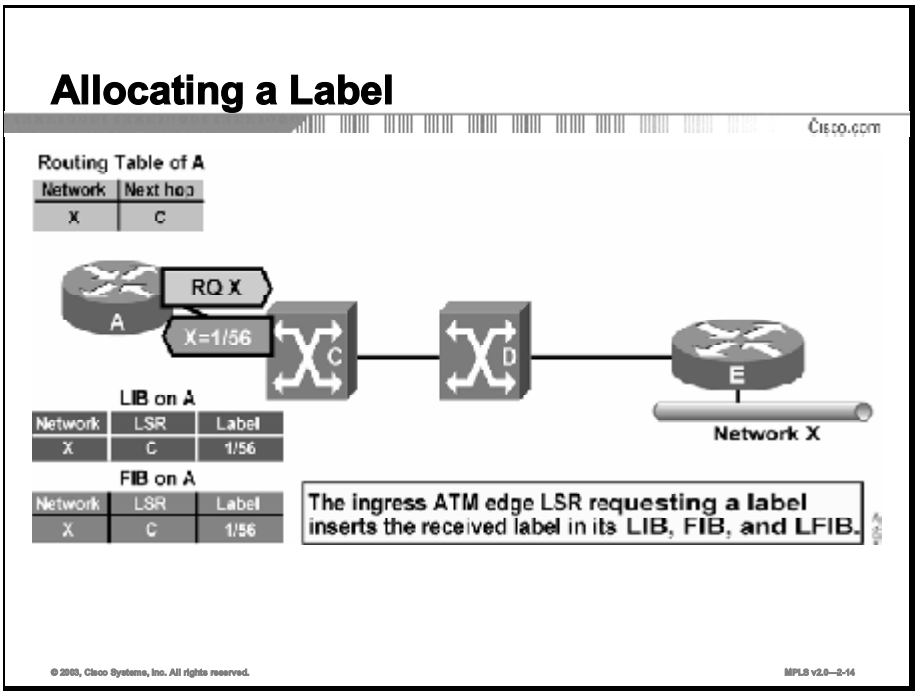

Processing of LDP replies on router A (also an ATM edge LSR) is similar to processing in frame-mode MPLS; the received label is stored in the LIB, FIB, and LFIB tables.

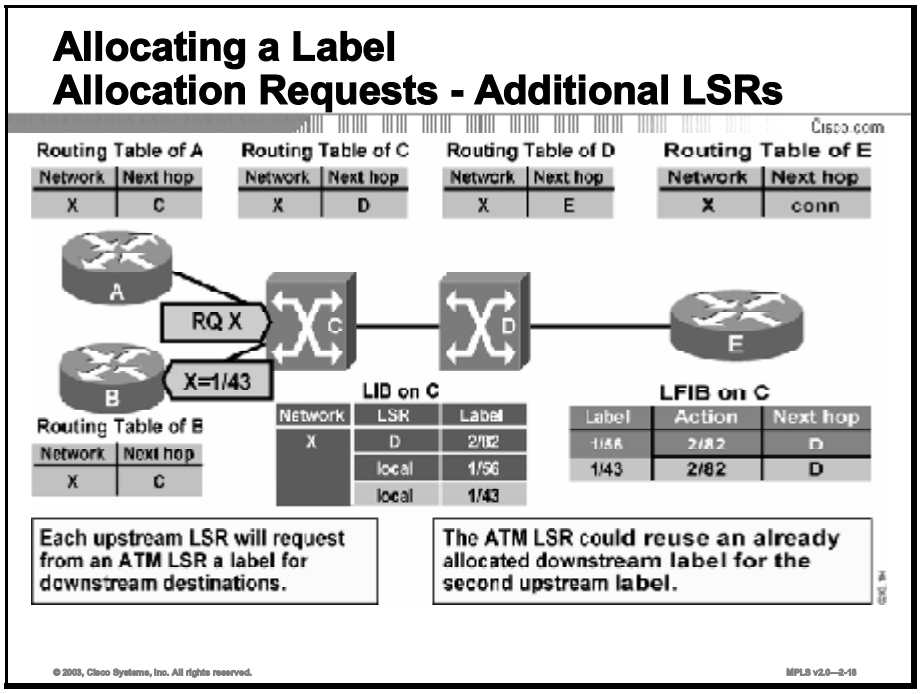

The figure shows how another router, router B, requests a label for the same destination that router A has previously requested. The ATM switch C already has a next-hop label for network X and therefore can immediately reply to router B.

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

Cell Interleave Issues

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

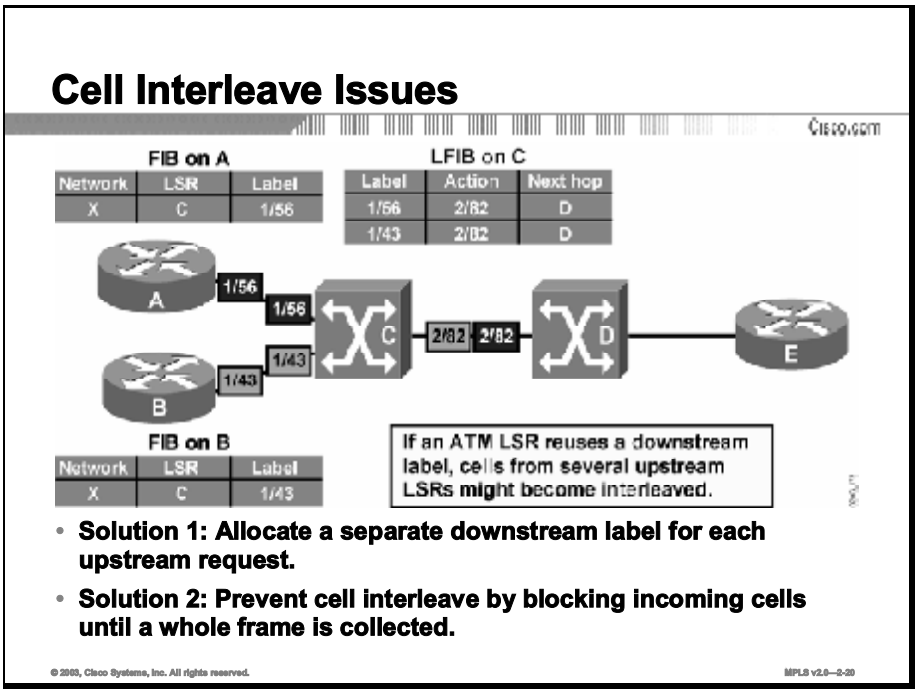

Analysis of the previous two figures reveals that an unusual situation has developed. Two virtual circuits from routers A and B (1/56 and 1/43) merge into one (2/82).

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.

There are 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.

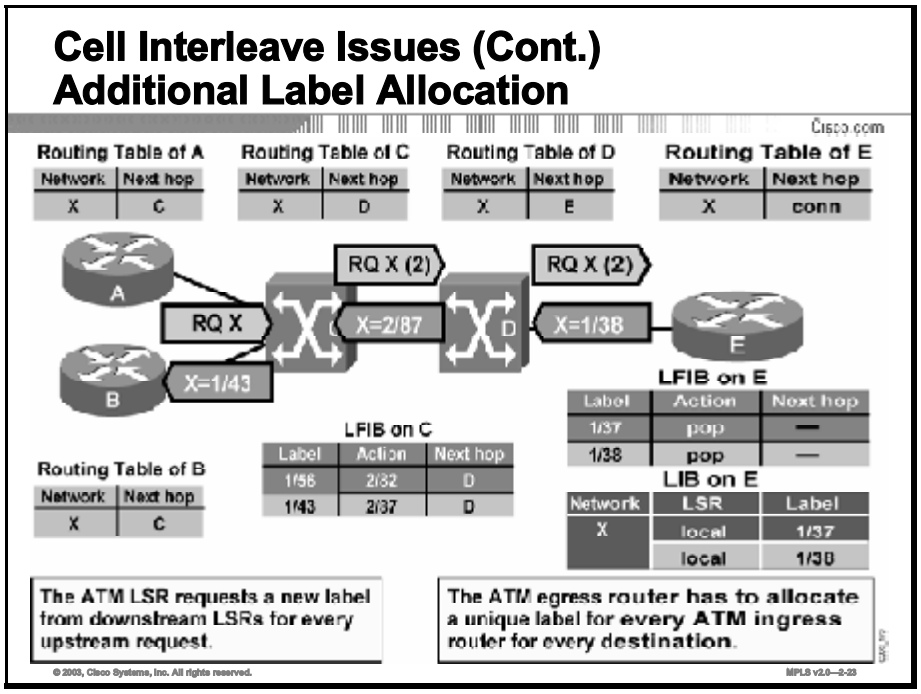

The figure illustrates the first option, where an additional LSP tunnel is created for the same destination network X for every upstream ATM edge LSR.

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

VC Merge

This topic describes what VC merge is, and its benefits and drawbacks.

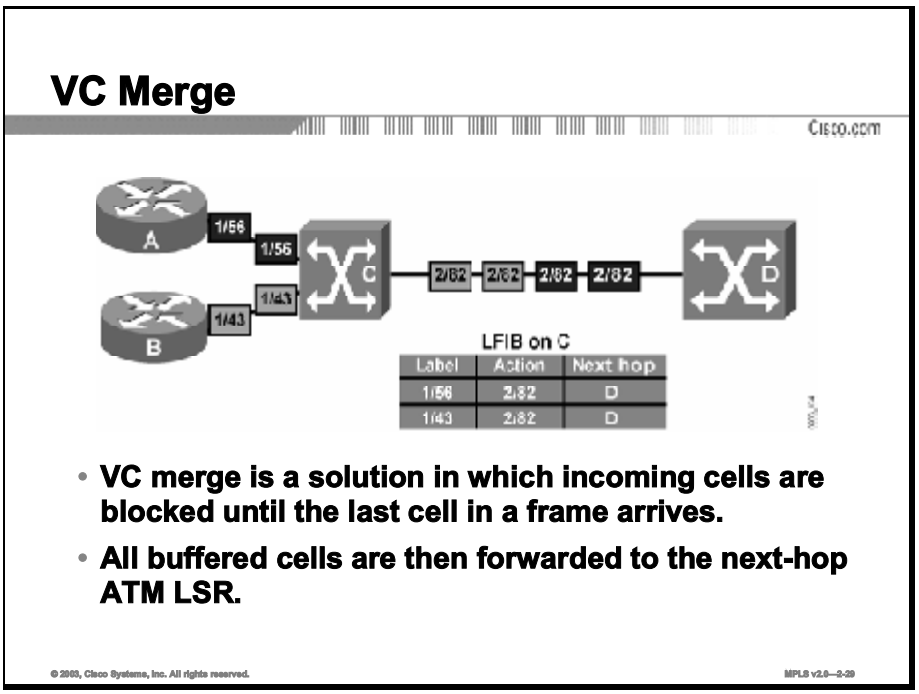

The figure illustrates the second option, where the ATM switch C buffers cells coming from router B until the last cell of the packet coming from router A is forwarded.

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

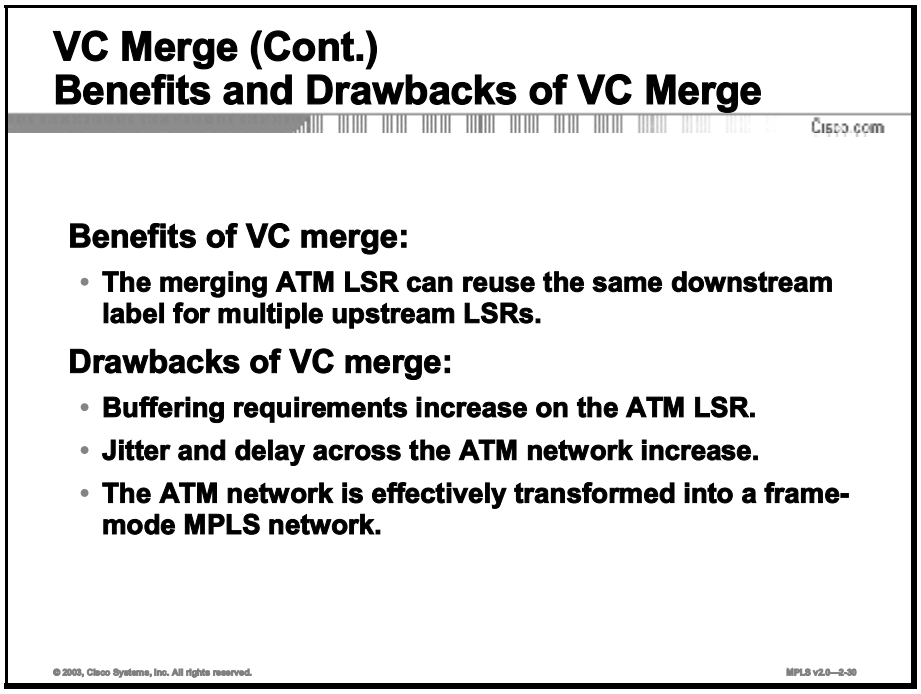

The major benefit of VC merge is that it minimizes the number of labels (VPI/VCI values) needed in the ATM part of the network. As identified in the first topic in this lesson, labels are a scarce resource in cell-mode MPLS networks.

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.

Loop Detection in Cell-Mode MPLS Networks

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

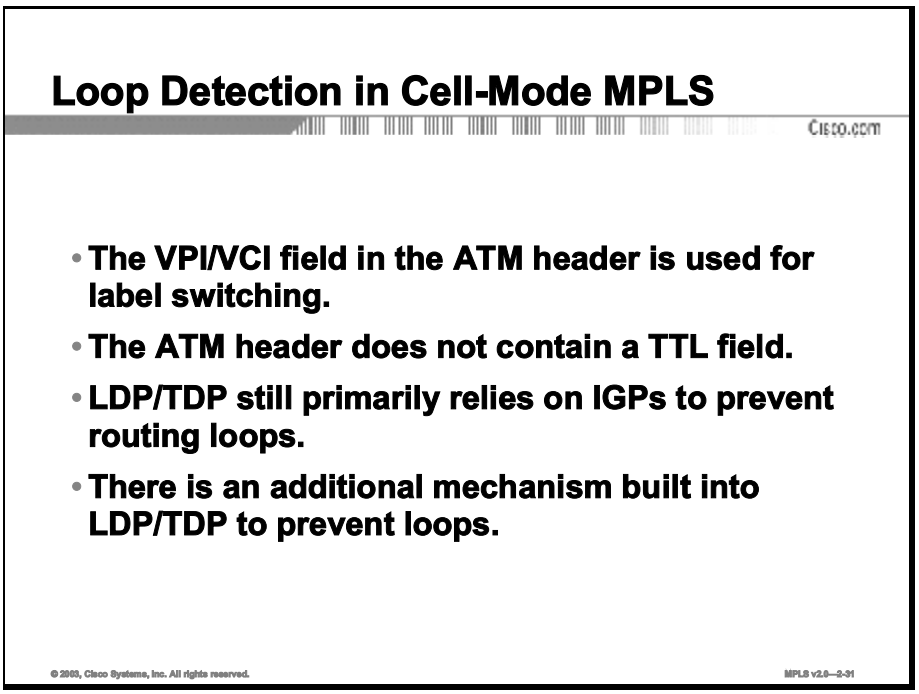

Cell-mode MPLS uses the VPI/VCI fields in the ATM header to encode labels. These two fields do not include a TTL field. Therefore, cell-mode MPLS must use other ways of preventing routing loops.

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.

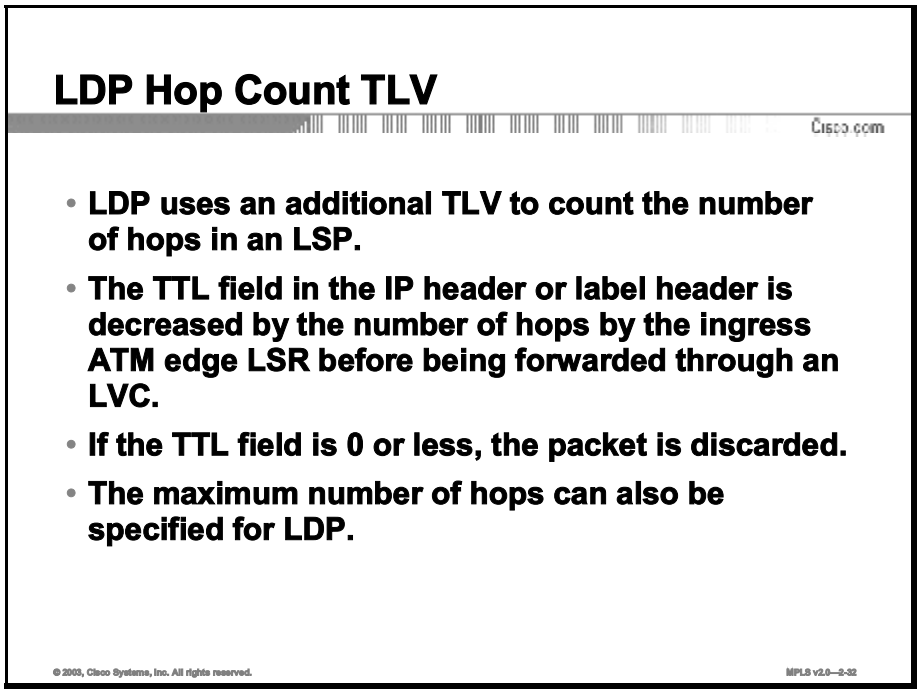

LDP uses a hop-count TLV (type, length, value) attribute to count hops in the ATM part of the MPLS domain.

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.

LDP Hop Count Example

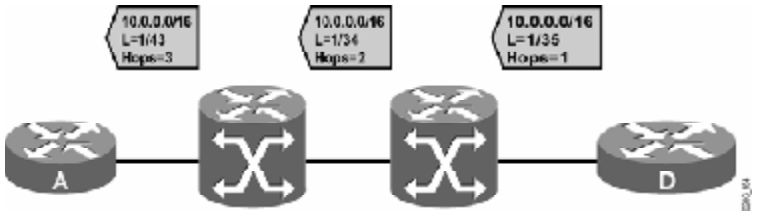

- **LSR A discovers the length of the LSP across the ATM domain to LSR D through LDP.**

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

The next example shows how traceroute is affected by this functionality.

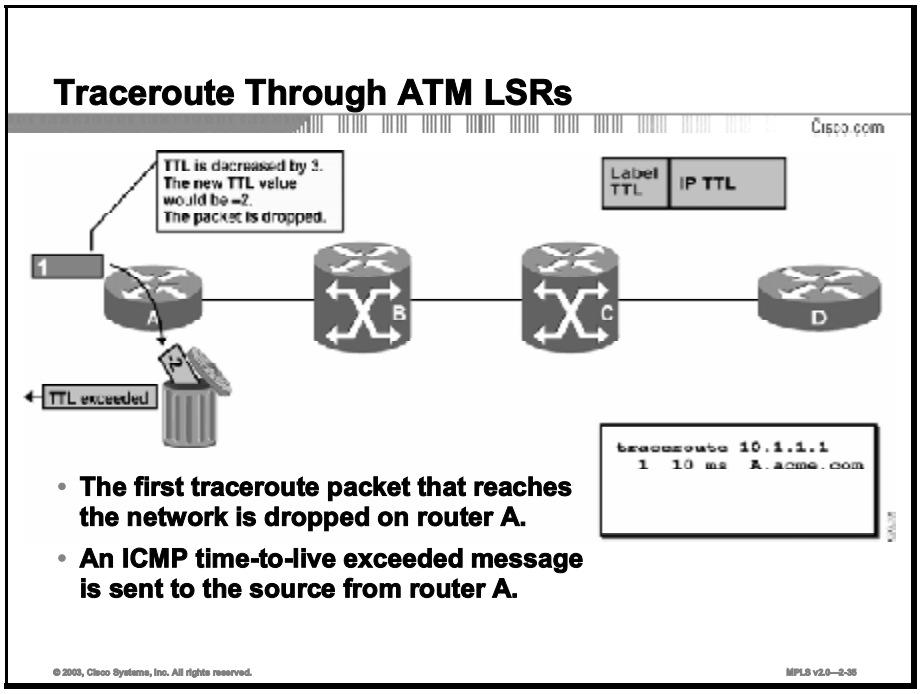

The following figures illustrate how traceroute works across an IP-aware ATM network that is not capable of using the TTL field and generating ICMP replies.

The example here 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.

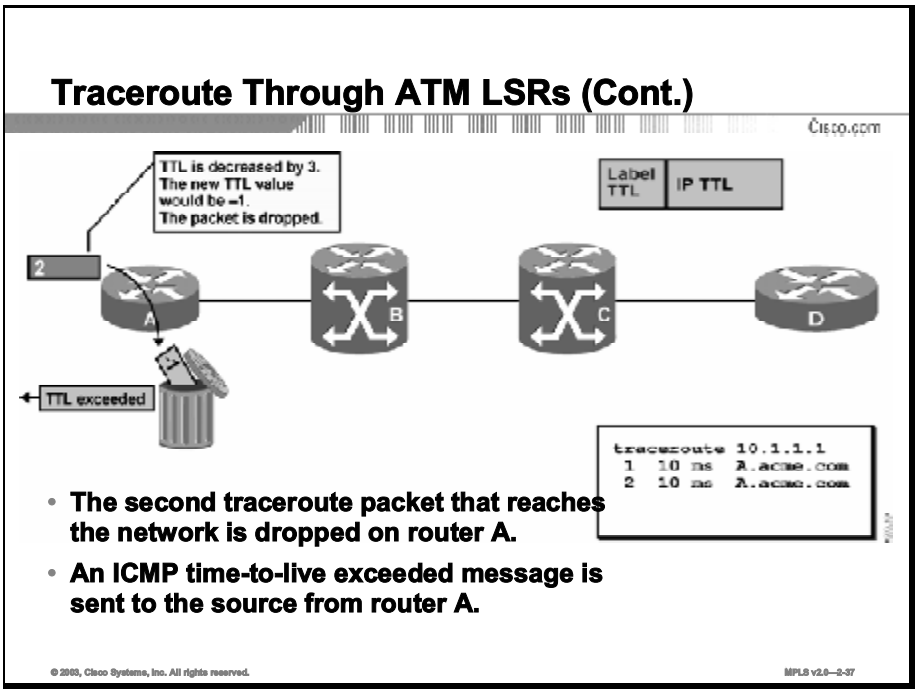

The second packet is also dropped, and another ICMP reply is sent from router A on behalf of ATM switch B, which cannot identify the TTL field and send ICMP replies itself.

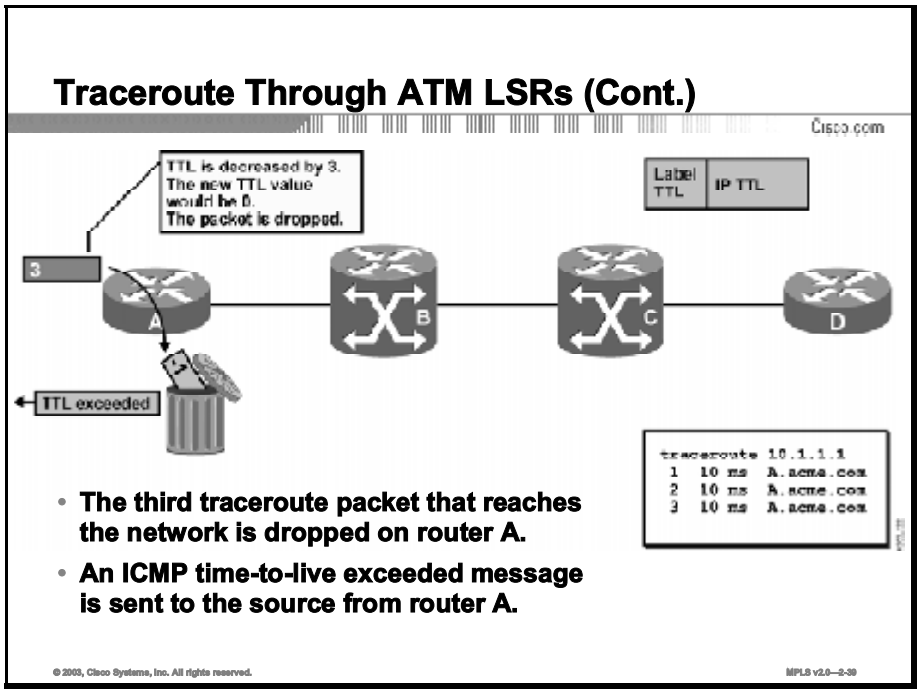

The third packet is also dropped, and the third ICMP reply is sent from router A on behalf of the ATM switch C.

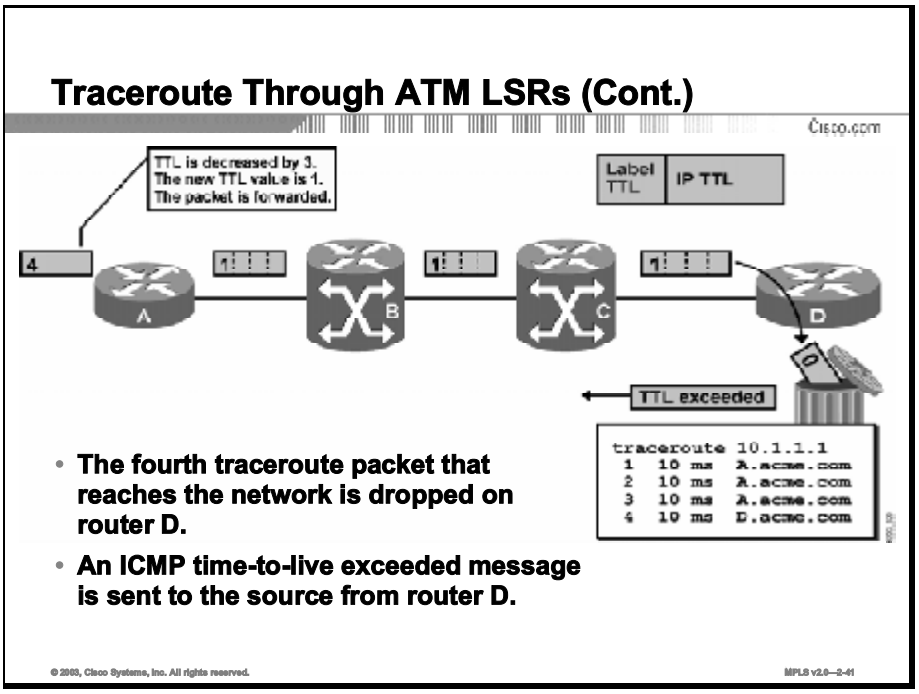

The fourth packet can reach the other edge ATM LSR (a router), which is capable of identifying the TTL field and sending ICMP replies.

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.

Per-Interface Label Allocation

This topic describes per-interface label allocation.

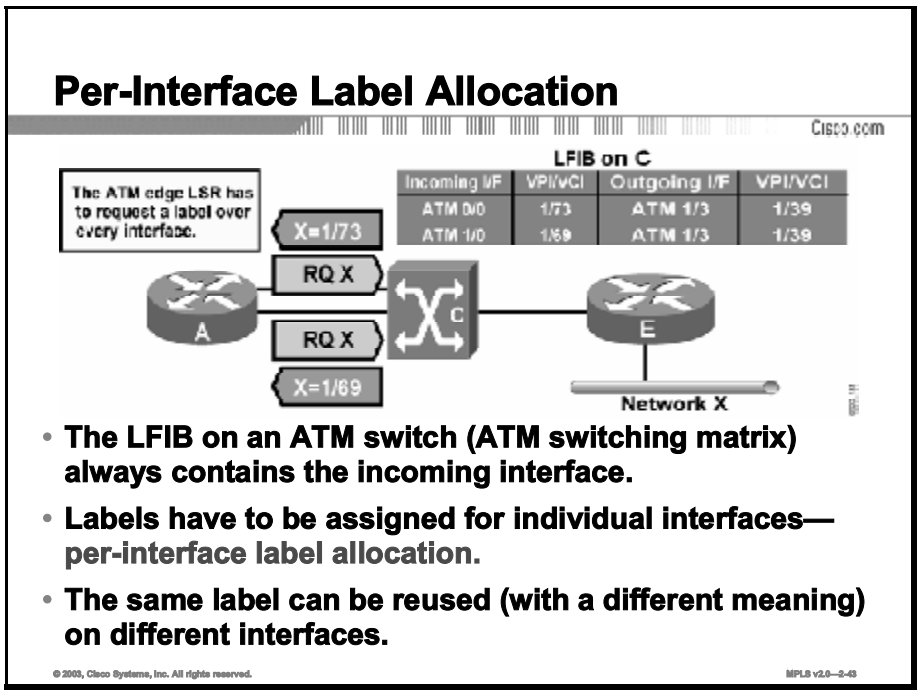

Cell-mode MPLS defaults to using per-interface label space because ATM switches support per-interface VPI/VCI values to encode labels.

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.

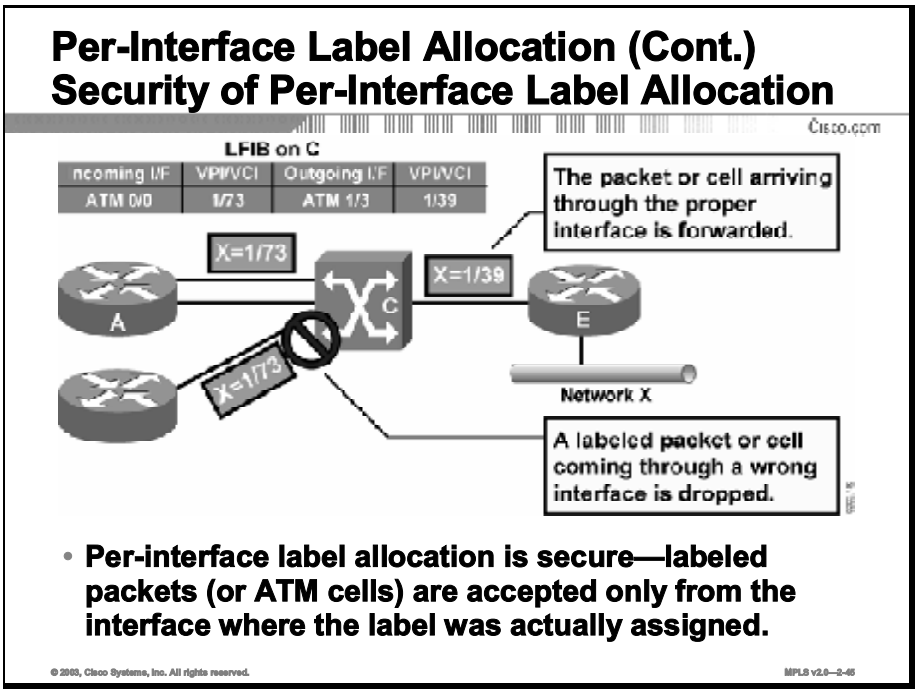

One benefit of per-interface label space is that it prevents label spoofing. In the figure, for example, the bottom router has tried to send a cell with a label that was advertised only to router A. The switch has failed to forward the cell because it came in through the wrong interface.

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.

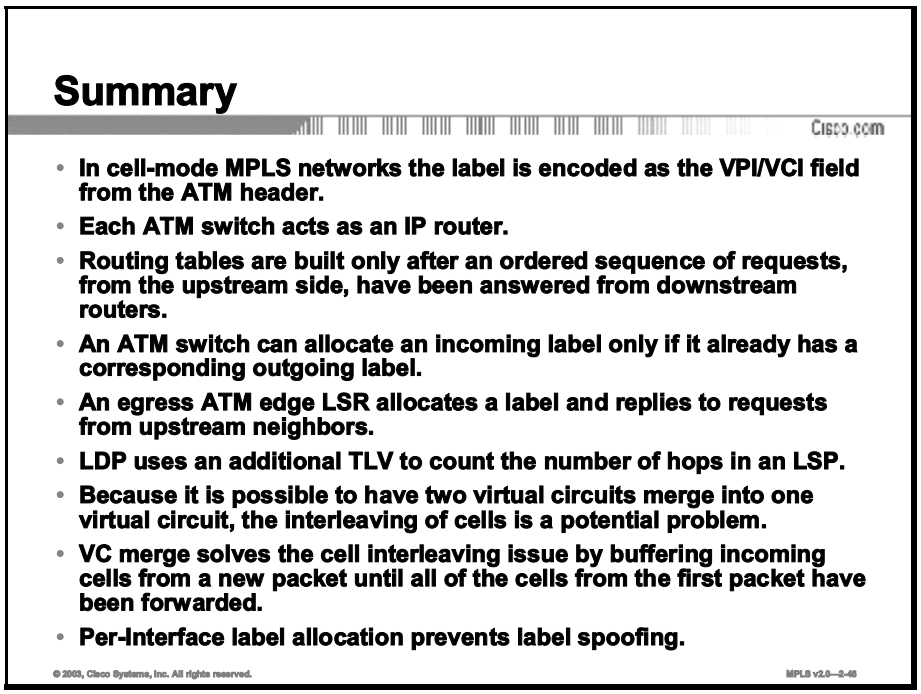

References

For additional information, refer to these resources:

- RFC 3031, “Multiprotocol Label Switching Architecture”
- RFC 3036, “LDP Specification”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are two possible solutions to the interleaving of cells in cell-mode MPLS? (Choose two.)

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.

Q2) In cell-mode MPLS networks, where are labels inserted?

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.

Q3) With regard to VC merge, which statement is NOT true?

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.

Q4) Which of the following pertains to the IP routing table?

A) NOT built on ATM LSRs

B) built on the data plane of each ATM switch

C) built on the control plane of each ATM switch

D) built on the forwarding plane of each ATM switch

Q5) Which of the following pertains to the IP forwarding table?

A) built as in a frame-mode MPLS network

B) built only after the label requests have been answered (with labels) from upstream LSRs

C) 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.

Q6) Which of the following statements is NOT true?

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.

Q7) An ATM switch will respond to a request for a label in what situation?

A) It will respond when it knows the next-hop label.

B) It will always reply to downstream label requests.

C) It will always reply to upstream label requests.

D) ATM switches do not use MPLS labels.

Q8) ATM switches perform which of the following?

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.

Q9) Which of the following is used in cell-mode loop detection?

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

Quiz Answer Key

Q1) A, C

Relates to: Cell Interleave Issues

Q2) B

Relates to: Cell-Mode MPLS Network Issues

Q3) C

Relates to: VC Merge

Q4) C

Relates to: Building the IP Routing Table

Q5) B

Relates to: Building the IP Forwarding Table

Q6) C

Relates to: Per-Interface Label Allocation

Q7) A

Relates to: Requesting a Label

Q8) B

Relates to: Allocating a Label

Q9) C

Relates to: Loop Detection in Cell-Mode MPLS Networks

MPLS Label Allocation, Distribution, and Retention Modes

Overview

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

Relevance

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

Objectives

This lesson describes the MPLS label allocation, distribution, and retention modes used in Cisco MPLS networks.

Upon completing this lesson, you will be able to:

- Describe the parameters used in Cisco MPLS label distribution and allocation
- Describe the two types of label space used in Cisco MPLS label distribution and allocation
- Describe the two types of label distribution used in Cisco MPLS
- Describe the two types of label allocation used in Cisco MPLS
- Describe the two types of label retention used in Cisco MPLS
- Describe the standard parameter sets in Cisco IOS platform MPLS implementation

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Typical Label Distribution over LC-ATM Interfaces and VC Merge” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Label Distribution Parameters
- Label Space
- Label Distribution
- Label Allocation
- Label Retention
- Standard Parameter Sets in Cisco IOS Platform MPLS Implementation
- Summary
- Quiz

Label Distribution Parameters

This topic describes the options in label allocation, propagation, and retention.

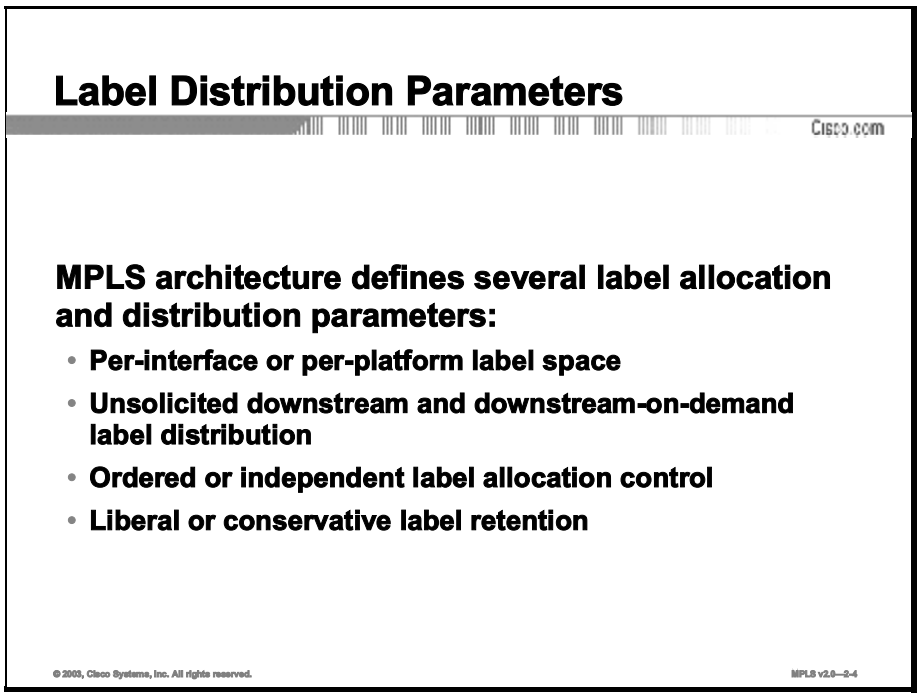

The two label space options are:

- 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 them 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:

- 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 they 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:

- 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 keeps only labels that it previously requested. This mode is called conservative label retention mode.

Label Space

This topic describes how labels are generated, either on a per-interface or per-platform basis.

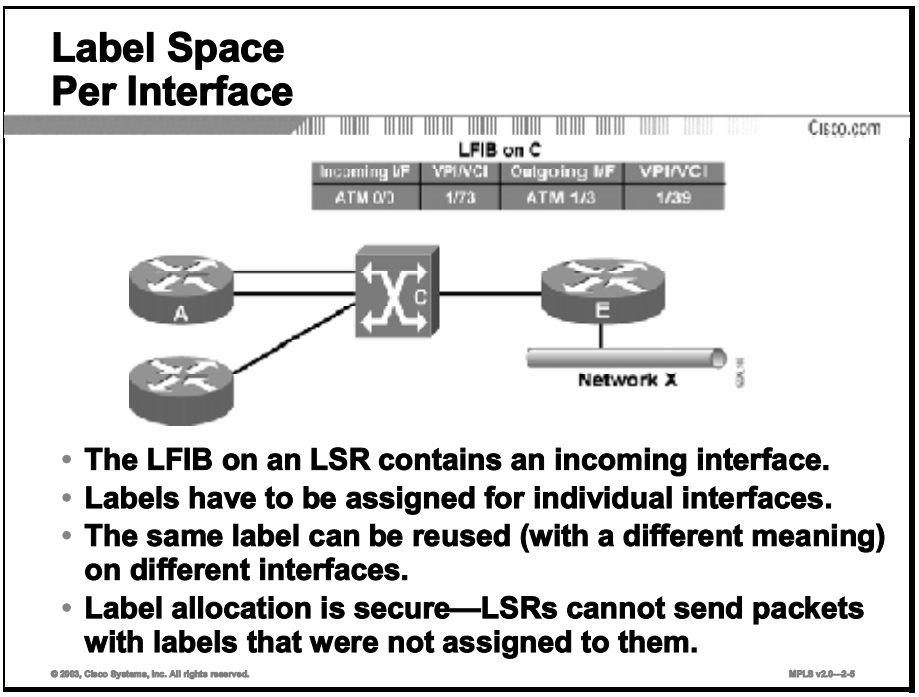

The LFIB table used with cell-mode MPLS maps a local label bound to an input interface to a next-hop label pointing to the outgoing interface. The label assigned to an input interface can be reused on another interface, and it can have a different meaning (assigned to a different destination).

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 (Cont.)
Per Platform

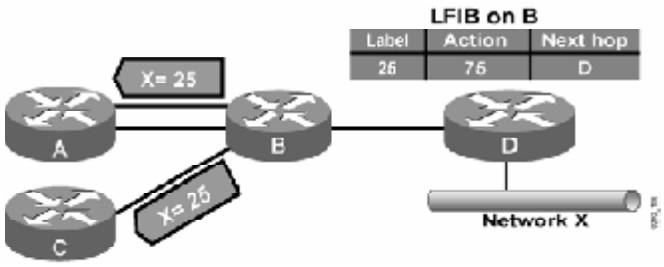

- **The LFIB on an LSR does not contain an incoming interface.**
- **The same label can be used on any interface and is announced to all adjacent LSRs.**
- **The label is announced to adjacent LSRs only once and can be used on any link.**
- **Per-platform label space is less secure than per-interface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

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 them. 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.

Label Distribution

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

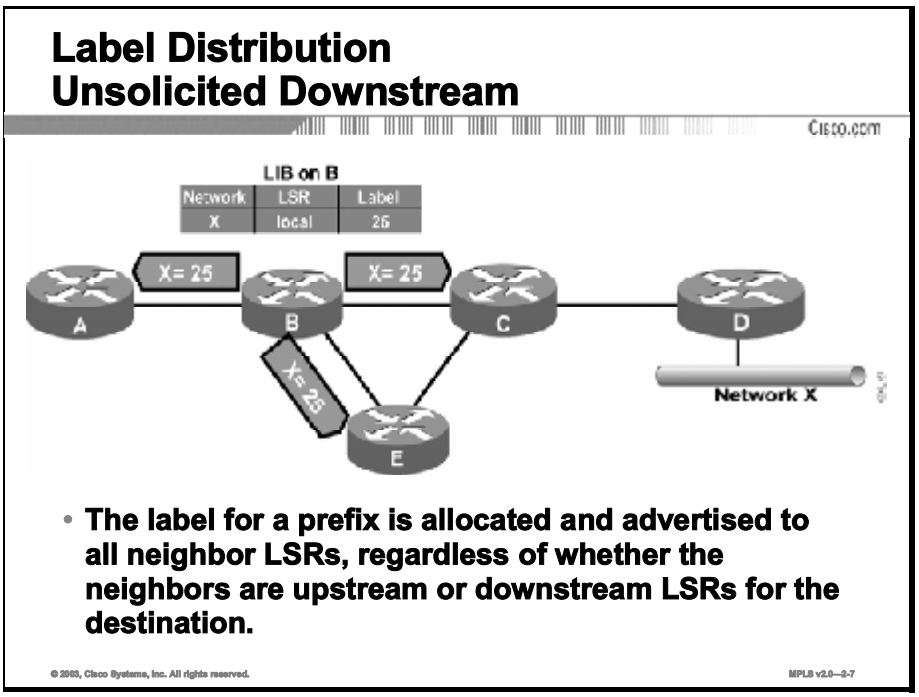

Unsolicited downstream distribution of labels is a method where each router independently assigns a label to each destination IP prefix in its routing table. This mapping is stored in the LIB table, which sends it to all LDP peers. There is no control mechanism to govern the propagation of labels in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

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 it in the LIB table
- ignore it

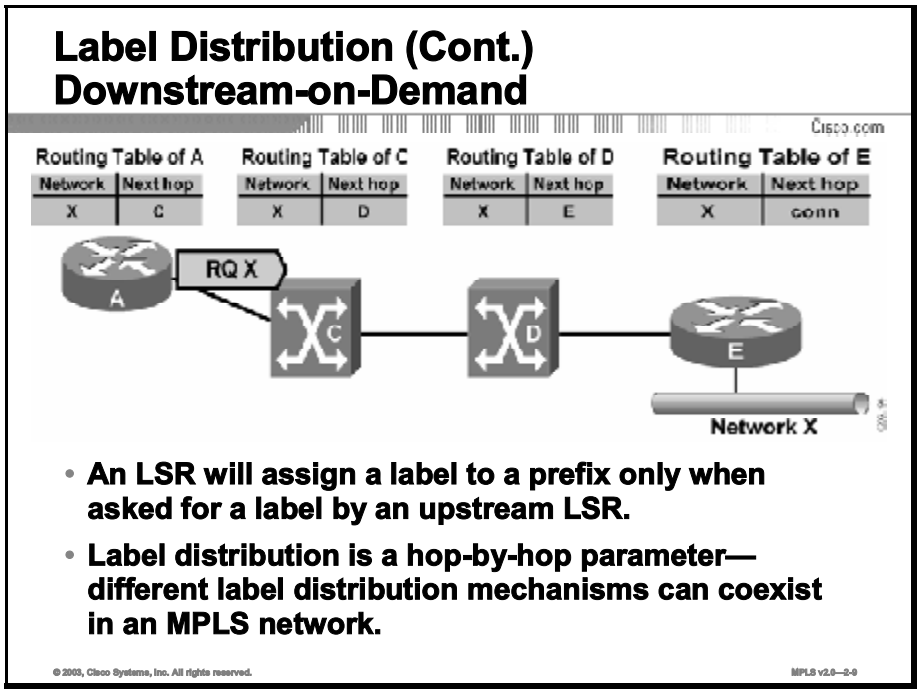

Downstream-on-demand distribution of labels requires each LSR to specifically request a label from its downstream neighbor. The figure shows how router A requests a next-hop label from its downstream LDP peer.

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.

Label Allocation

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

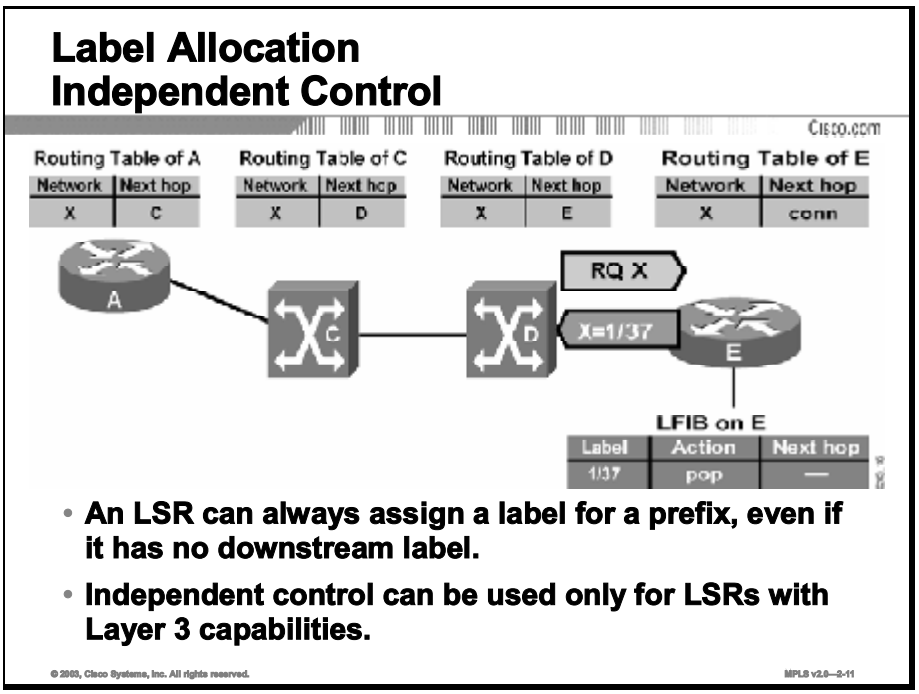

Independent control mode is usually combined with unsolicited downstream propagation of labels, where labels can be created and propagated independently of any other LSR. When independent control mode is used, an LSR might be faced with an incoming labeled packet where there is no corresponding outgoing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

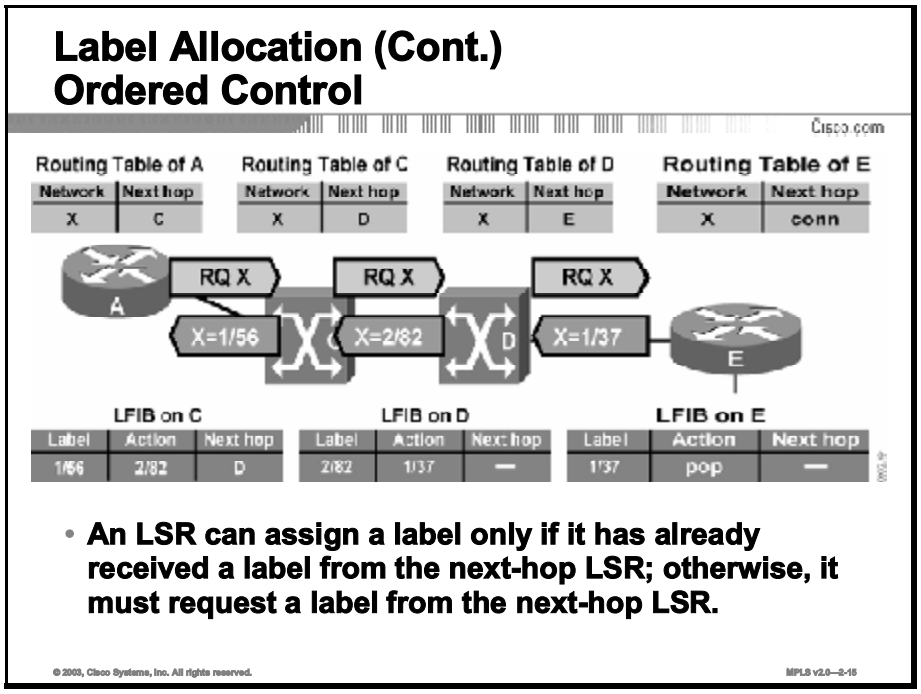

Ordered control mode is usually combined with downstream-on-demand propagation of labels, where a local label can be assigned and propagated only if a next-hop label is available. This requirement results in an ordered sequence of downstream requests until an LSR is found that already has a next-hop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

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.

Label Retention

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

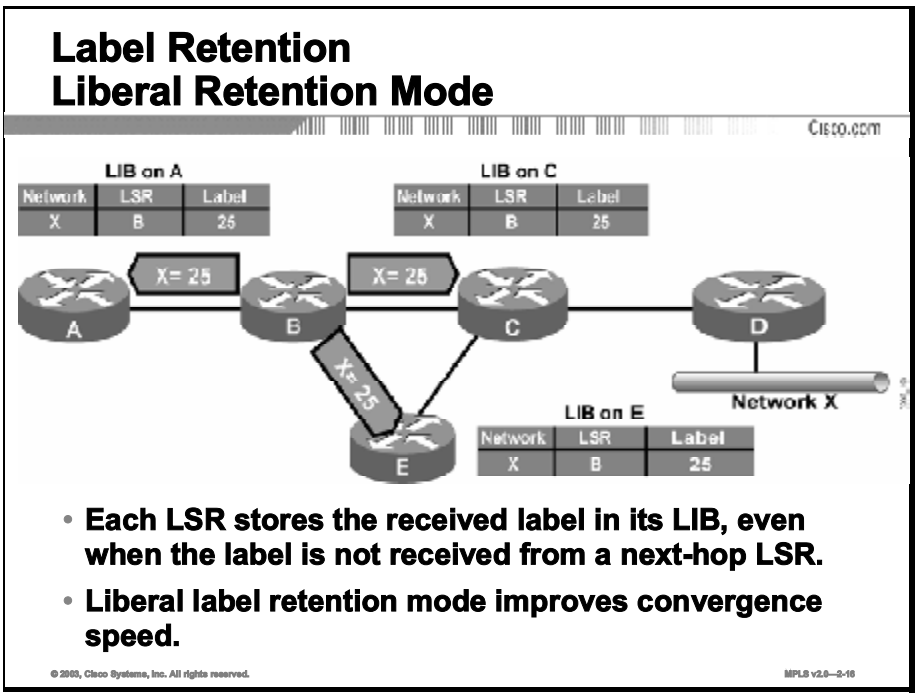

Liberal label retention mode dictates that each LSR keep all labels received from LDP peers even if they are not the downstream peers for network X.

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.

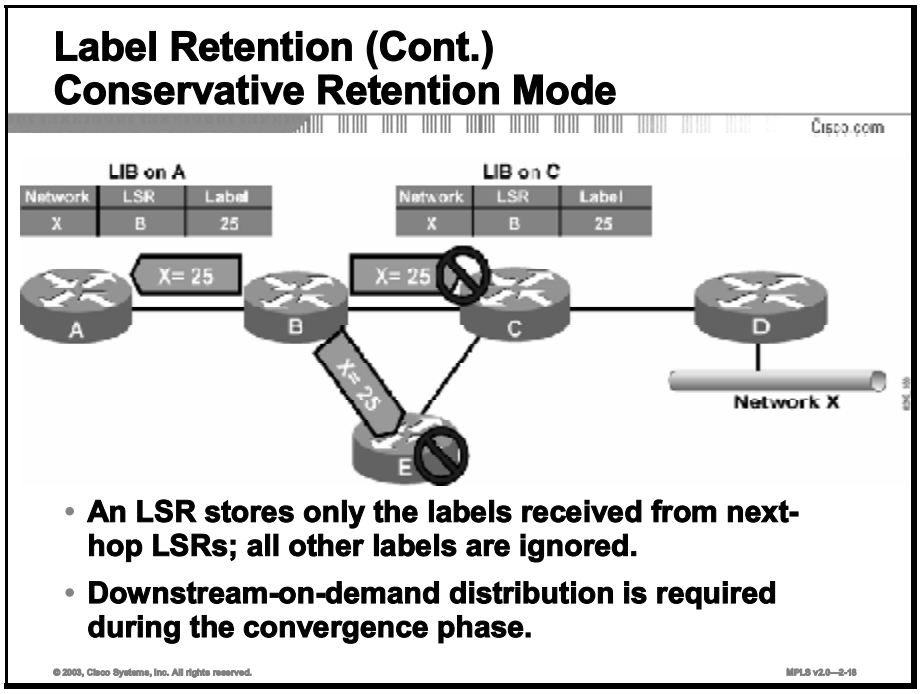

Conservative label retention mode keeps only labels that can immediately be used.

The figure illustrates how routers C and E do not consider router B to be the next hop for network X and therefore drop the labels received from router B.

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

Standard Parameter Sets in Cisco IOS Platform MPLS Implementation

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

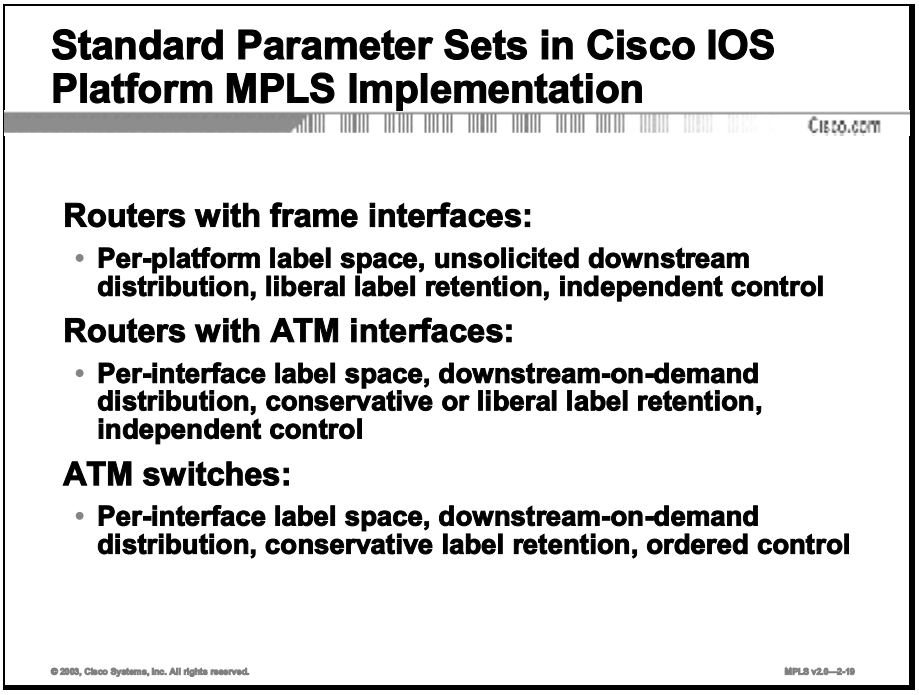

The following default operation applies to routers using frame-mode MPLS (LSRs):

- **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:** 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:** Implicitly achieved by using the downstream-on-demand propagation of labels; no label is received unless it is requested.
- **Ordered control mode:** Used in combination with downstream-on-demand propagation of labels to make sure every ATM LSR has all 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 those of ATM switches except that they use independent control mode, because they are the endpoints of the virtual circuits.

Summary

This topic summarizes the key points discussed in this lesson.

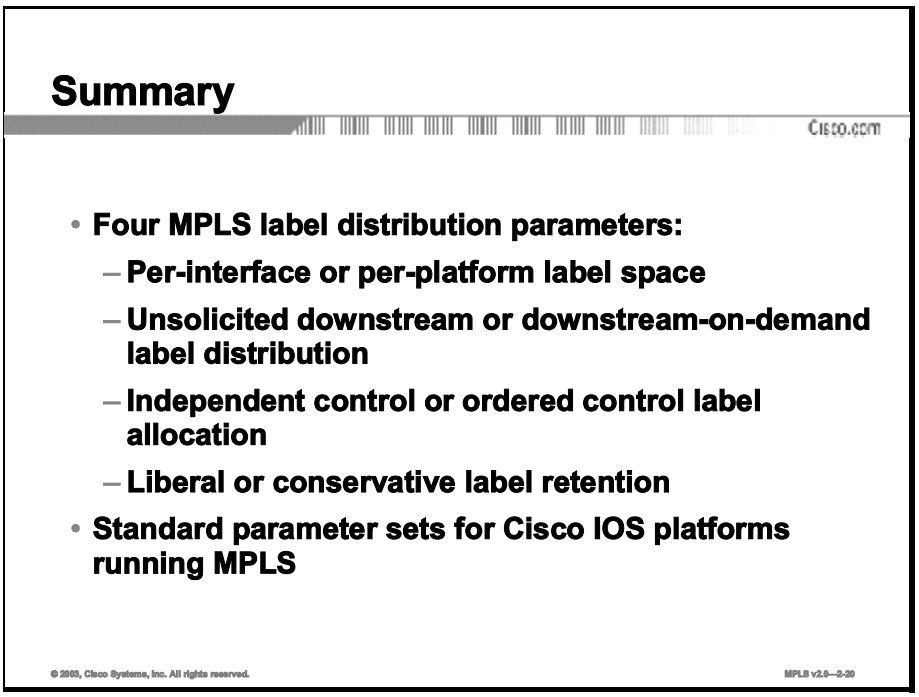

References

For additional information, refer to these resources:

- RFC 3031, “Multiprotocol Label Switching Architecture”
- RFC 3036, “LDP Specification”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following is NOT a label distribution parameter?

A) label space

B) label quality

C) label retention

D) label allocation and distribution

Q2) Cell-mode MPLS uses _____ label space, and frame-mode uses _____ label space.

Q3) Which two types of label distribution are used in Cisco MPLS networks? (Choose two.)

A) downstream-on-demand

B) unsolicited downstream

C) solicited downstream-on-demand

D) unsolicited downstream-on-demand

Q4) The two modes of label allocation are _____ control and _____ control.

Q5) Which two of the following are the label retention modes used in Cisco MPLS networks? (Choose two.)

A) total

B) light

C) liberal

D) conservative

Q6) Which of the following is correct?

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.

Quiz Answer Key

Q1) B

Relates to: Label Distribution Parameters

Q2) per-interface, per-platform

Relates to: Label Space

Q3) A, B

Relates to: Label Distribution

Q4) independent, ordered (or ordered, independent)

Relates to: Label Allocation

Q5) C, D

Relates to: Label Retention

Q6) D

Relates to: Standard Parameter Sets in Cisco IOS Platform MPLS Implementation

LDP Neighbor Discovery

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.

Relevance

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 where the neighbor is not directly connected to a peer. This information will provide a further understanding of the MPLS technology.

Objectives

This lesson describes how LDP neighbors are discovered.

Upon completing this lesson, you will be able to:

- Describe how LDP sessions are established
- Describe the LDP hello message
- Describe how the label space of an LDP neighbor is announced
- Describe how LDP neighbors are discovered
- Describe the LDP session negotiation between LDP neighbors
- Describe LDP sessions between ATM LSRs
- Describe how LDP discovers nonadjacent neighbors

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Label Allocation, Distribution, and Retention Modes” lesson of this module

Outline

This lesson includes these topics:

- Overview
- LDP Session Establishment
- LDP Hello Message
- Label Space
- LDP Neighbor Discovery
- LDP Session Negotiation
- LDP Sessions Between ATM LSRs
- LDP Discovery of Nonadjacent Neighbors
- Summary
- Quiz

LDP Session Establishment

This topic describes how LDP sessions are established between neighbors.

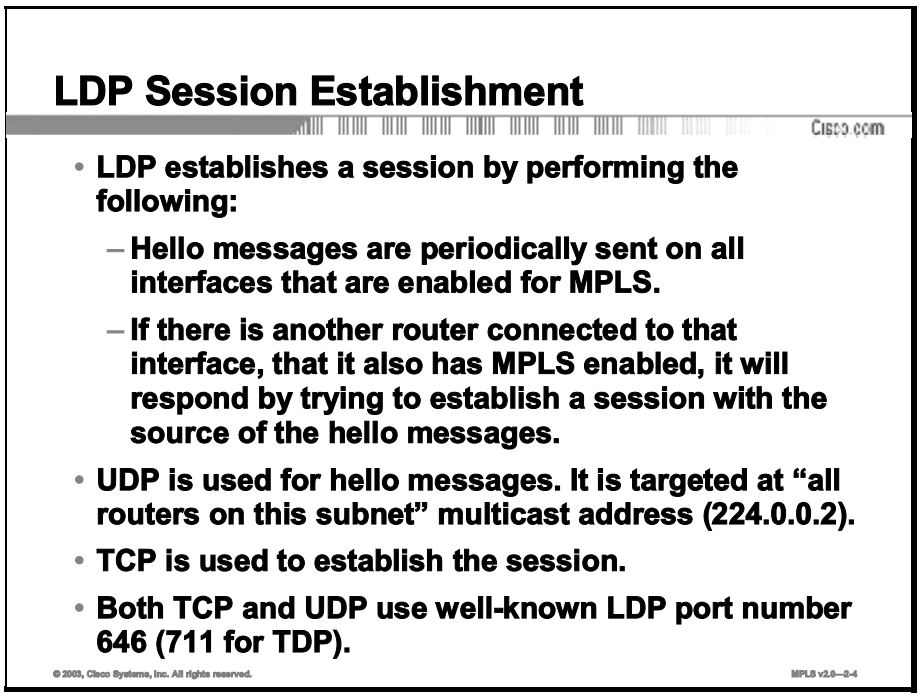

LDP is a standard protocol used to exchange labels between adjacent routers. TDP) is a Cisco proprietary protocol that has the same functionality as LDP.

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).

LDP Hello Message

This topic describes the LDP hello message.

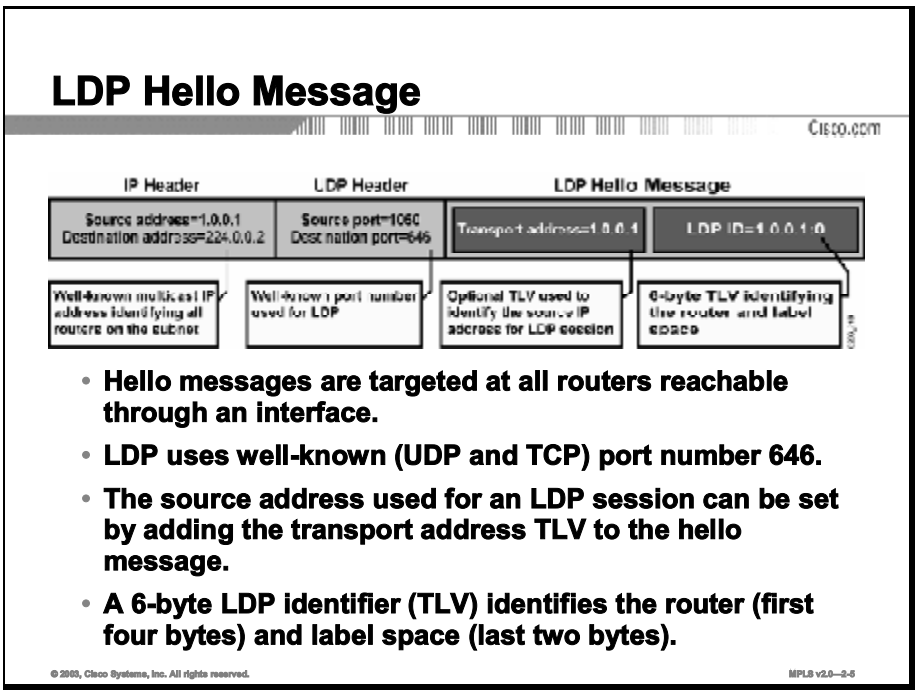

The contents of a hello message are as follows:

- Destination IP address (224.0.0.2), which targets all routes on the subnet
- 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.

Label Space

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

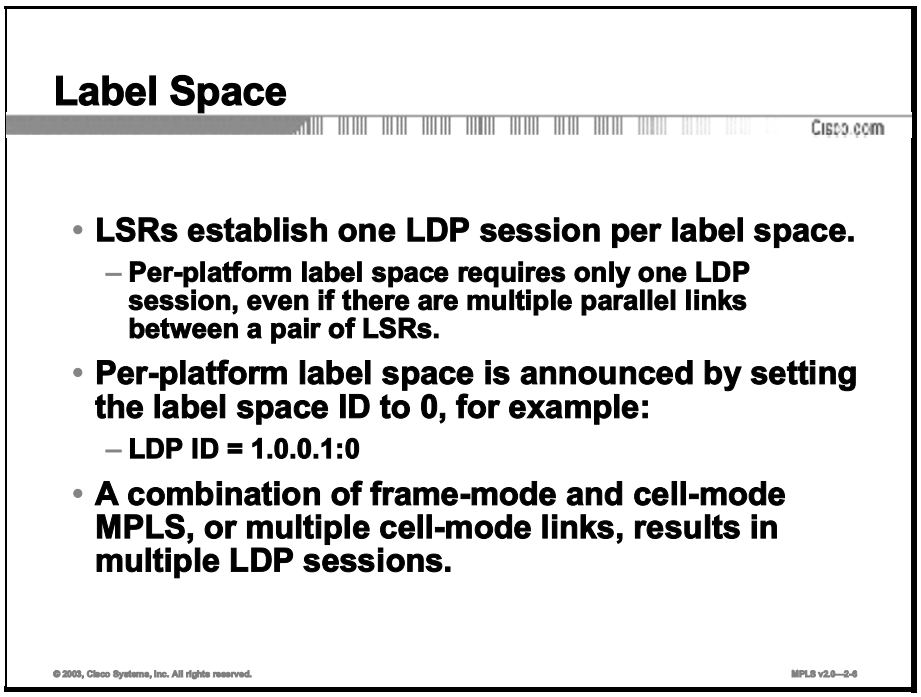

If a pair of routers are connected over two or more parallel links and use frame-mode MPLS, they try to establish multiple sessions by using the same LDP identifier. Because they are using per-platform label space, this action will result in only one session remaining; the other session will be broken.

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.

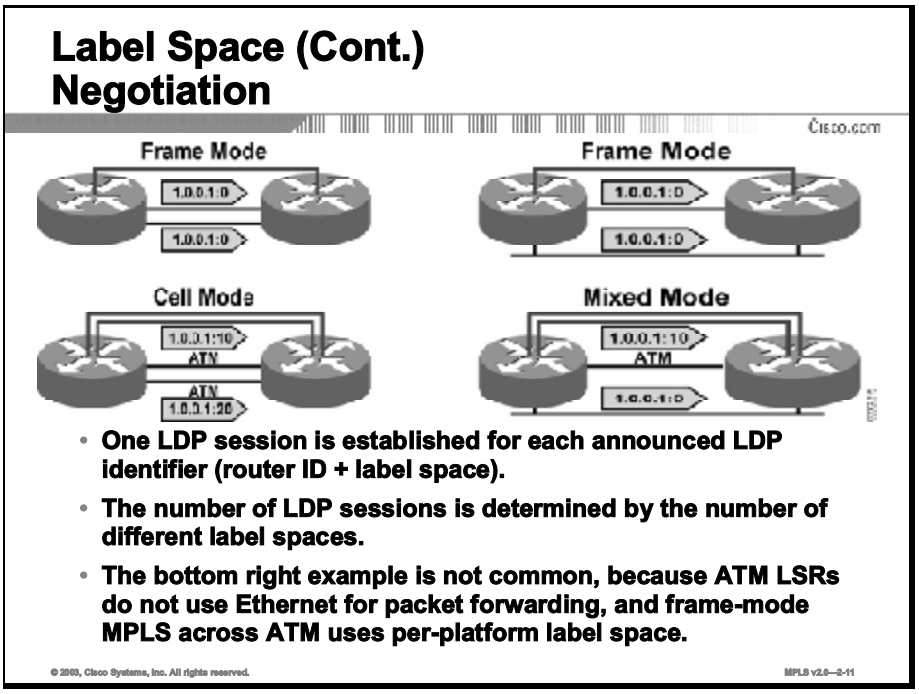

The figure illustrates four different combinations with two parallel links between a pair of routers. The top routers are frame-mode routers.

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.

LDP Neighbor Discovery

This topic describes the LDP neighbor discovery process.

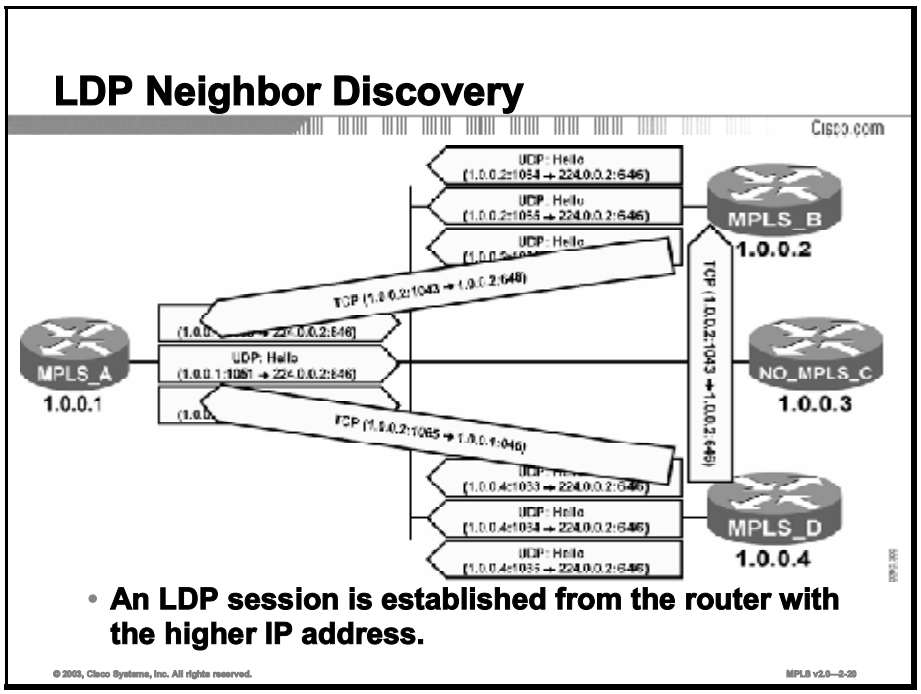

In the figure, three out of four routers periodically send out LDP hello messages (the fourth router is not MPLS-enabled).

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, then 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.

LDP Session Negotiation

This topic describes LDP neighbor session negotiation.

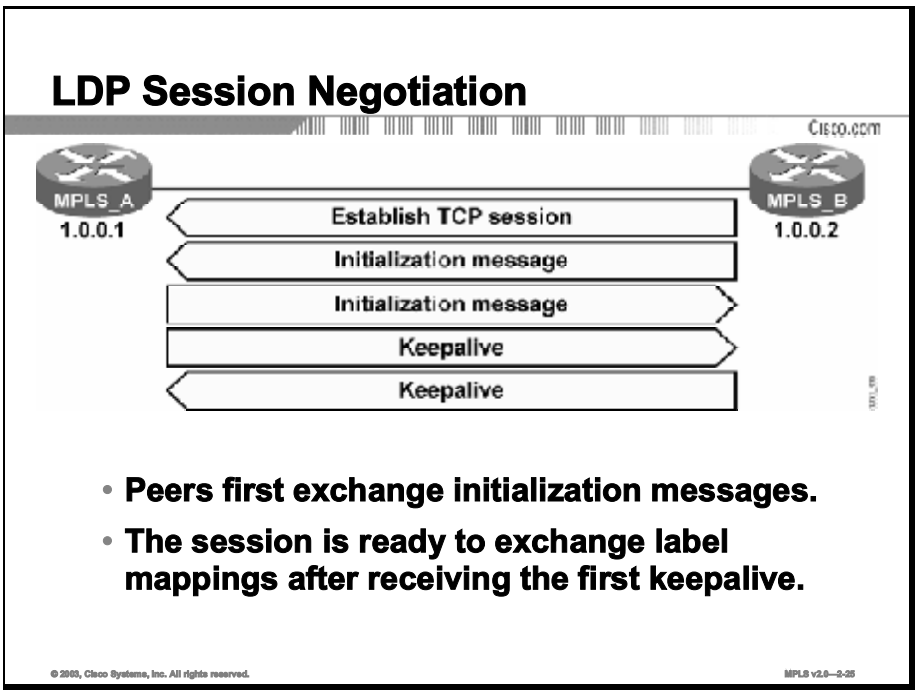

LDP session negotiation is a three-step process:

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.

LDP Sessions Between ATM LSRs

This topic describes how sessions are established between ATM LSRs.

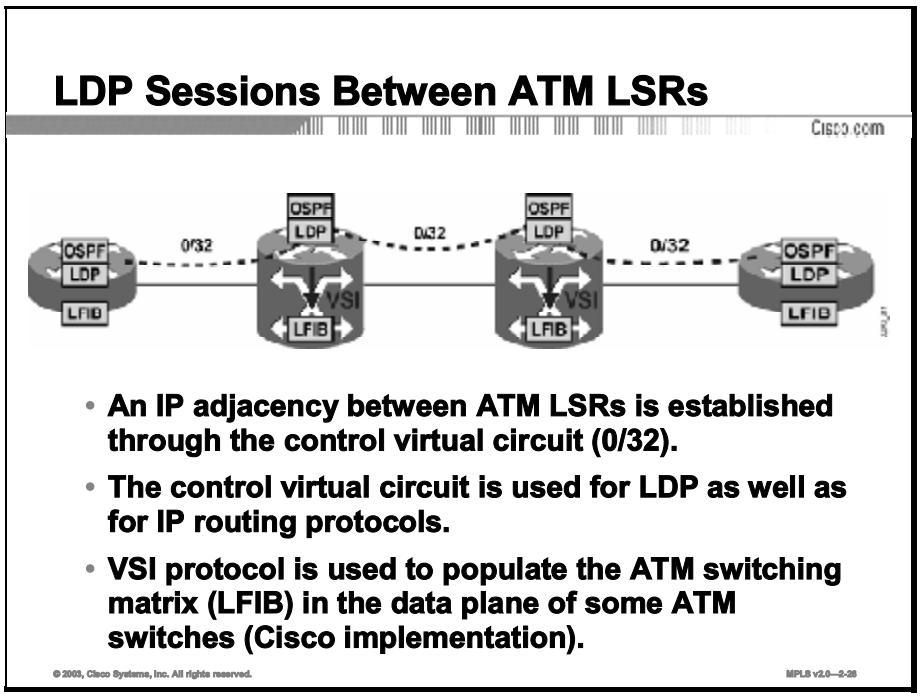

The figure illustrates the operation of LDP in ATM networks. ATM LSRs establish the IP adjacency across the MPLS control virtual circuit, which by default has a VPI/VCI value of 0/32.

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.

LDP Discovery of Nonadjacent Neighbors

This topic describes how LDP discovers nonadjacent neighbors.

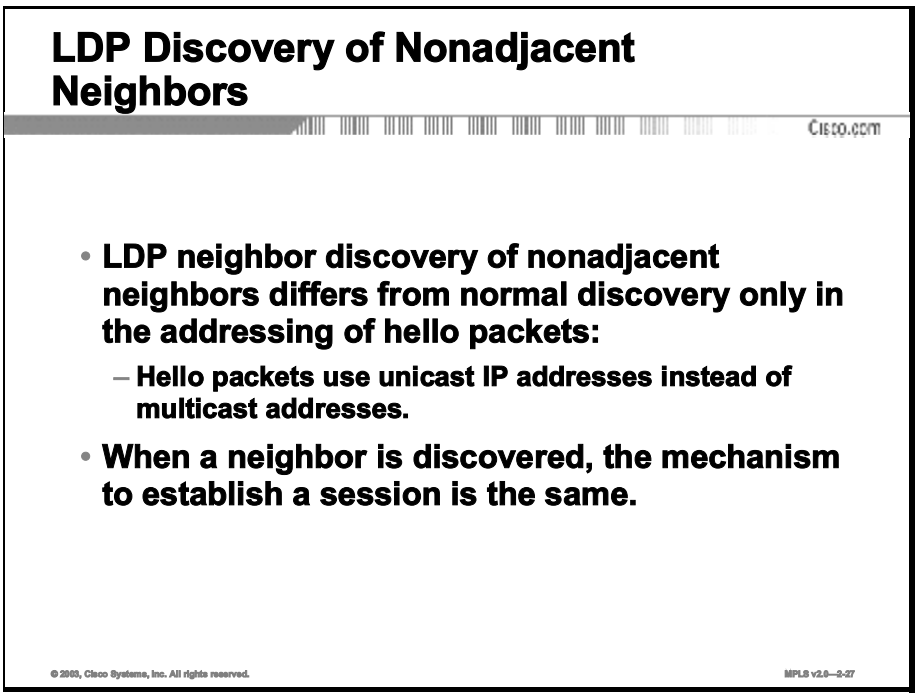

LDP can also be used between nonadjacent routers. However, LDP hello messages use unicast IP addresses instead of multicast. The rest of the session negotiation is the same as for adjacent routers.

Summary

This topic summarizes the key points discussed in this lesson.

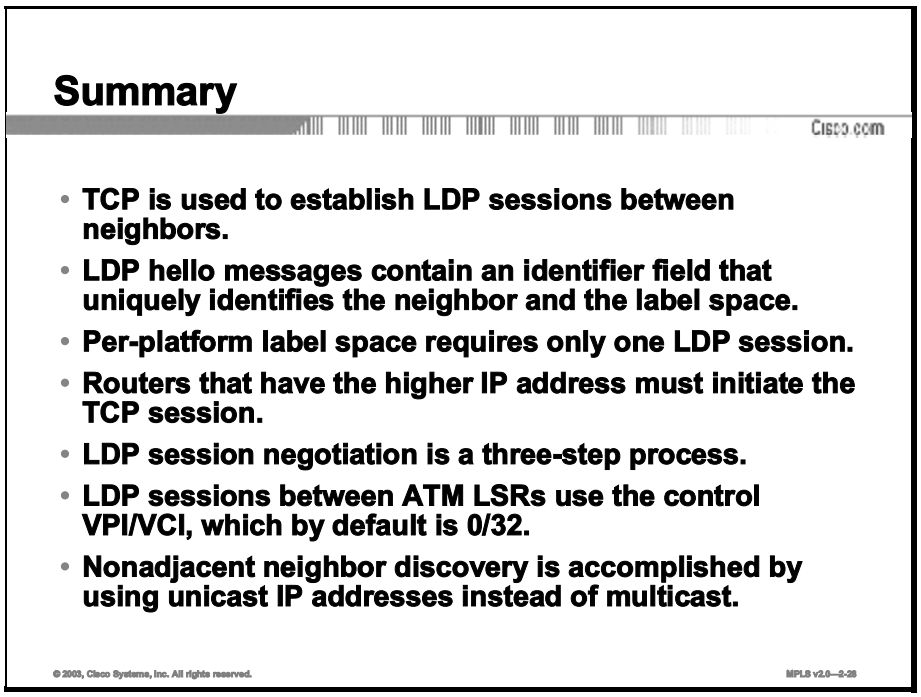

References

For additional information, refer to these resources:

- RFC 3031, “Multiprotocol Label Switching Architecture”
- RFC 3036, “LDP Specification”

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What multicast address does LDP use to send hello messages?

A) 224.0.0.1

B) 224.0.0.2

C) 224.0.0.12

D) 224.0.20.0

Q2) Per-platform label space requires which of the following?

A) Oonly one LDP session

B) one session per interface

C) multiple sessions for parallel links

D) “Per-platform” is not a proper term in MPLS terminology

Q3) What is the purpose of the LDP identifier in a hello message?

A) contains the source address

B) contains the multicast address

C) contains the TCP destination port

D) uniquely identifies the neighbor and the label space

Q4) LDP sessions are initiated by using the _____ IP address.

Q5) Exchanging initialization messages is what step in the LDP session negotiation process?

E) first step in LDP session negotiation

F) second step in LDP session negotiation

G) third step in LDP session negotiation

H) not required in LDP session negotiation

Q6) By default, ATM LSRs establish IP adjacency across what VPI/VCI virtual circuit?

A) 0/32

B) 1/32

C) 32/0

D) 32/1

Q7) LDP discovers nonadjacent neighbors by broadcasting _____ IP addresses.

Quiz Answer Key

Q1) B

Relates to: LDP Session Establishment

Q2) A

Relates to: Label Space

Q3) D

Relates to: LDP Hello Message

Q4) higher

Relates to: LDP Neighbor Discovery

Q5) B

Relates to: LDP Session Negotiation

Q6) A

Relates to: LDP Sessions Between ATM LSRs

Q7) unicast

Relates to: LDP Discovery of Nonadjacent Neighbors

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: Label Assignment and Distribution

Complete the quiz to assess what you have learned in this module.

Objectives

This activity tests your knowledge on how 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 label allocation, distribution, and retention modes
- Describe how LDP neighbors are discovered

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key at the end of this quiz.

Step 3 Review the related lesson for each question that you answered incorrectly.

Quiz

Answer these questions:

Q1) Which table holds all labels assigned by an LSR and their mapping to labels that have been received from the neighbors of the LSR?

A) FIB

B) LIB

C) FLIB

D) LFIB

Q2) What does the term "pop" mean when you are describing penultimate hop popping?

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

Q3) Which of the following best describes convergence in a frame-mode MPLS network after a link failure has occurred and been restored?

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.

Q4) A solution for cell interleaving that could occur in ATM MPLS networks is which of the following?

A) PHS

B) VC merge

C) PC merge

D) PSS

Q5) Which two statements are correct? (Choose two.)

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.

Q6) LDP and TDP use what two well-known port-numbers? (Choose two.)

A) LDP uses 464.

B) LDP uses 646.

C) LDP uses 711.

D) TDP uses 171.

E) TDP uses 646.

F) TDP uses 711.

Q7) In frame-mode MPLS networks the number of LDP sessions that are required between neighbors is determined by?

A) number of interfaces

B) number of different label spaces

C) number of LDP processes running a router

D) the information contained in the source address field of the hello message response

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) B

Relates to: Typical Label Distribution in Frame-Mode MPLS

Q2) C

Relates to: Typical Label Distribution in Frame-Mode MPLS

Q3) A

Relates to: Convergence in Frame-Mode MPLS

Q4) B

Relates to: Typical Label Distribution over LC-ATM Interfaces and VC Merge

Q5) B, C

Relates to: MPLS Label Allocation, Distribution, and Retention Modes

Q6) B, F

Relates to: LDP Neighbor Discovery

Q7) B

Relates to: LDP Neighbor Discovery

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). It 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, it 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 LC-ATM Cisco IOS platforms. This includes being able to do the following:

- Explain the basics 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

Module Outline

The module contains these lessons:

- CEF Switching Review
- Configuring Frame-Mode MPLS on Cisco IOS Platforms
- Monitoring Frame-Mode MPLS on Cisco IOS Platforms
- Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms
- Configuring LC-ATM MPLS
- Configuring LC-ATM MPLS over ATM Virtual Path
- Monitoring LC-ATM MPLS on Cisco IOS Platforms

CEF Switching Review

Overview

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

Relevance

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

This lesson describes the basics of CEF switching.

Upon completing this lesson, you will be able to:

- Describe the Cisco IOS platform switch mechanisms
- Describe standard IP switching
- Describe CEF switching
- Configure IP CEF switching
- Monitor IP CEF switching

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Label Assignment and Distribution” module of this course

Outline

This lesson includes these topics:

- Overview
- Cisco IOS Platform Switching Mechanisms
- Standard IP Switching Review
- CEF Switching Review
- Configuring IP CEF
- Monitoring IP CEF
- Summary
- Quiz

Cisco IOS Platform Switching Mechanisms

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

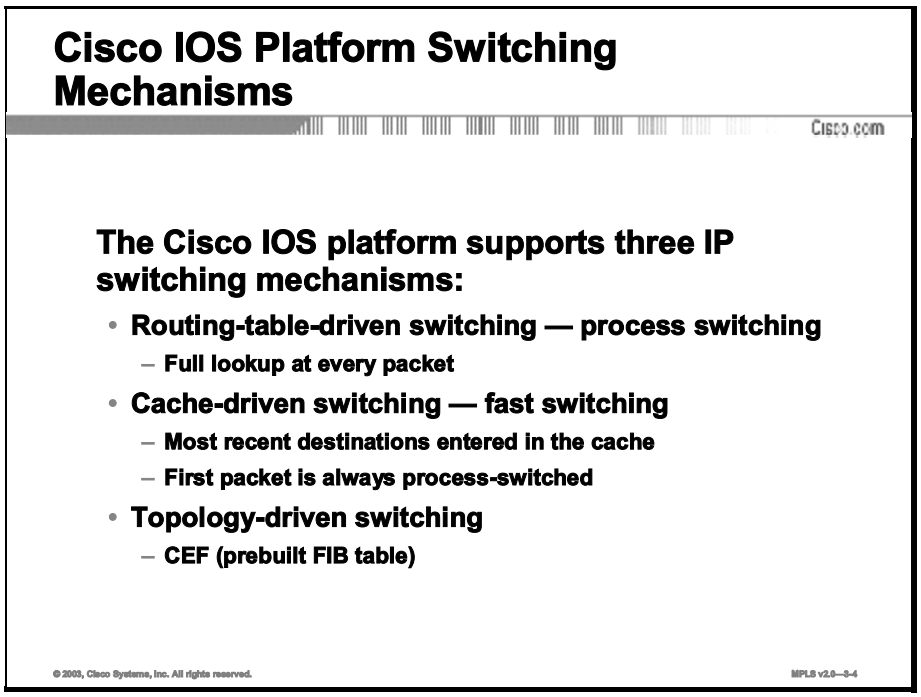

The first and the oldest switching mechanism available in Cisco routers is process switching. Because it has to find a destination in the routing table (possibly a recursive lookup) and construct a new Layer 2 frame header for every packet, it is very slow and is normally not used.

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 following 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. It 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.

Standard IP Switching Review

This topic presents a review of standard IP switching on Cisco IOS platforms.

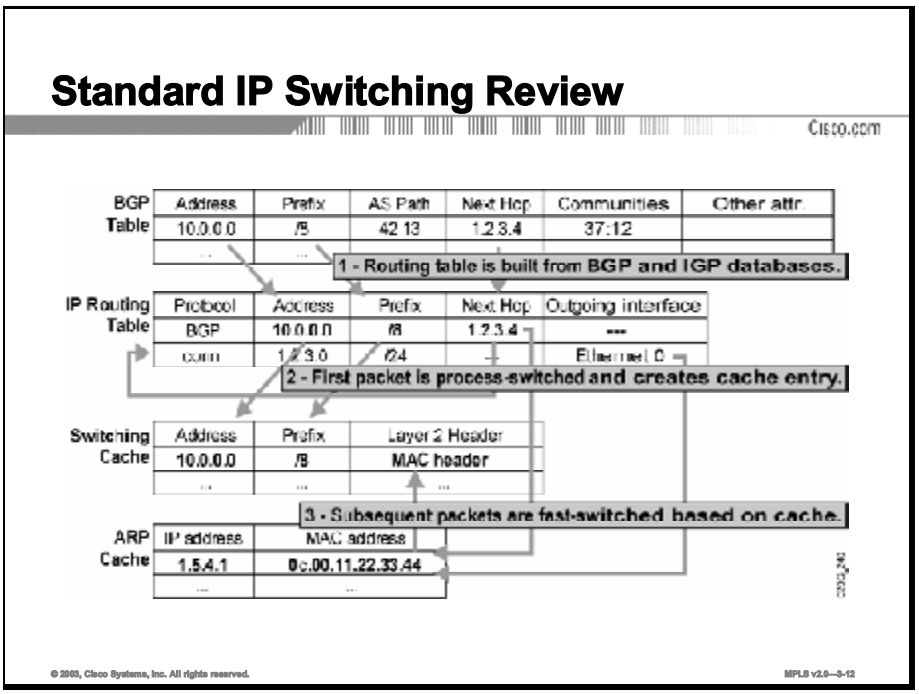

The figure shows a sequence of events when process switching and fast switching are used for destinations learned through Border Gateway Protocol (BGP):

- 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 it is not there, process switching has to switch the packet when the process is run. The process performs a recursive lookup to find the outgoing interface. It 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:
  - 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 as well as 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.

CEF Switching Review

This topic presents a review of CEF switching on Cisco IOS platforms.

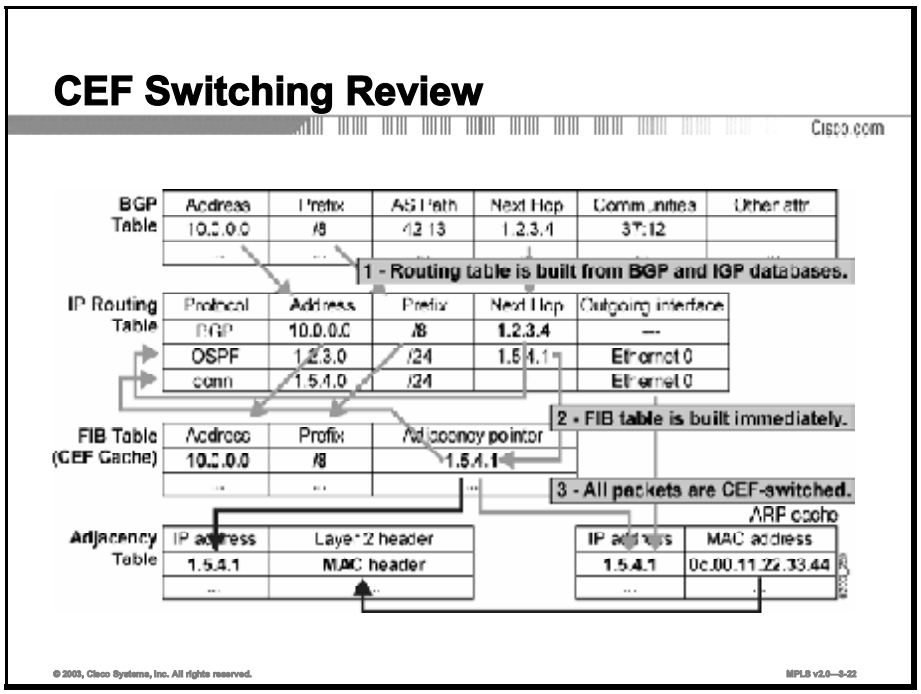

CEF uses a different architecture from process switching or any other cache-based switching mechanism. CEF uses a complete IP switching table, the FIB table, which holds the same information as the IP routing table. The generation of entries in the FIB table is not packet-triggered but change-triggered.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

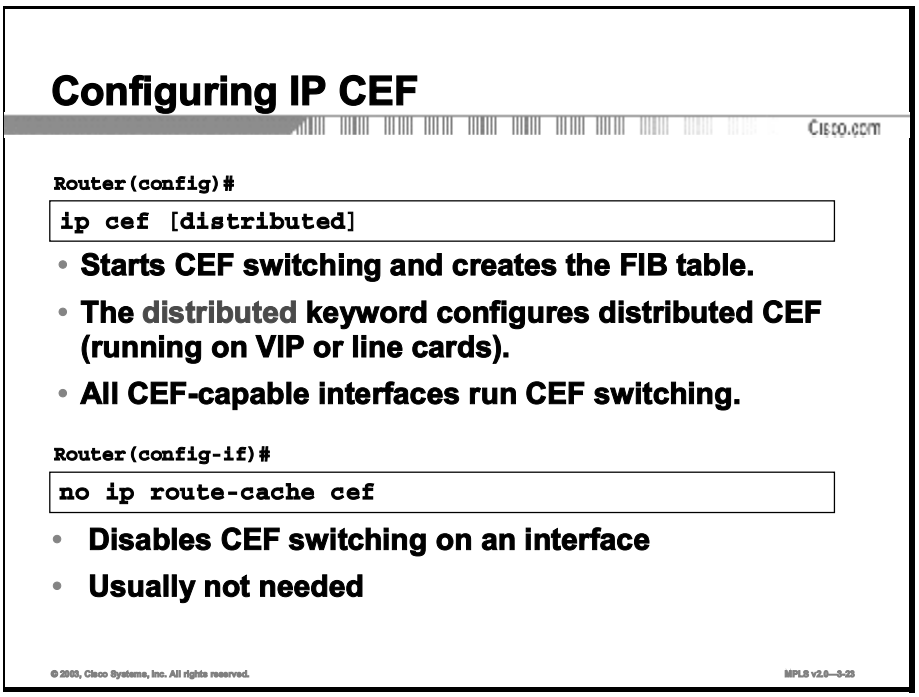

ip cef

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:

- **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.

Defaults

On this platform...	The default is...
Cisco 7000 series Route Switch Processor (RSP7000)	CEF is not enabled.
Cisco 7200 series router	CEF is not enabled.
Cisco 7500 series router	CEF is not enabled.
Cisco 12000 series Gigabit Switch Router (GSR)	CEF is not enabled.

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:

- **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.

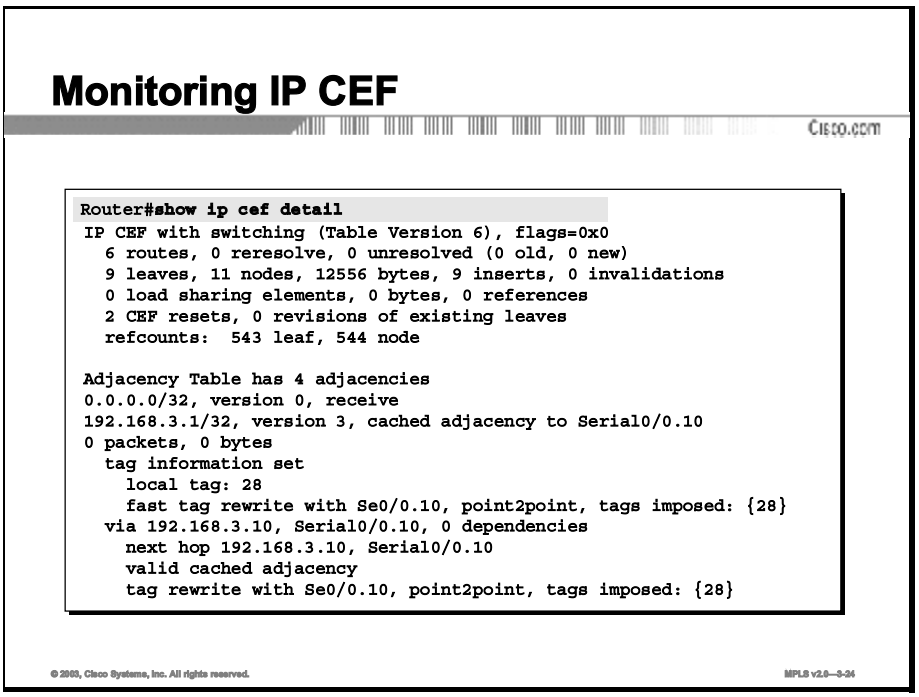

show ip cef

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**]

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.

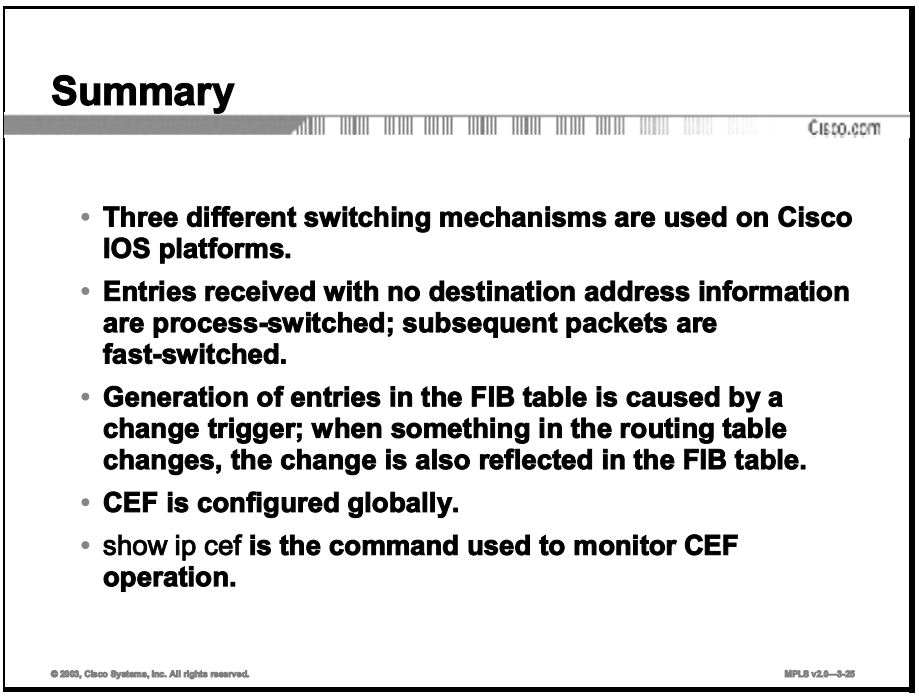

References

For additional information, refer to this resource:

- Search www.cisco.com for “CEF switching” for additional information.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Topology-driven switching is also known as?

A) CEF

B) fast switching

C) cache switching

D) process switching

Q2) The command to monitor CEF is?

A) **Router#show cef**

B) **Router>show ip cef**

C) **Router#show ip cef**

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

Q3) The command to enable CEF on a Cisco router is?

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?

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?

A) The packet is dropped.

B) The packet is cache-switched.

C) The packet is process-switched.

D) CEF switching is used.

Quiz Answer Key

Q1) A

Relates to: Cisco IOS Platform Switching Mechanisms

Q2) C

Relates to: Monitoring IP CEF

Q3) D

Relates to: Configuring IP CEF

Q4) A

Relates to: CEF Switching Review

Q5) C

Relates to: Standard IP Switching Review

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.

Relevance

It is important to understand how to enable and configure MPLS, because it will be required to successfully complete the lab for this lesson.

Objectives

This lesson describes how to configure frame-mode MPLS on Cisco IOS platforms.

Upon completing this lesson, you will be able to:

- Describe the MPLS configuration tasks
- Configure the MPLS ID on a router
- Enable and 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

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “CEF Switching Review” lesson of this module

Outline

This lesson includes these topics:

- Overview
- MPLS Configuration Tasks
- Configuring MPLS on a Frame-Mode Interface
- Configuring a Label-Switching MTU
- Configuring IP TTL Propagation
- Conditional Label Distribution
- Configuring Frame-Mode MPLS on Switched WAN Media
- Summary
- Quiz

MPLS Configuration Tasks

This topic describes the mandatory MPLS configuration tasks and those that are optional.

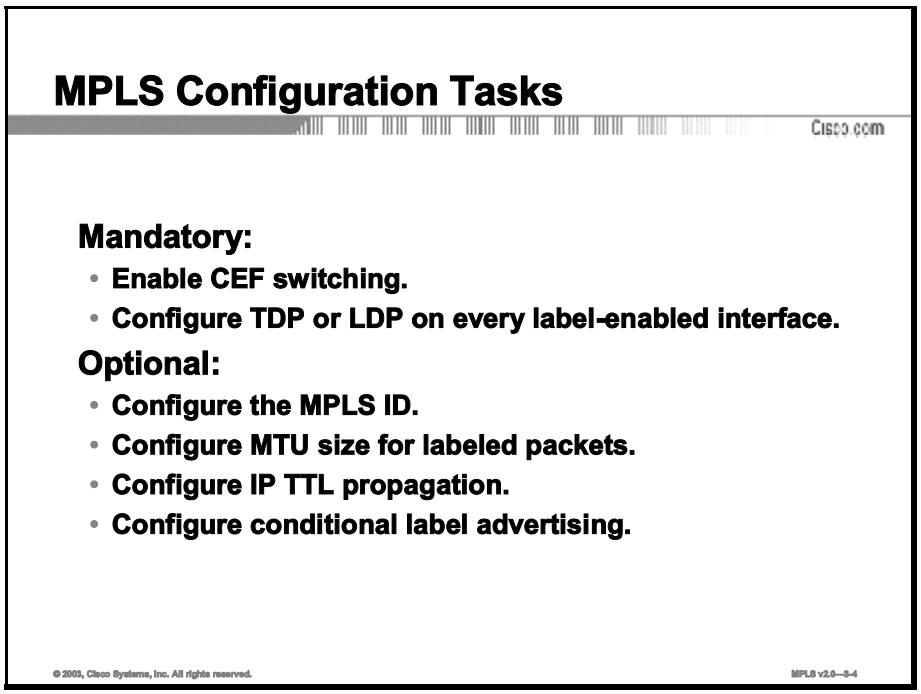

To enable MPLS, you must first enable CEF switching. Depending on the Cisco IOS software version, you may need to establish the range for the label pool.

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

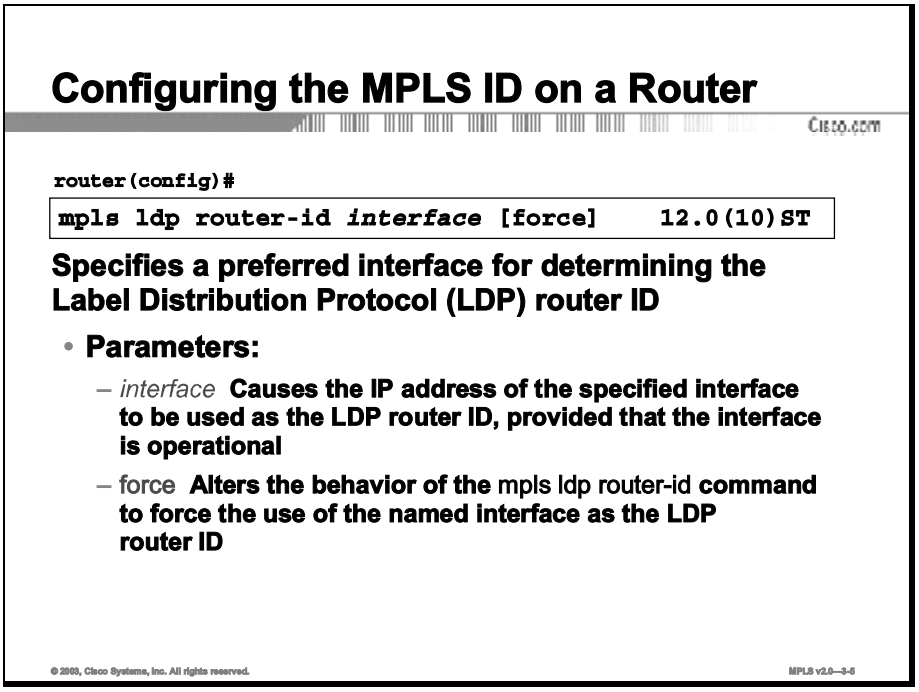

mpls ldp router-id

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:

- **mpls ldp router-id** *interface* [**force**]
- **no mpls ldp 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 enable and configure MPLS on a frame-mode interface.

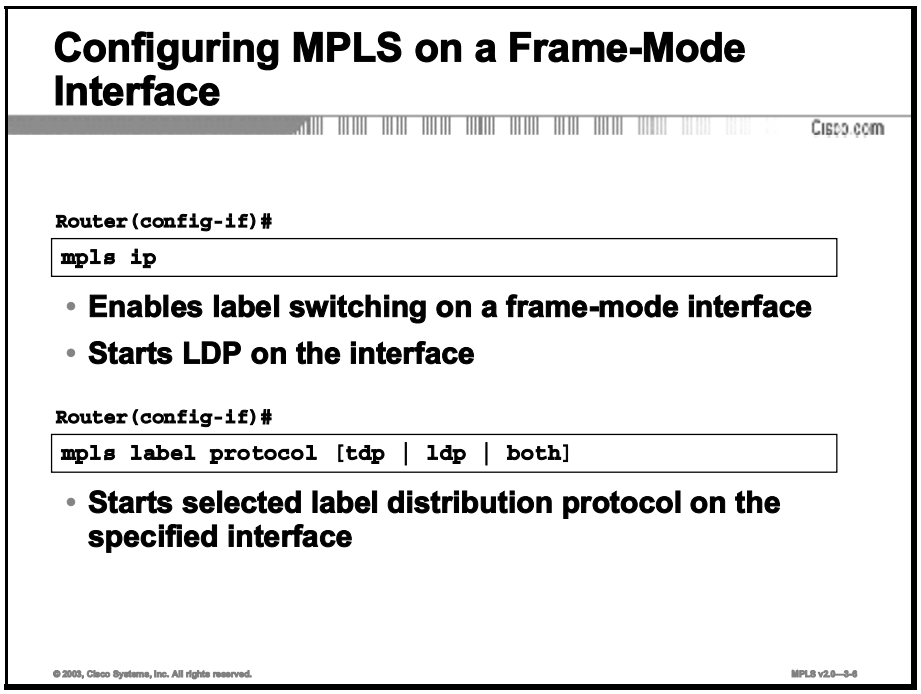

mpls ip

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:

- **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:

- **mpls label protocol <protocol>**
- **no mpls label protocol <protocol>**

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 in the tag-switching syntax in the configuration by the IOS software.

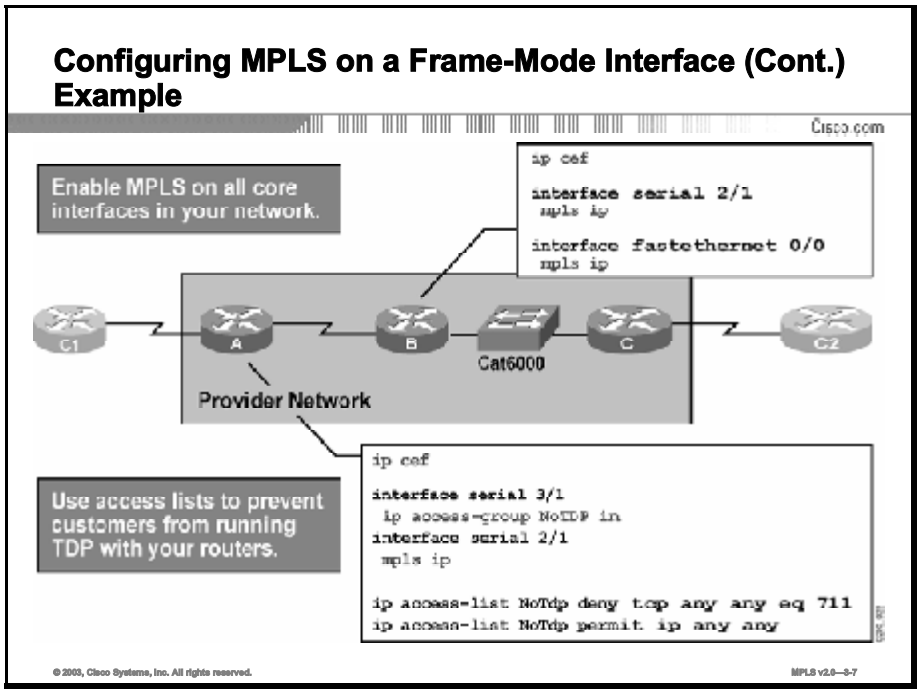

The figure shows the configuration steps needed to enable MPLS on an edge label switch router (LSR). The configuration includes an access list that denies any attempt to establish a TDP session from an interface that is not enabled for MPLS.

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 access list that denies TCP sessions on the well-known port number 711 (TDP).

Configuring MPLS on a Frame-Mode Interface (Cont.) Example

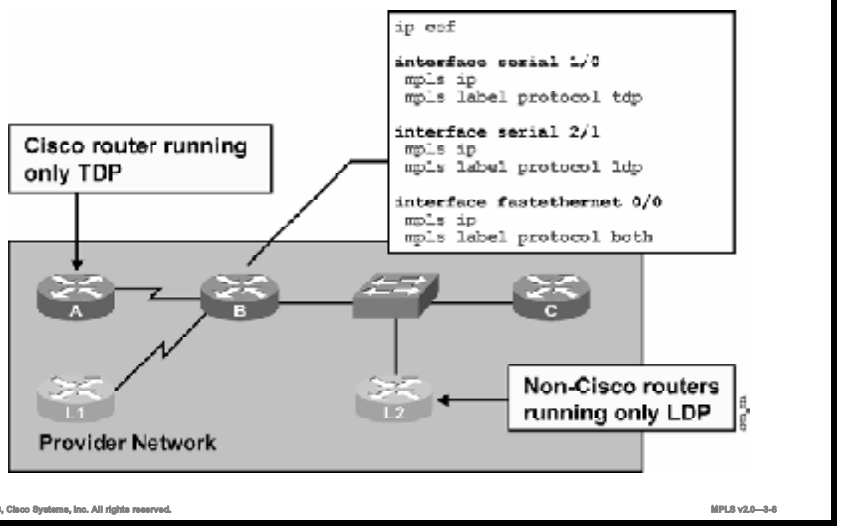

When combining Cisco routers with equipment of other vendors, you may need to use standard LDP (MPLS). TDP (tag switching) can be replaced by LDP on point-to-point interfaces. However, you can also use both protocols on shared media if some devices do not support TDP.

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 change the MTU size in label switching.

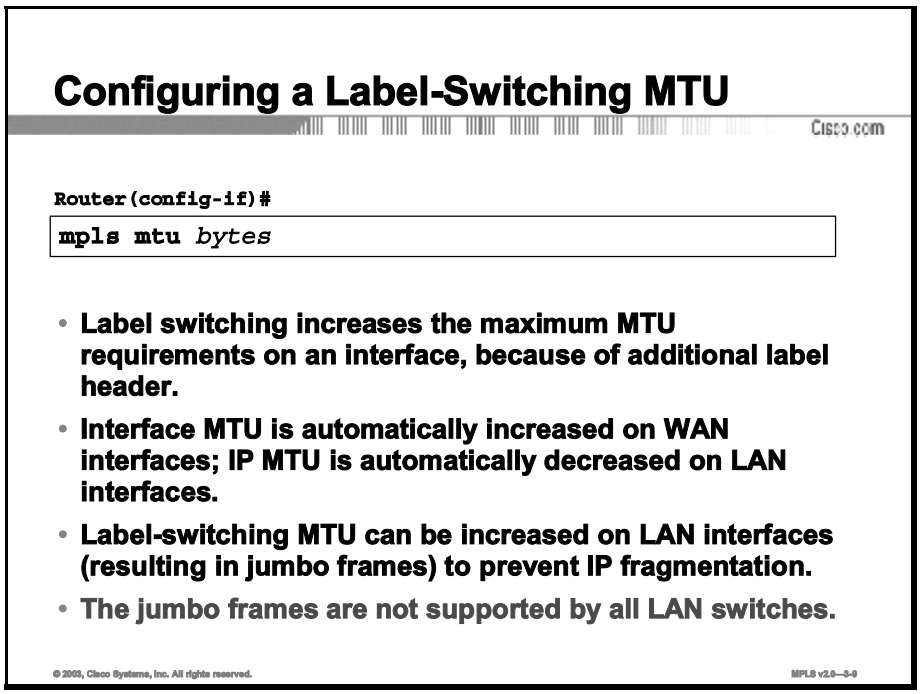

mpls mtu

To set the per-interface MTU for labeled packets, use the **mpls mtu** interface configuration command:

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

Syntax Description

Parameter	Description
bytes	MTU in bytes.

Defaults

The minimum MTU is 128 bytes. The maximum depends on the type of interface medium.

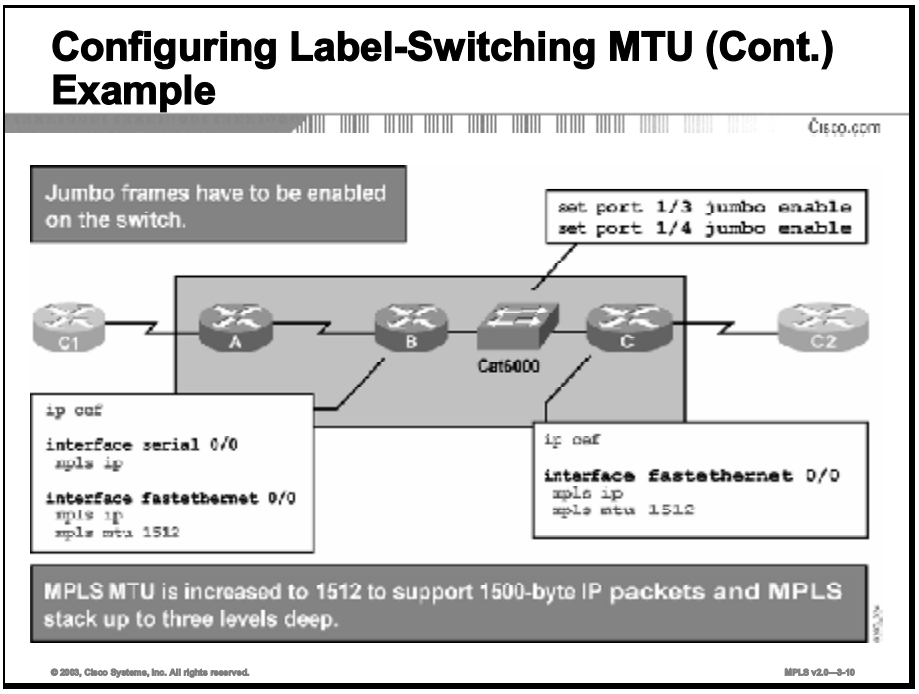

One way of preventing labeled packets from exceeding the maximum size (and being fragmented as a result) is to increase the MTU size of labeled packets for all segments in the label switched path (LSP) tunnel. The problem will typically occur on LAN switches, where it is more likely that a device does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

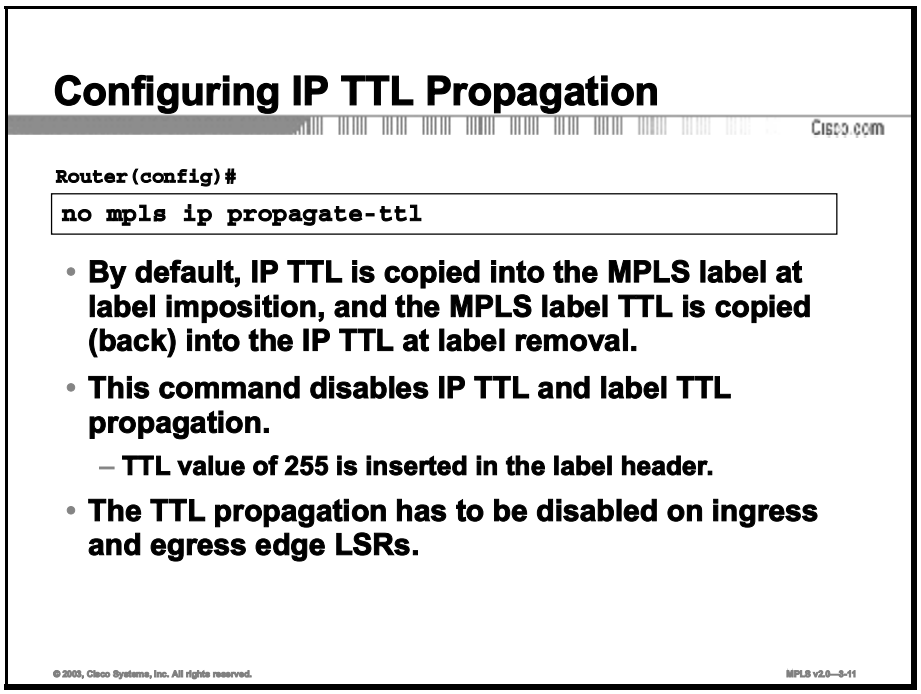

mpls ip propagate-ttl

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:

- **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.

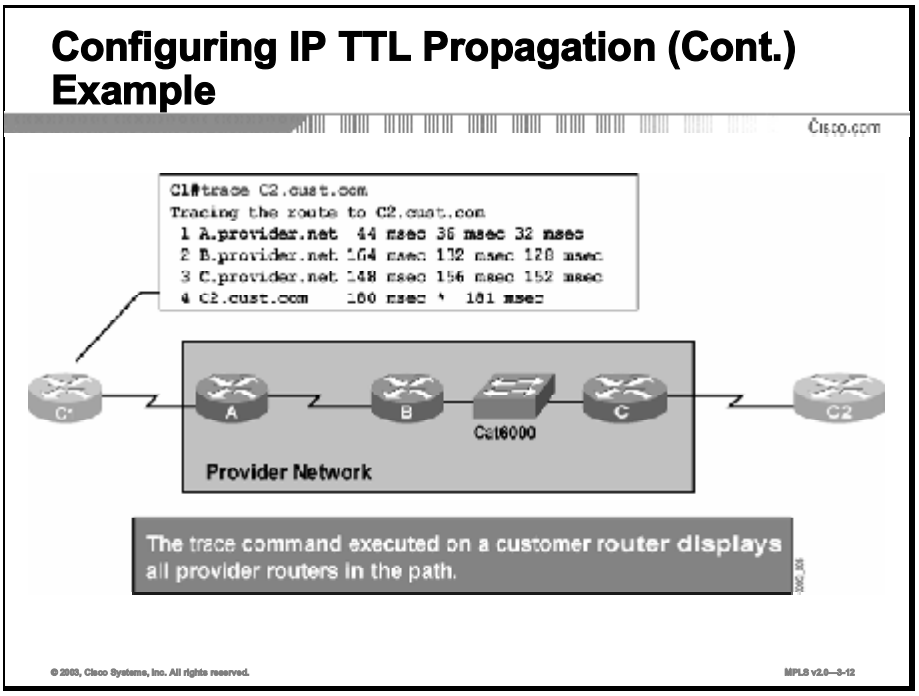

The figure illustrates typical traceroute behavior in an MPLS network. Because the label header of a labeled packet carries the TTL value from the original IP packet, the routers in the path can drop packets when the TTL is exceeded. Traceroute will therefore show all the routers in the path. This is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

In the example, router C1 is executing a **trace** command that results in this behavior:

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.

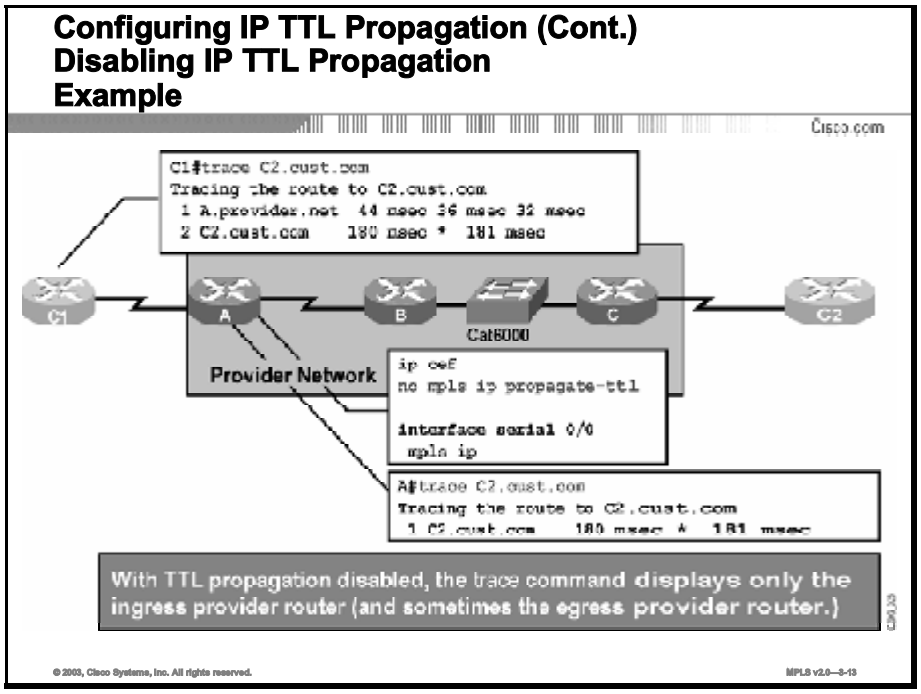

If TTL propagation is disabled, the TTL value is not copied into the label header. Instead, the label TTL field is set to 255. The probable result is that no router in the TTL field in the label header will be decreased to 0 inside the MPLS domain (unless there is a forwarding loop inside the MPLS network).

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

In the example, 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).

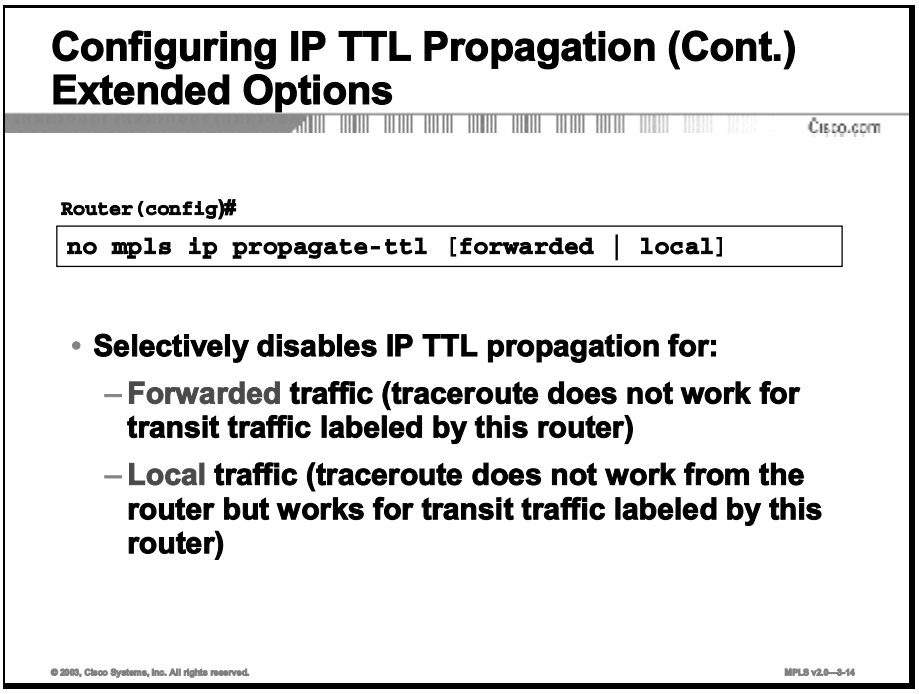

mpls ip propagate-ttl

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 the TTL field is copied from the IP header and inserted into the MPLS label. This aspect allows a **trace** command to show all 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. This is the most common application of the command:

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

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 the hops in the network.

Command Modes

Global configuration

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 the provider.

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

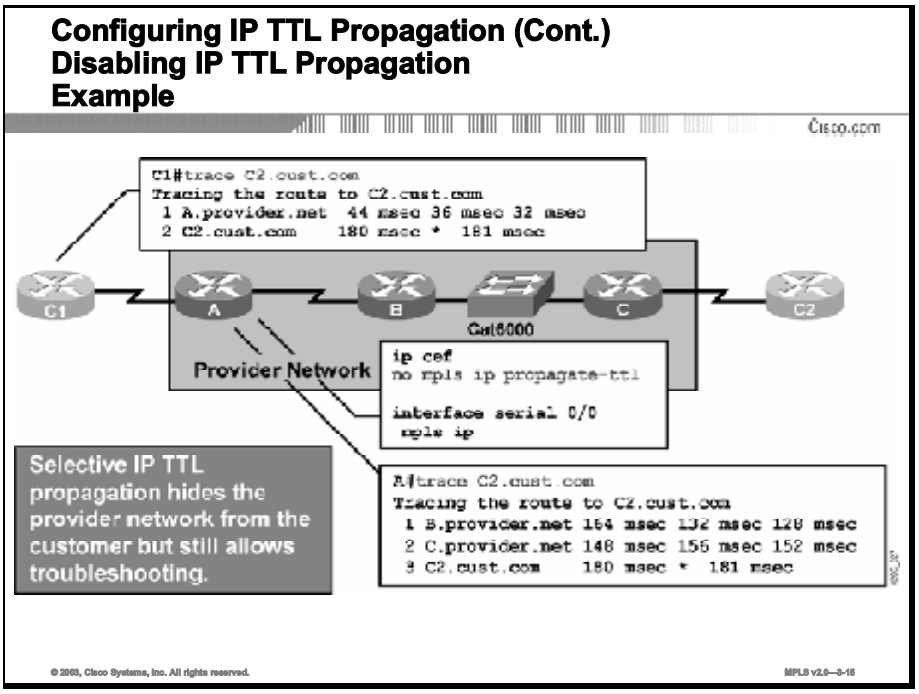

The third option for TTL propagation is the **no mpls ip propagate-ttl forwarded** command. Typically, a service provider likes to hide the backbone network from outside users but allow inside traceroute to work for easier troubleshooting of the network.

This goal can be achieved by disabling TTL propagation for forwarded packets only:

- 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 router will see all the backbone routers. Others 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.

Conditional Label Distribution

This topic discusses configuring conditional label distribution.

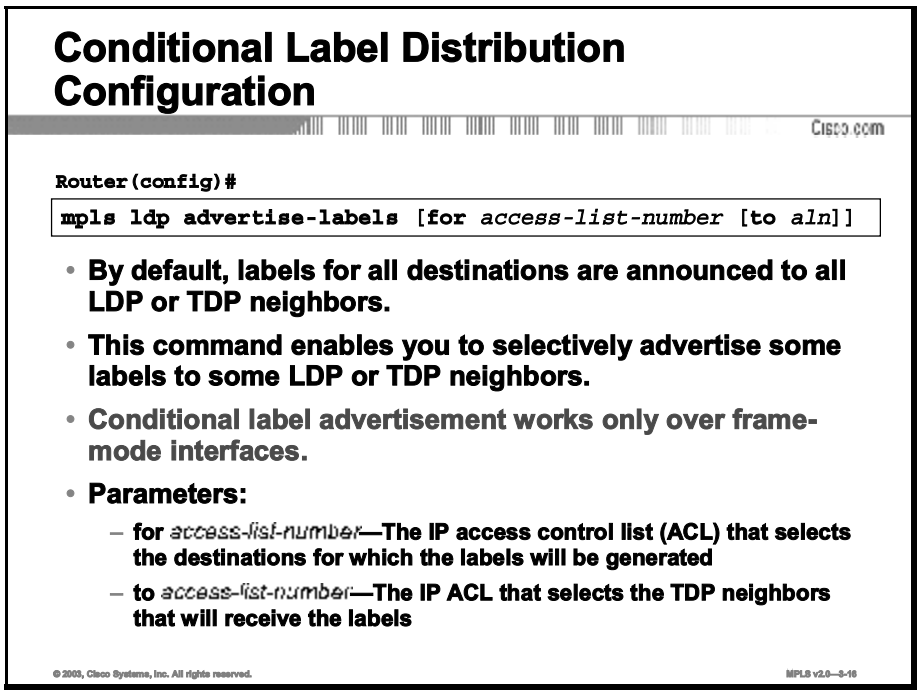

mpls ldp advertise-labels

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:

- **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]]

Syntax Description

Parameter	Description
for *access-list number*	(Optional) Specifies which destinations should have their labels advertised.
to *access-list number*	(Optional) Specifies which LSR neighbors should receive label advertisements. A label switch router (LSR) is identified by its router ID, which consists of the first four bytes of its six-byte LDP identifier.

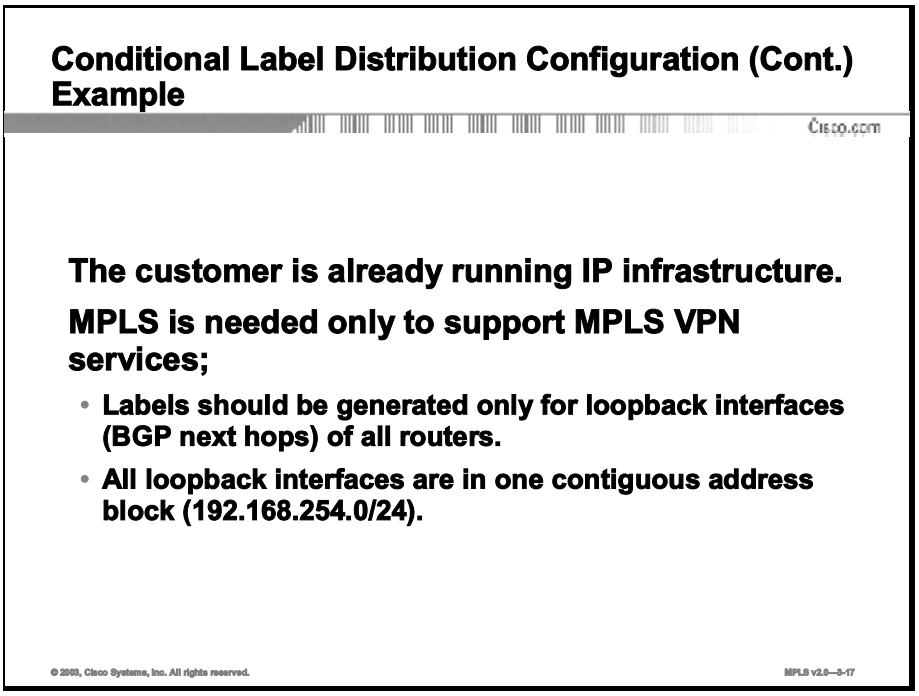

The example here describes where conditional label advertising can be used. The existing network still performs normal IP routing, but the MPLS LSP tunnel between the loopback interfaces of the LSR routers is needed to enable MPLS Virtual Private Network (VPN) functionality.

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

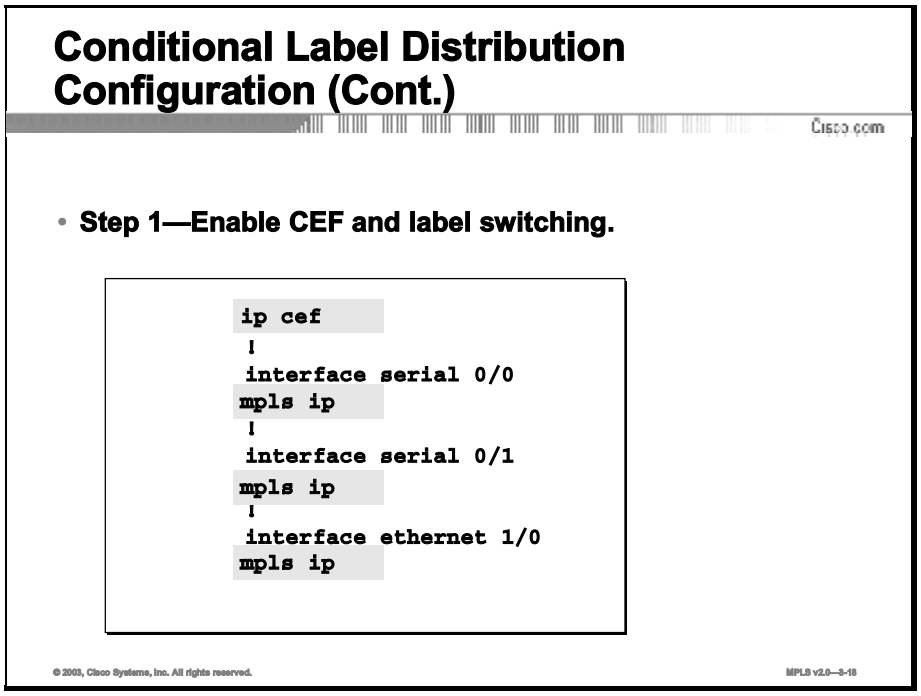

In the first step, CEF switching and MPLS have to be enabled on all core interfaces. The MPLS MTU size may be adjusted on the LAN interfaces.

Conditional Label Distribution Configuration (Cont.)

Cisco.com

- **Step 2— Enable conditional label advertisement.**

```
!
! Disable default advertisment mechanism
!
no mpls ldp advertise-labels
!
! Configure conditional advertisments
!
mpls ldp advertise-labels for 90 to 91
!
access-list 90 permit 192.168.254.0 0.0.0.255
access-list 91 permit any
```


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.

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

Configuring Frame-Mode MPLS on Switched WAN Media

This topic discusses the configuration of frame-mode MPLS on switched WAN media.

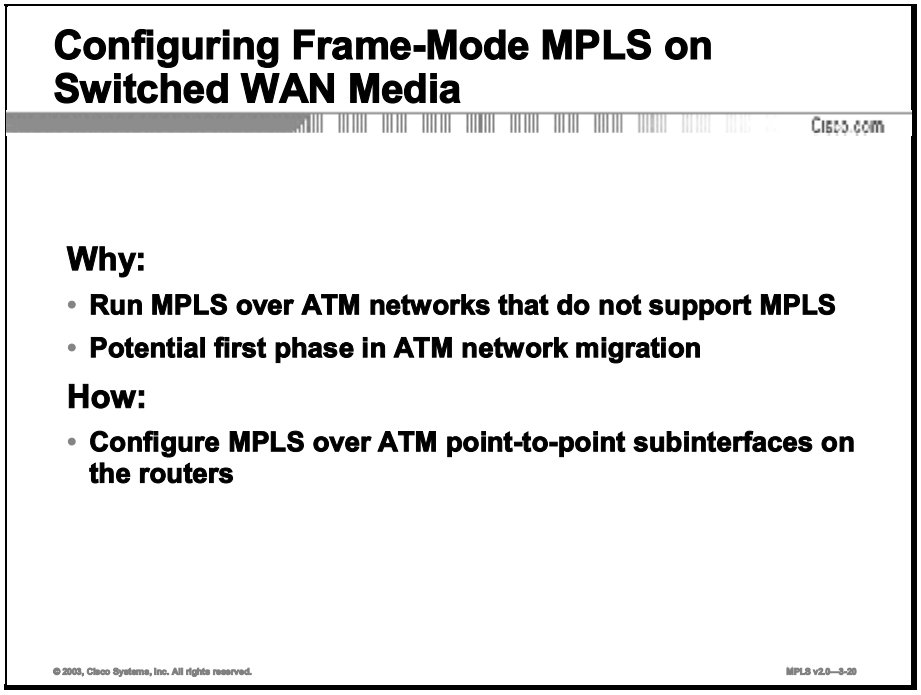

When an underlying ATM infrastructure that does not support cell-mode MPLS is used, MPLS can still be used across point-to-point permanent virtual circuits (PVCs). The MPLS configuration is equal to that on any other Layer 2 media.

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

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

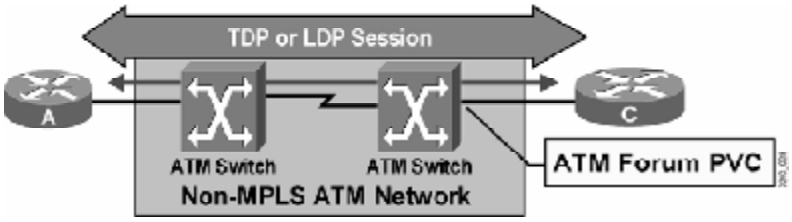

- **Routers view the ATM PVC as a frame-mode MPLS interface.**
- **TDP or LDP is run between the adjacent routers.**
- **Many LSPs can be established over one ATM PVC.**
- **The ATM network is not aware of MPLS between the routers.**

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

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

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

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

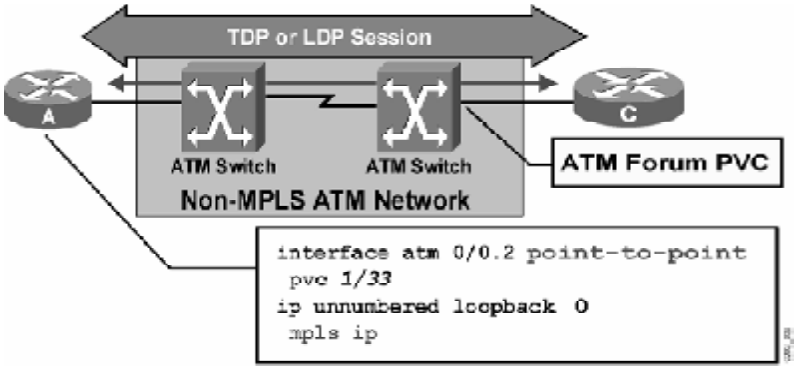

- **Create a point-to-point ATM subinterface.**
- **Configure ATM PVC on the subinterface.**
- **Start label switching and LDP or TDP on the interface.**

Configuring frame-mode MPLS on an ATM interface involves using the same commands as described previously.

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

Configuring Frame-Mode MPLS on Switched WAN Media (Cont.)
MPLS over Frame Relay Networks

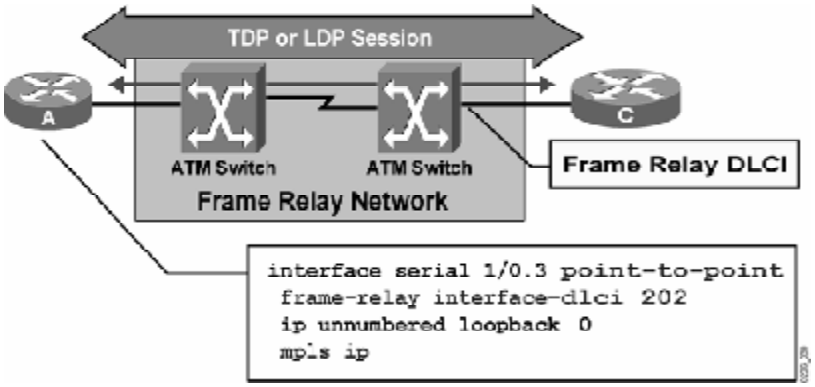

- **Create a point-to-point or multipoint Frame Relay subinterface.**
- **Configure Frame Relay DLCI on the subinterface.**
- **Start label switching and LDP or TDP on the interface.**

Enabling MPLS on a Frame Relay PVC, also called a data-link connection identifier, or 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.

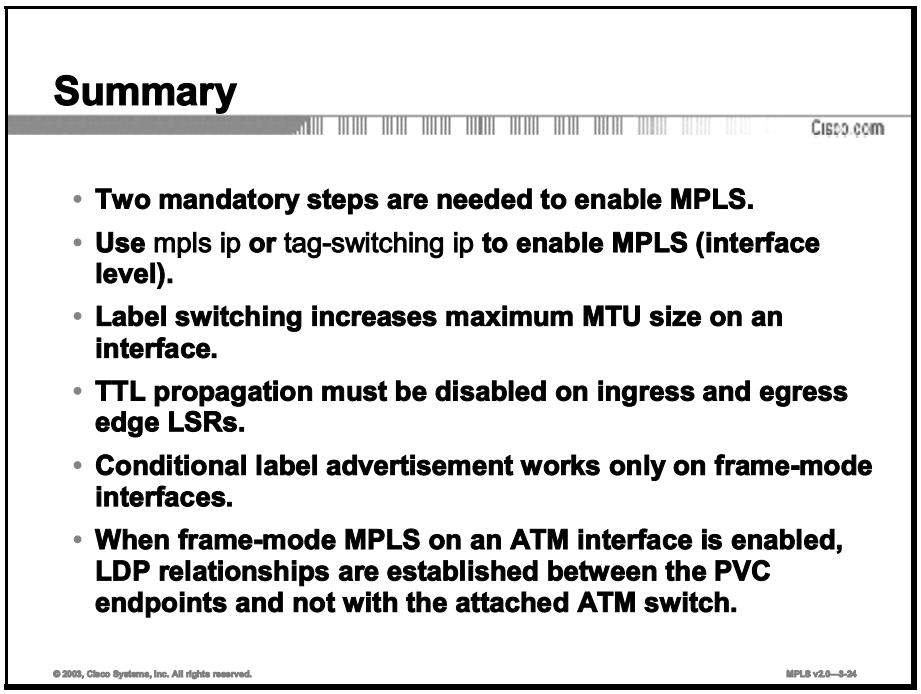

References

For additional information, refer to this resource:

- Search www.cisco.com for further information regarding the topics covered in this lesson.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) If IP TTL propagation is allowed what is the value that is placed in the MPLS header?

A) 0

B) 1

C) 254

D) 255

Q2) The MPLS MTU is increased to _____ to support 1500-byte IP packets and MPLS stacks up to three levels deep.

Q3) Which of the following is the correct command to enable MPLS in Cisco IOS software?

A) **Router#mpls ip**

B) **Router>mpls ip**

C) **Router(config)#mpls ip**

D) **Router(config-if)#mpls ip**

Q4) Which of the following is NOT a mandatory step to enable MPLS?

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

Q5) To specify which neighbors would selectively receive label advertisements, what would need to be configured?

A) controlled label distribution

B) conditional label distribution

C) unsolicited label distribution

D) All neighbors will receive all labels.

Q6) If frame-mode MPLS is run on ATM interfaces, LDP or LDP neighbor relationships are established between the _____routers.

Quiz Answer Key

Q1) D

Relates to: Configuring IP TTL Propagation

Q2) 1512

Relates to: Configuring a Label-Switching MTU

Q3) D

Relates to: Configuring MPLS on a Frame-Mode Interface

Q4) C

Relates to: MPLS Configuration Tasks

Q5) B

Relates to: Conditional Label Distribution

Q6) PVC endpoint routers

Relates to: Configuring Frame-Mode MPLS on Switched WAN Media

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.

Relevance

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

This lesson describes how to use monitoring commands in frame-mode MPLS on Cisco IOS platforms.

Upon completing this lesson, you will be able to:

- Describe and use the basic MPLS monitoring commands
- Describe and use the commands to monitor LDP
- Describe and use the commands to monitor label switching
- Describe and use the commands to debug MPLS and LDP

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Configuring Frame-Mode MPLS on Cisco IOS platforms” lesson of this module

Outline

This lesson includes these topics:

- Overview
- MPLS Monitoring Commands
- LDP Monitoring Commands
- Monitoring Label Switching
- Debugging MPLS and LDP
- Summary
- Quiz

MPLS Monitoring Commands

This topic describes the basic commands that are used to monitor MPLS.

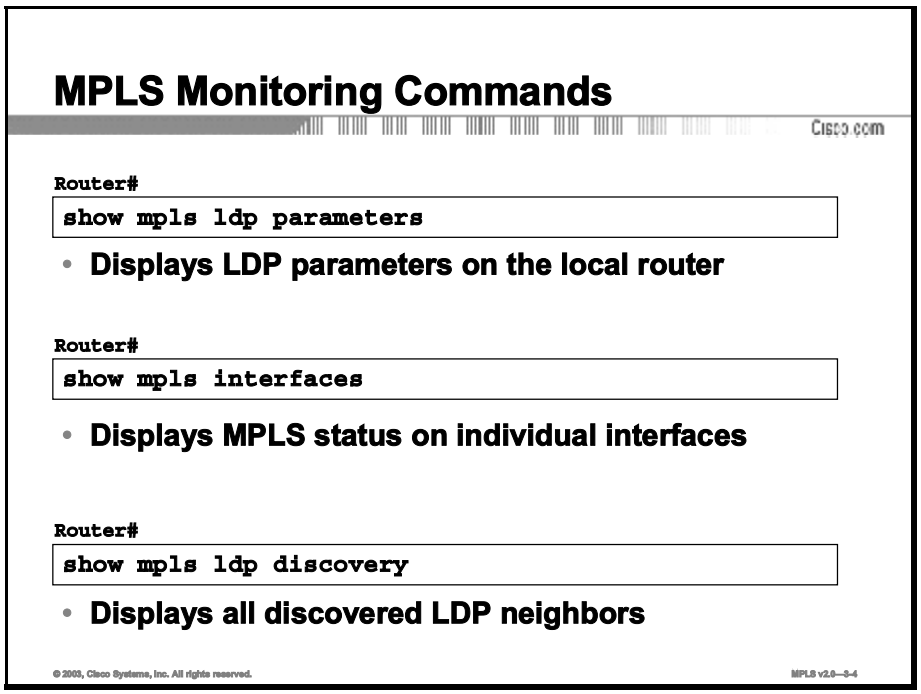

show mpls ldp parameters

To display available LDP parameters, use the **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 **show mpls interfaces** command in EXEC mode:

- **show mpls interfaces** [*interface*] [**detail**]

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.

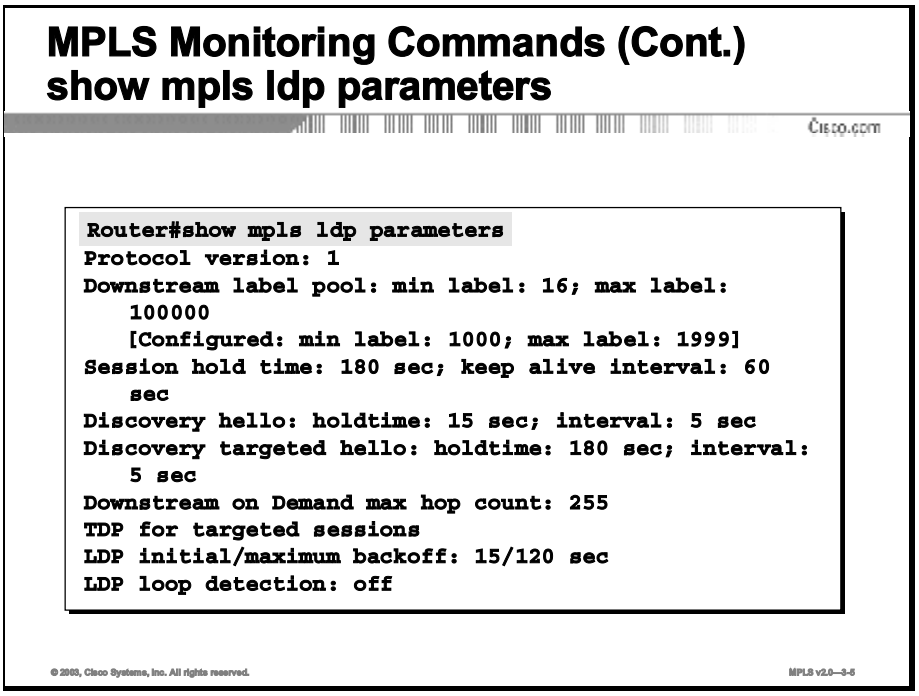

show mpls ldp parameters

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

- **show mpls ldp parameters**

Syntax Description

This command has no arguments or keywords.

The following table describes the significant fields in the display.

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.

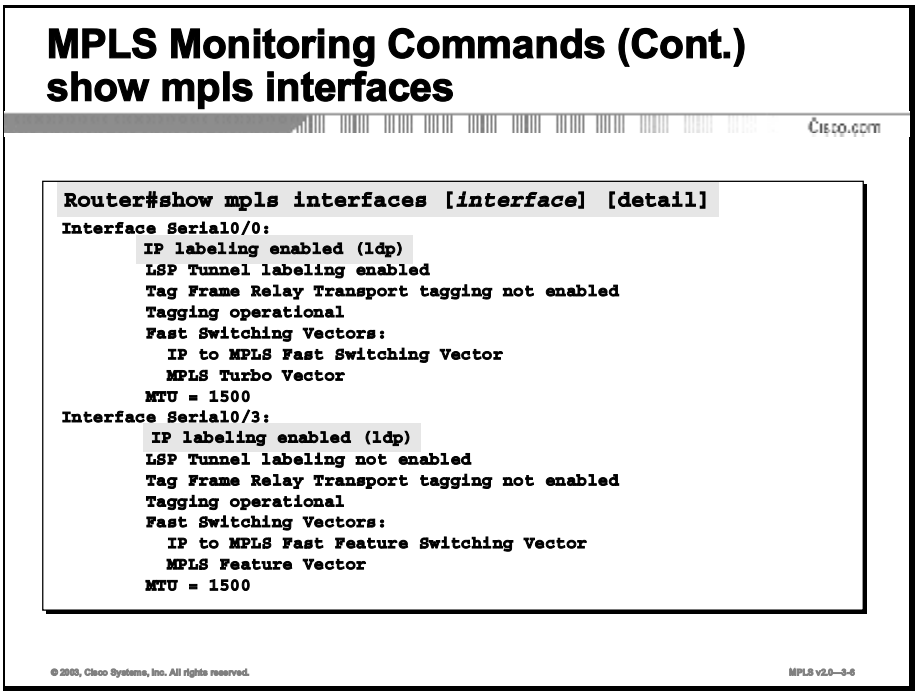

The **show mpls interfaces** command will show only those interfaces on which MPLS has been configured.

show mpls interfaces

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

- **show mpls interfaces** [**all**]

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.

The following table describes the significant fields in the display.

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 required Layer 2 negotiations have occurred.

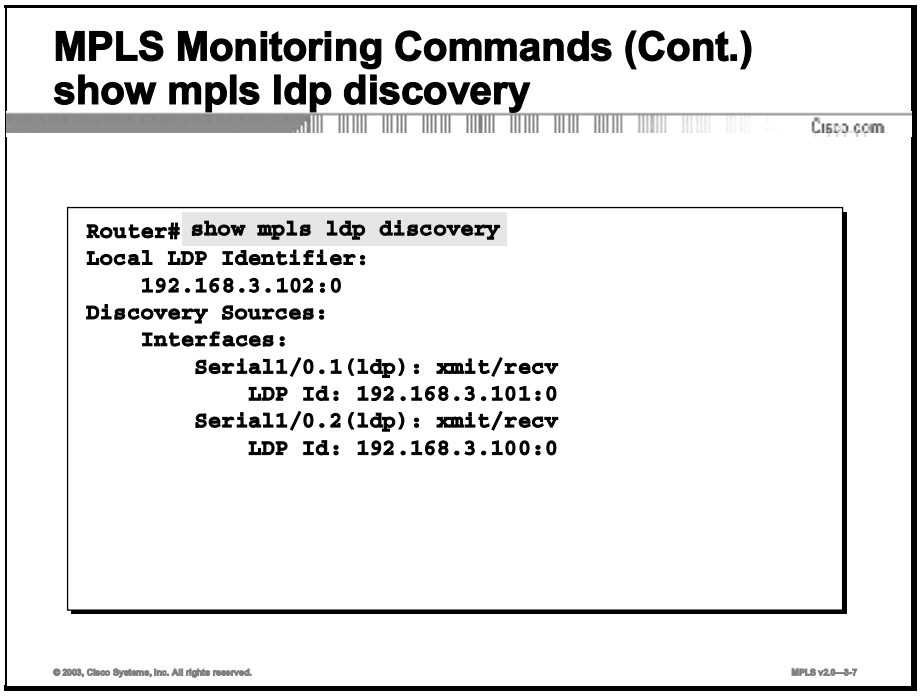

show mpls ldp discovery

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:

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

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) When the **all** keyword is specified alone in this command, the command displays LDP discovery information for all VPNs, including those in the default routing domain.

The following table describes the significant fields in the display.

Field	Description
`Local LDP Identifier`	The LDP identifier for the local router. An LDP identifier is a 6-byte construct displayed in the form "IP address:number." By convention, the first four bytes of the LDP identifier constitute the router ID; integers, starting with 0, constitute the final two 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.

Note The entry for a given target platform may indicate both active and passive.

LDP Monitoring Commands

This topic describes the commands used to monitor LDP.

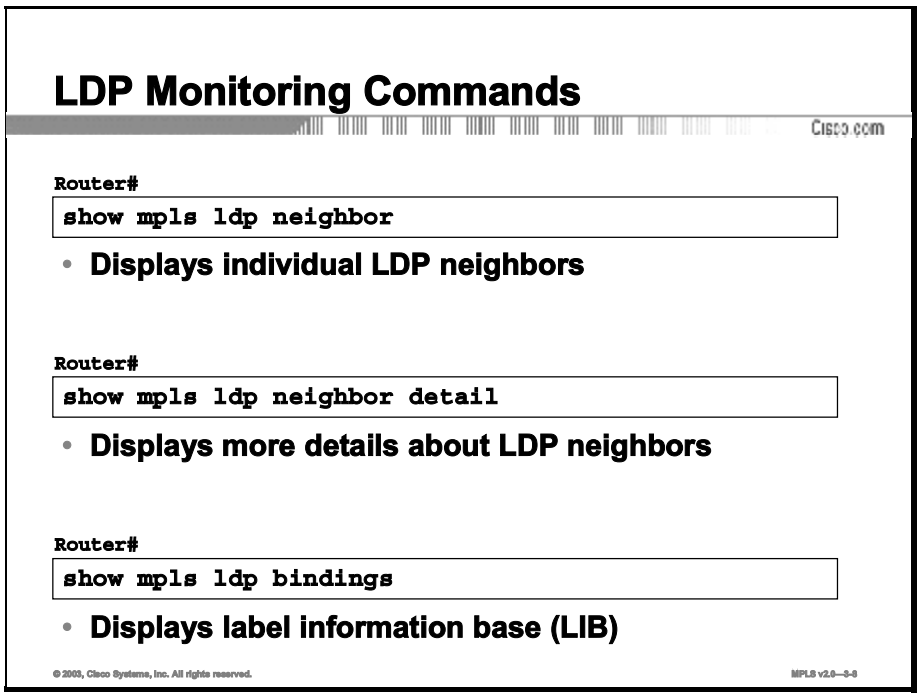

show mpls ldp neighbor

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

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

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) When the **all** keyword is specified alone in this command, the command displays LDP neighbor information for all VPNs, including those in the default routing domain.

show mpls ldp bindings

To display the contents of the LIB, use the **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**]

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.

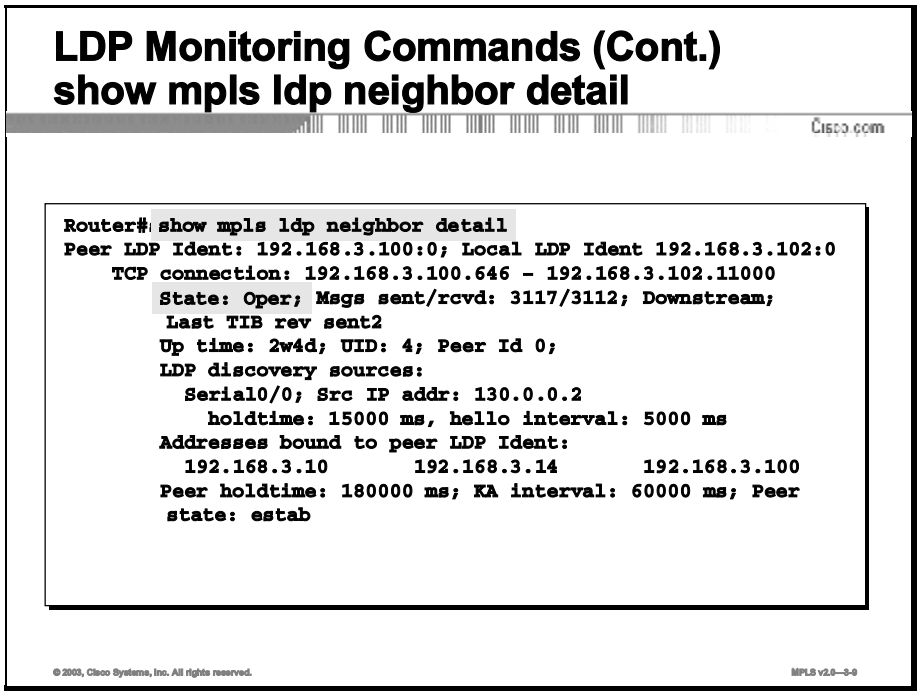

The status of the LDP (TDP) session is indicated by "State: Oper" (operational).

show mpls ldp neighbor

To display the status of LDP sessions, issue the **show mpls ldp neighbor** command 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

The following table describes the significant fields in the display.

Field	Description
Peer LDP Ident	LDP identifier of the neighbor (peer) for this session.
Local LDP Ident	LDP identifier for the local LSR for this session.
TCP connection	TCP connection used to support the LDP session, shown in the following format: ■ peer IP address.peer port ■ local IP address.local port
State	State of the LDP session. Generally, this is “Oper” (operational), but “transient” is another possible state.
Msgs sent/rcvd	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	Length of time that the LDP session has existed.
LDP discovery sources	Source(s) of LDP discovery activity that led to the establishment of this LDP session.
Addresses bound to peer LDP Ident	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.

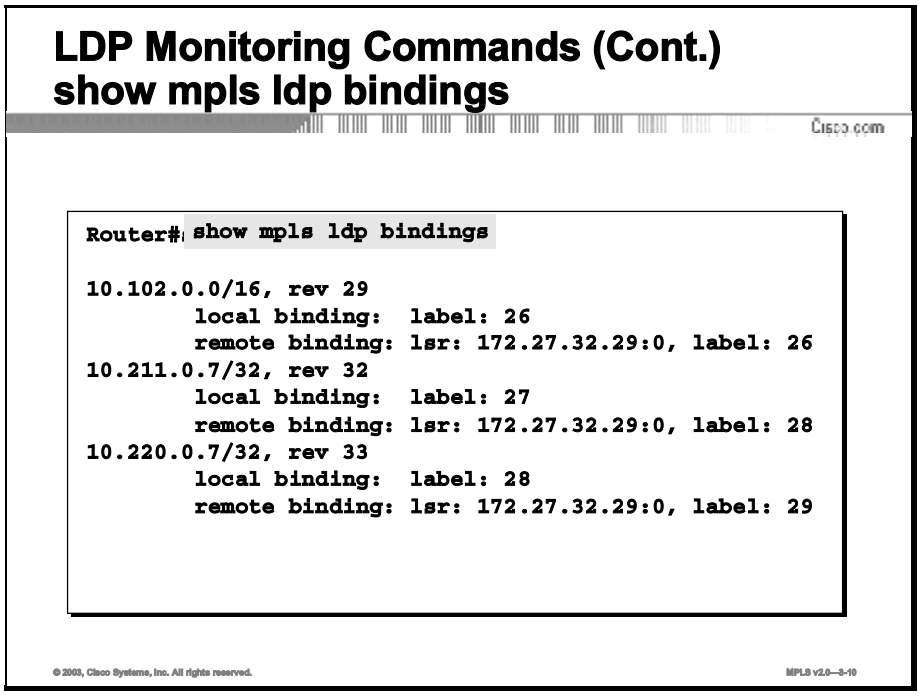

show mpls ldp bindings

To display the contents of the LIB, use the **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**]

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

The following 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 discusses commands used to monitor the LFIB and CEF tables.

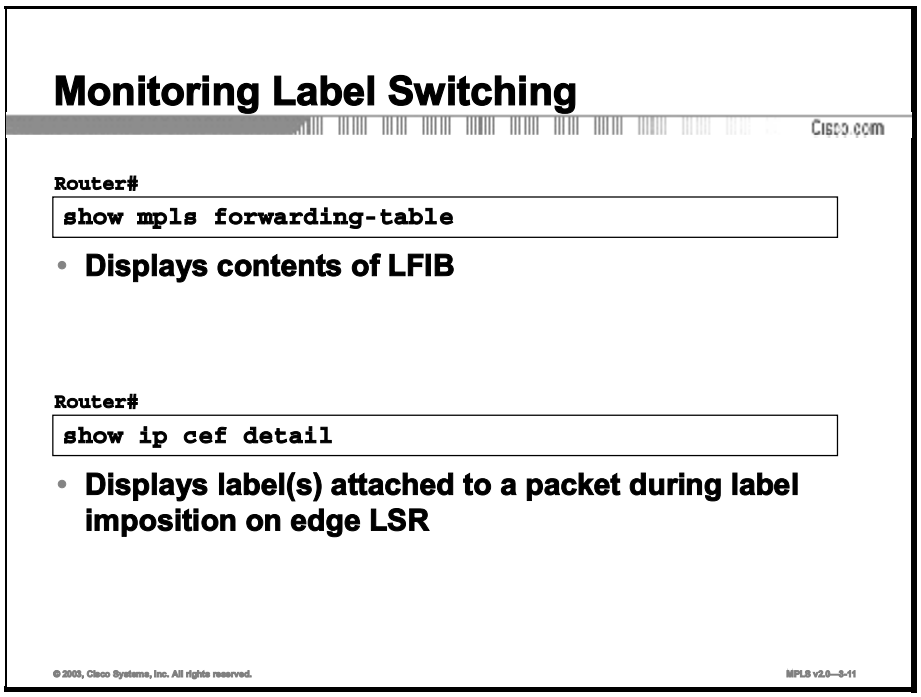

show mpls forwarding-table

To display the contents of the MPLS LFIB, use the **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 this 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 this 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 this form of the **show ip cef** in privileged EXEC mode:

- **show ip cef** [*type number*] [**detail**]

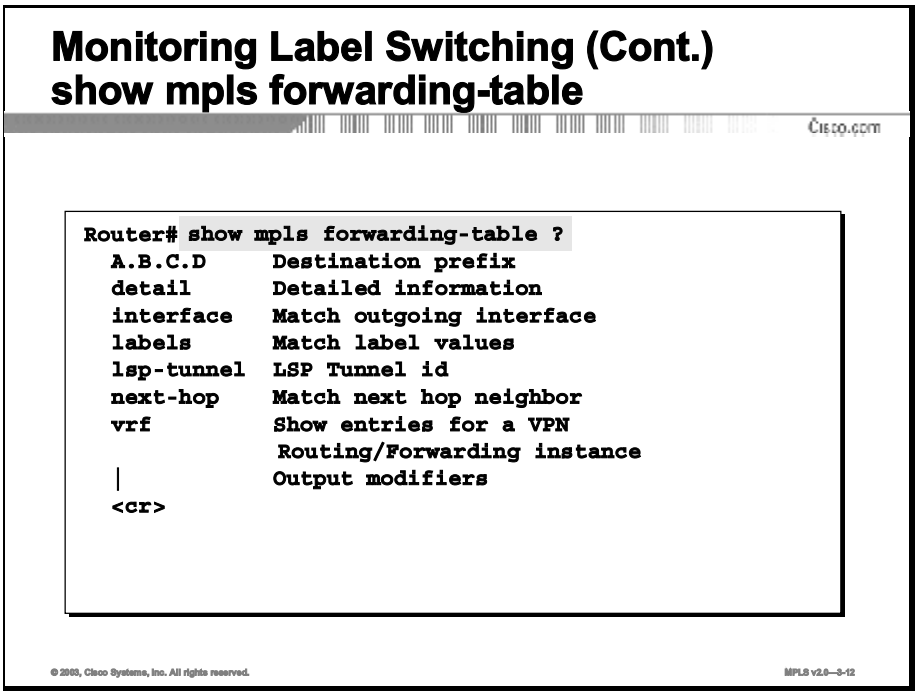

show mpls forwarding-table

To display the contents of the MPLS LFIB, use the **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**]

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

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

```
Router#show mpls forwarding-table

Local  Outgoing               Prefix                  Bytes
tag  Outgoing     Next Hop

tag    tag or VC or Tunnel Id             switched   interface

26     Untagged               10.253.0.0/16           0
Et4/0/0      172.27.232.6

28     1/33             10.15.0.0/16            0
AT0/0.1      point2point

29     Pop tag                10.91.0.0/16            0
Hs5/0        point2point

       1/36             10.91.0.0/16            0
AT0/0.1      point2point

30     32                     10.250.0.97/32          0
Et4/0/2      10.92.0.7

       32               10.250.0.97/32          0
Hs5/0        point2point

34     26               10.77.0.0/24            0
Et4/0/2      point2point

       26               10.77.0.0/24            0
Hs5/0        point2point

35     Untagged[T]            10.100.100.101/32       0
Tu1          point2point

36     Pop tag          168.1.0.0/16            0
Hs5/0        point2point

       1/37                   168.1.0.0/16            0
AT0/0.1      point2point
```

[T] = Forwarding through a LSP tunnel.

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

Monitoring Label Switching (Cont.)
show mpls forwarding-table detail

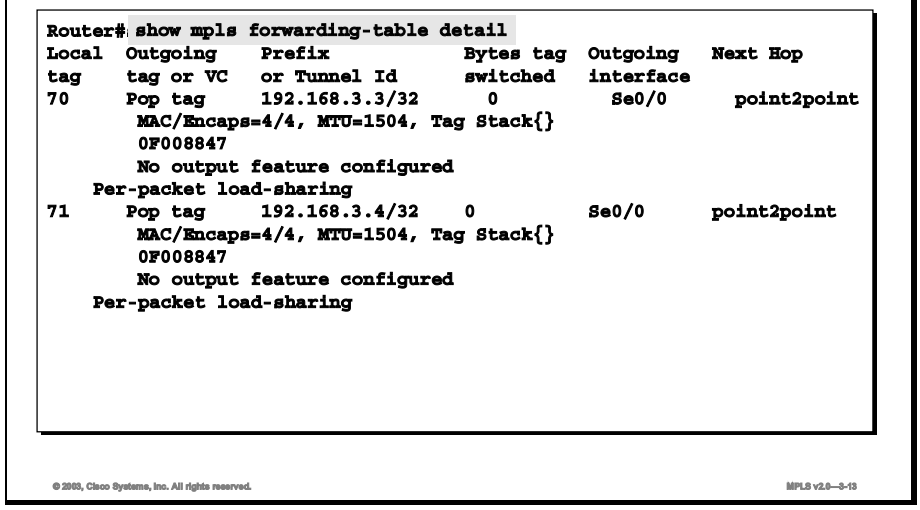

The following table describes the significant fields in the display.

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

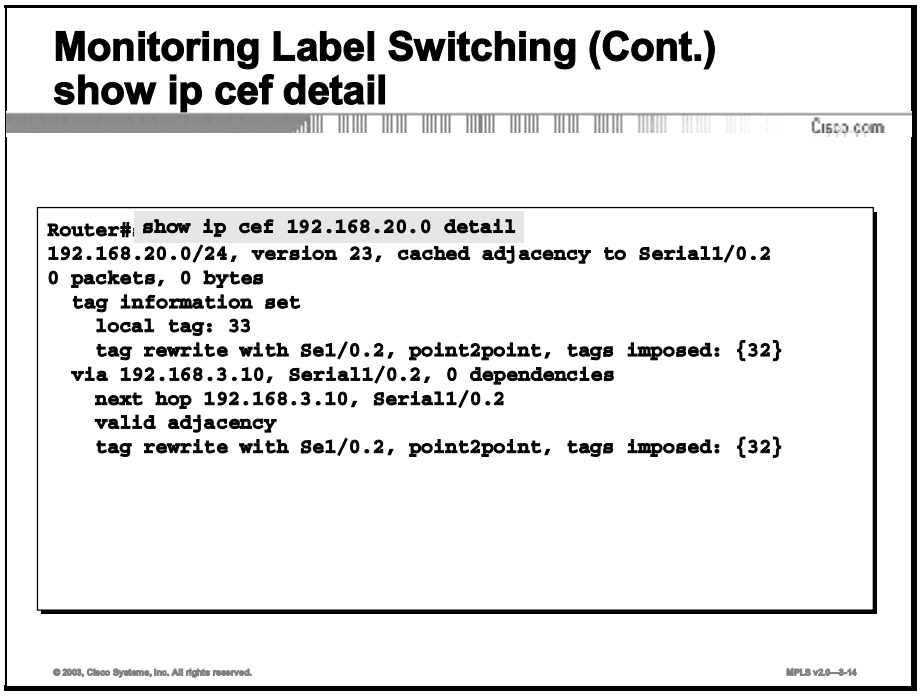

show ip cef detail

To display detailed FIB entry information for all FIB entries, use the **show ip cef detail** command in privileged EXEC mode:

- **show ip cef** [*type number*] [**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) 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 commands that are used to debug MPLS and LDP.

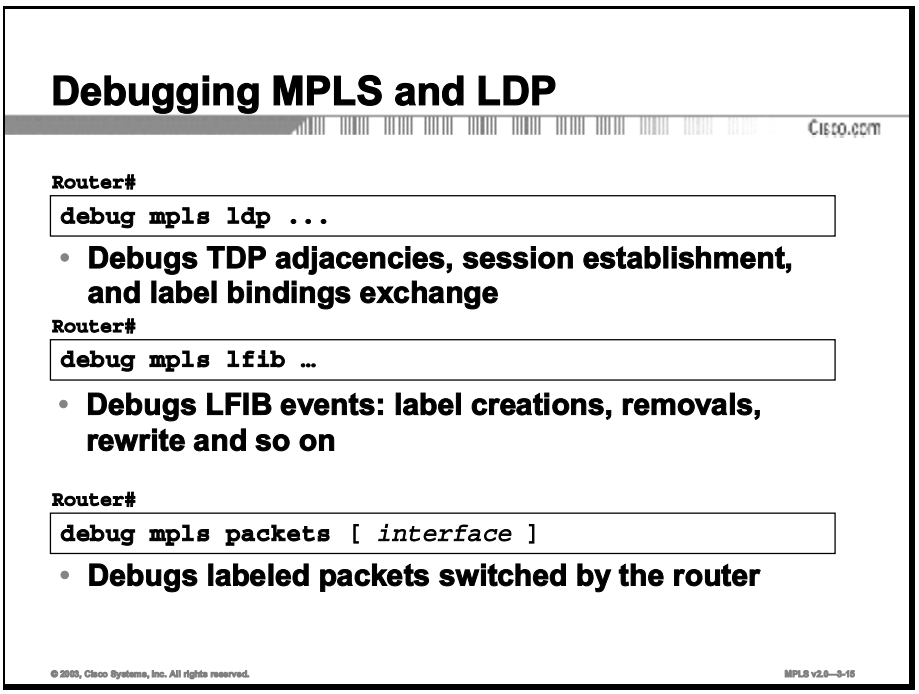

A large number of debug commands are associated with MPLS on Cisco IOS platforms. The **debug mpls ldp** set of commands debug various aspects of LDP protocol, from label distribution to exchange of the application-layer data between adjacent LDP-speaking routers.

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** in privileged EXEC mode. To disable debugging output, use the **no** form of this command:

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

Syntax Description

Field	Description
Hs0/0	The identifier for the interface on which the packet was received or transmitted.
Recvd	Packet received.
Xmit	Packet transmitted.
CoS	Class of service field from the packet label header.
TTL	Time-to-live field from the packet label header.
(no tag)	Last label popped off the packet and transmitted unlabeled.
Tag(s)	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.

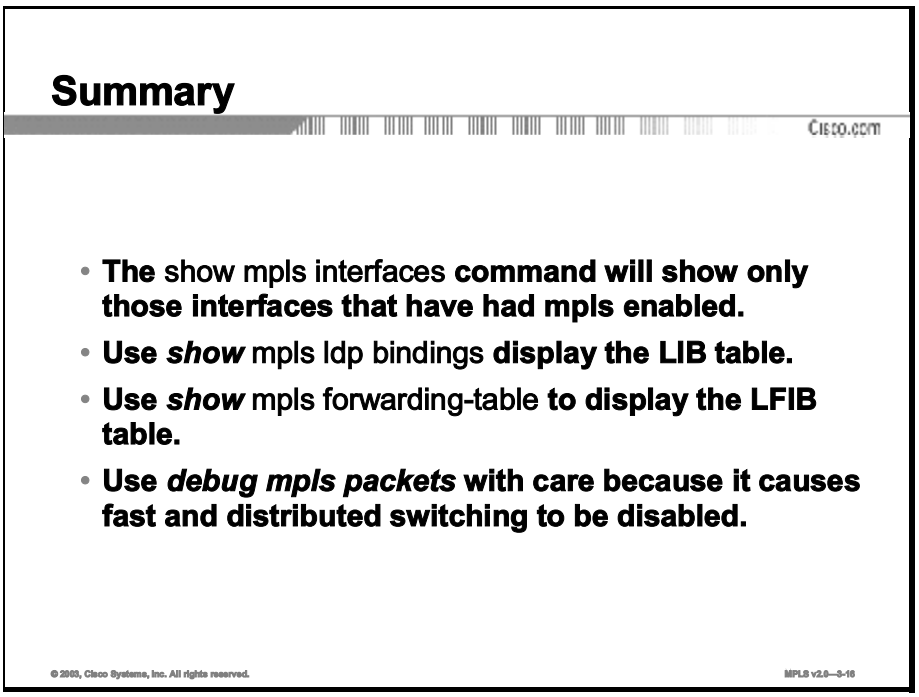

References

For additional information, refer to this resource:

- Search www.cisco.com for additional resources covering the topics discussed in this lesson.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which command is used to display information about the LDP hello protocol?

A) **show ip cef**

B) **show mpls ldp parameters**

C) **show ldp forwarding-table**

D) **show mpls ldp discovery**

Q2) Which command is used to display the contents of the LIB table?

A) **show mpls ldp labels**

B) **show mpls ldp bindings**

C) **show mpls ldp neighbors**

D) **show mpls forwarding-table**

Q3) Which command is used to display the contents of the LFIB table?

A) **show mpls ldp labels**

B) **show mpls ldp bindings**

C) **show mpls ldp neighbors**

D) **show mpls forwarding-table**

Q4) Which of the following commands would NOT be used to debug MPLS or LDP?

A) **debug mpls ldp**

B) **debug mpls lfib**

C) **debug mpls packets**

D) **debug mpls ldp neighbors**

Quiz Answer Key

Q1) D

Relates to: MPLS Monitoring Commands

Q2) B

Relates to: LDP Monitoring Commands

Q3) D

Relates to: Monitoring Label Switching

Q4) D

Relates to: Debugging MPLS and LDP

Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms

Overview

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

Relevance

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

This lesson describes how to troubleshoot frame-mode MPLS problems on Cisco IOS platforms.

Upon completing this lesson, you will be able to:

- Identify the common problem issues that are found in MPLS
- Identify and troubleshoot session startup issues
- Identify and troubleshoot label allocation issues
- Identify and troubleshoot label distribution issues
- Identify and troubleshoot packet labeling issues
- Identify and troubleshoot intermittent MPLS failures after interface failure
- Identify and troubleshoot packet propagation issues

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Monitoring Frame-Mode MPLS on Cisco IOS Platforms” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Common Frame-Mode MPLS Issues
- LDP Session Startup Issues
- Label Allocation Issues
- Label Distribution Issues
- Packet-Labeling Issues
- Intermittent MPLS Failures After Interface Failure
- Packet Propagation Issues
- Summary
- Quiz

Common Frame-Mode MPLS Issues

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

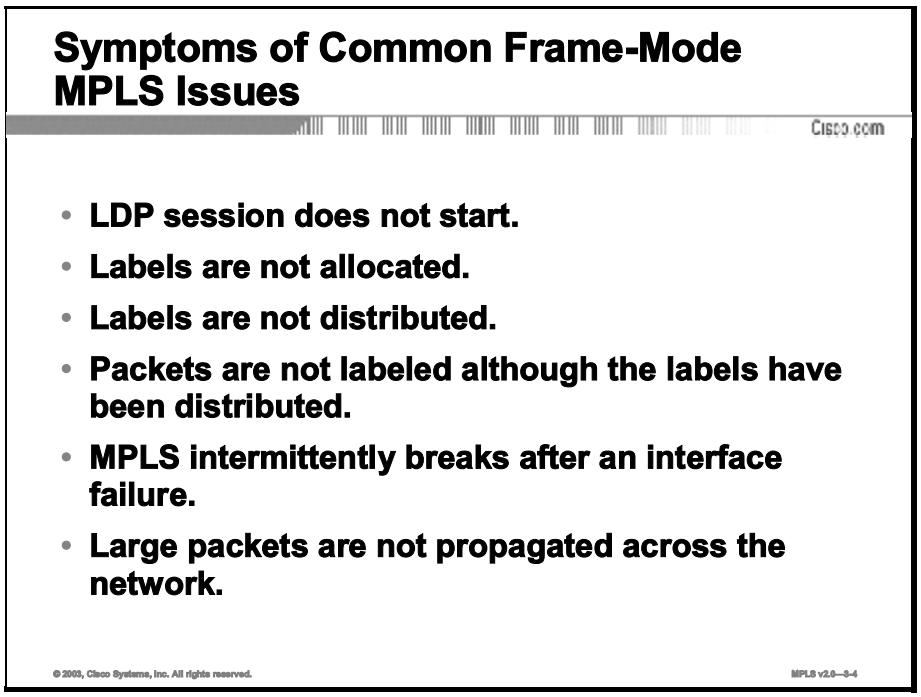

A number of common issues can be encountered while you are troubleshooting a frame-mode MPLS network:

- 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 they were successfully propagated across the pure IP backbone.

The following topics describe each of these issues and provide recommended steps for troubleshooting them.

LDP Session Startup Issues

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

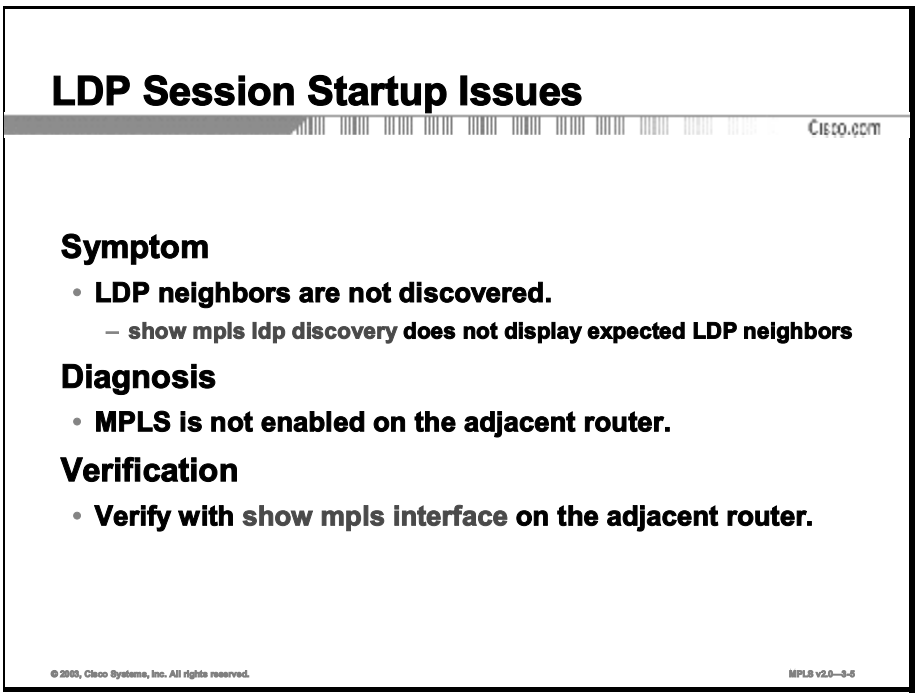

Diagnosis: If MPLS is enabled on an interface, but no neighbors are discovered, it is likely that MPLS is not enabled on the neighbor.

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.

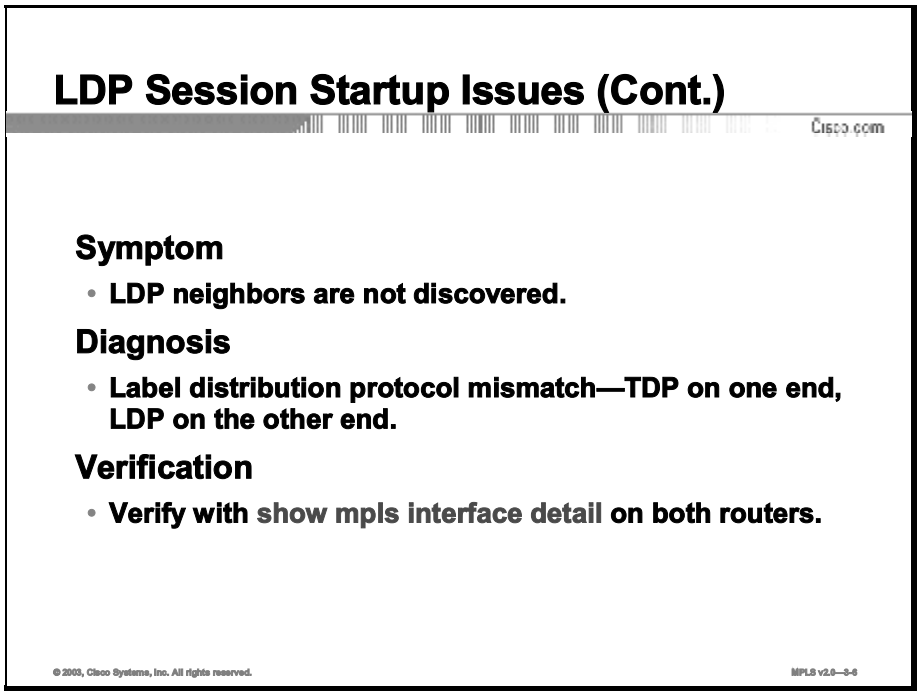

Diagnosis: Another possibility is that the neighbor has a different label distribution protocol enabled on the interface.

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.

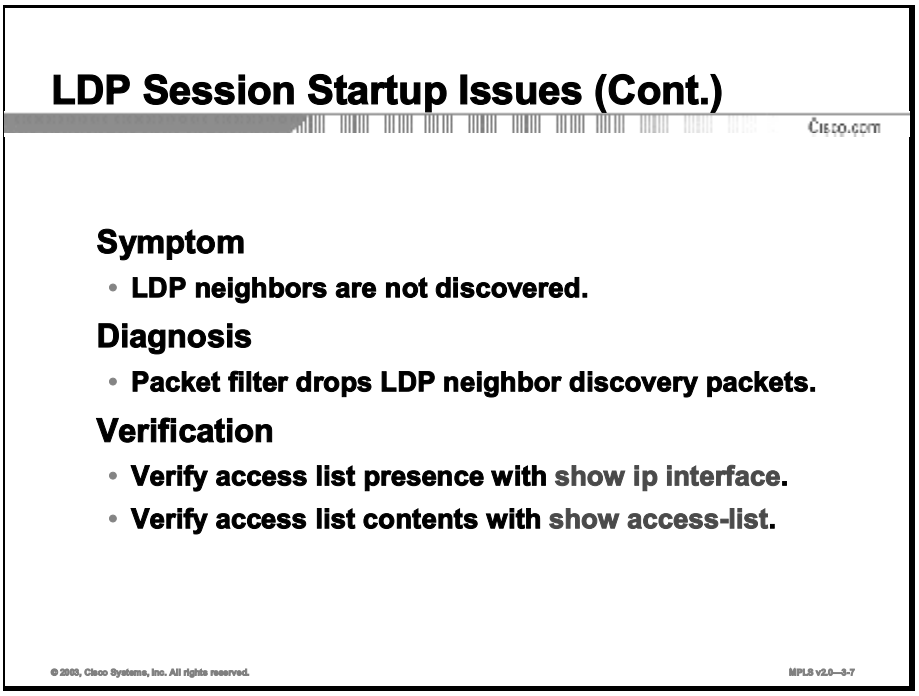

Diagnosis: MPLS configurations match on both ends, but the session still does not get established. Check whether there are any input access lists that deny discovery messages.

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).

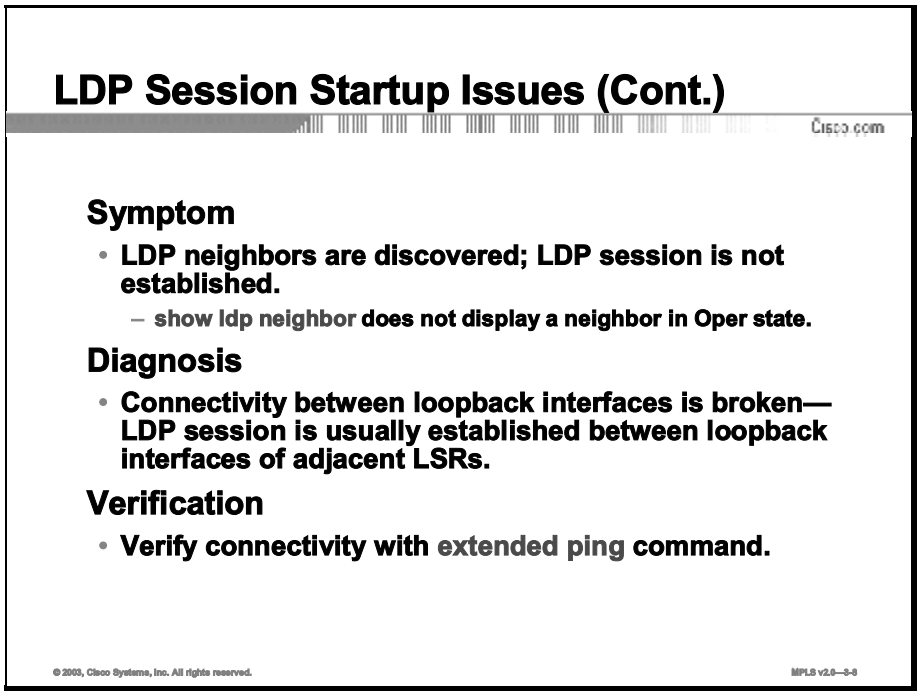

Diagnosis: LDP neighbors are exchanging hello packets, but the LDP session is never established.

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.

Label Allocation Issues

This topic describes issues that could arise in the allocation of labels in MPLS networks.

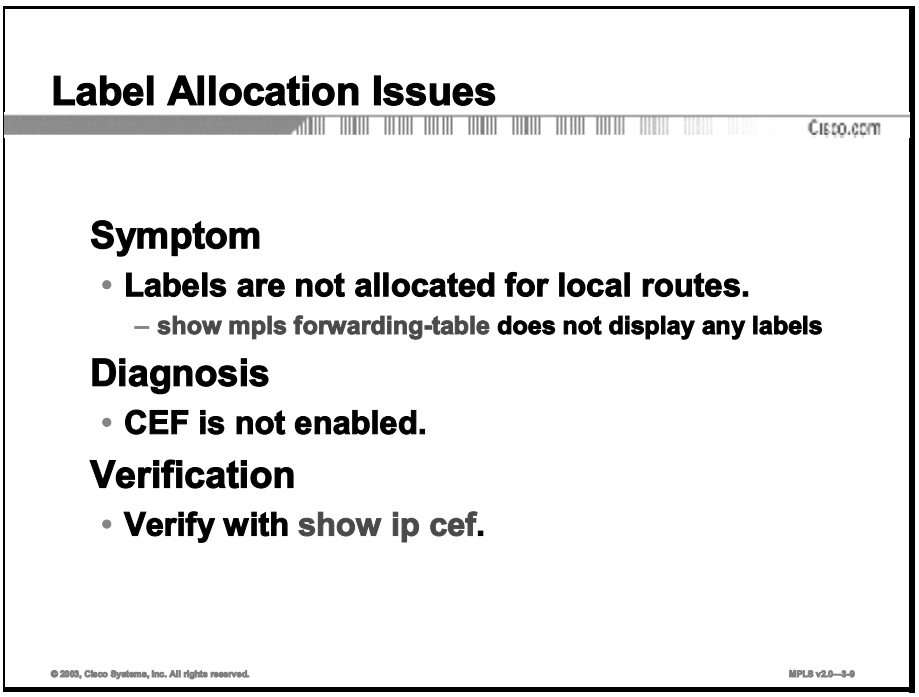

Diagnosis: Labels are not allocated for any or some of the local routes. Use the **show ip cef** command to verify whether CEF switching is enabled on all MPLS-enabled interfaces.

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

Label Distribution Issues

This topic describes issues that could arise in the distribution of labels in MPLS networks.

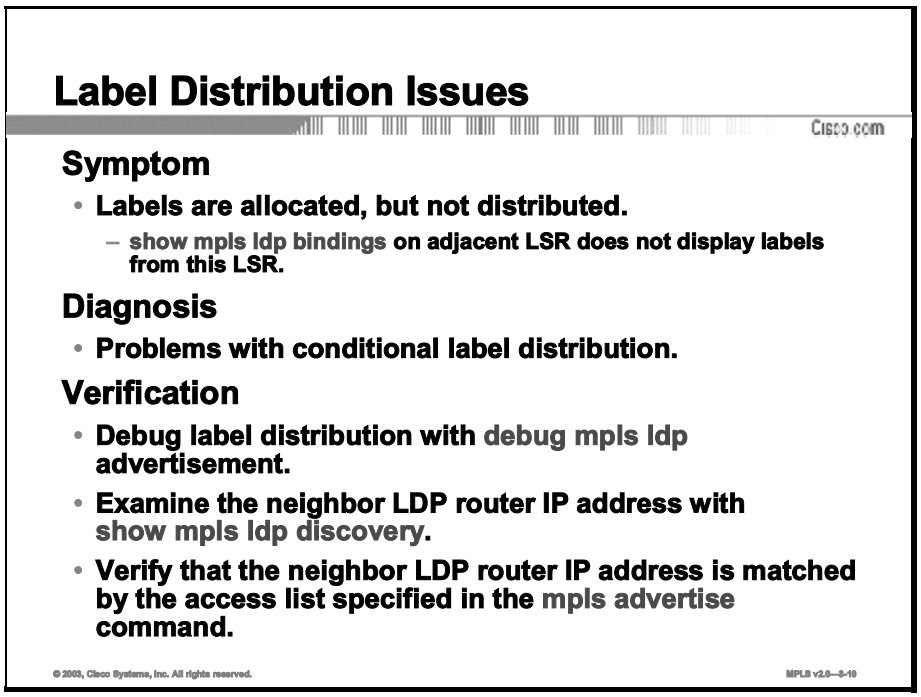

Symptom: Labels are generated for local routes but are not received on neighboring routers.

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

Packet-Labeling Issues

This topic describes issues that can arise in packet labeling in MPLS networks.

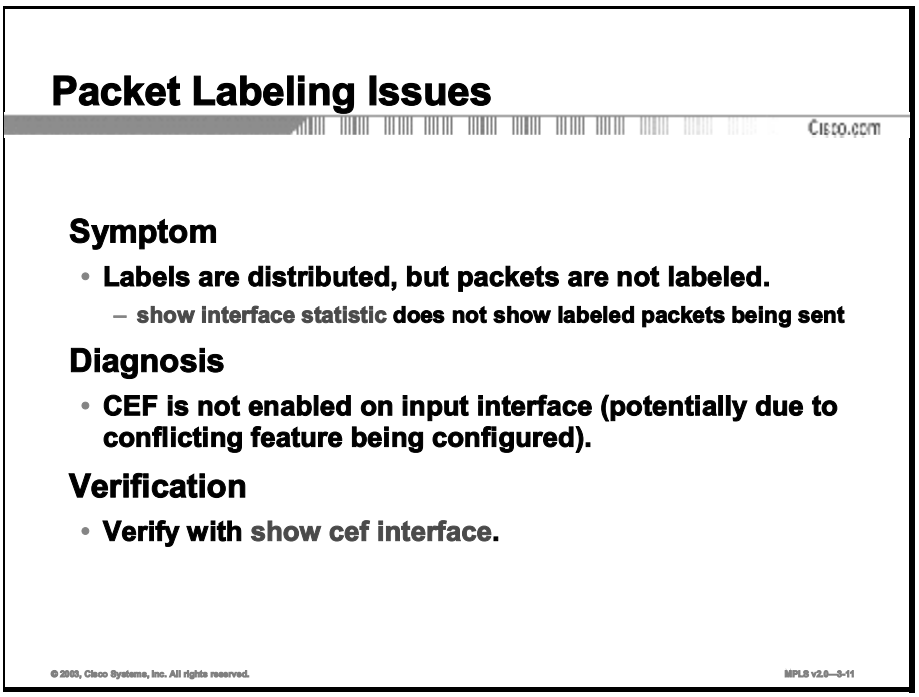

Symptom: Labels exist, but packets are not labeled.

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.

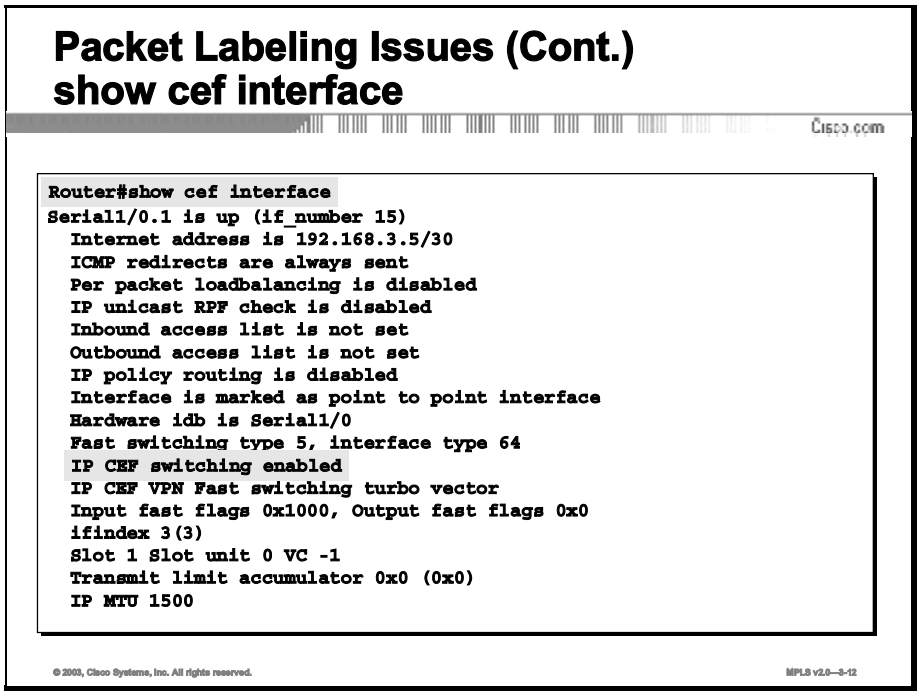

show cef interface

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

- **show cef interface** *type number* [**detail**]

Syntax Description

Parameter	Description
type number	Interface number and 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.

The following table describes the significant fields in the display.

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

Intermittent MPLS Failures After Interface Failure

This topic describes intermittent issues that can arise after an interface failure in MPLS networks.

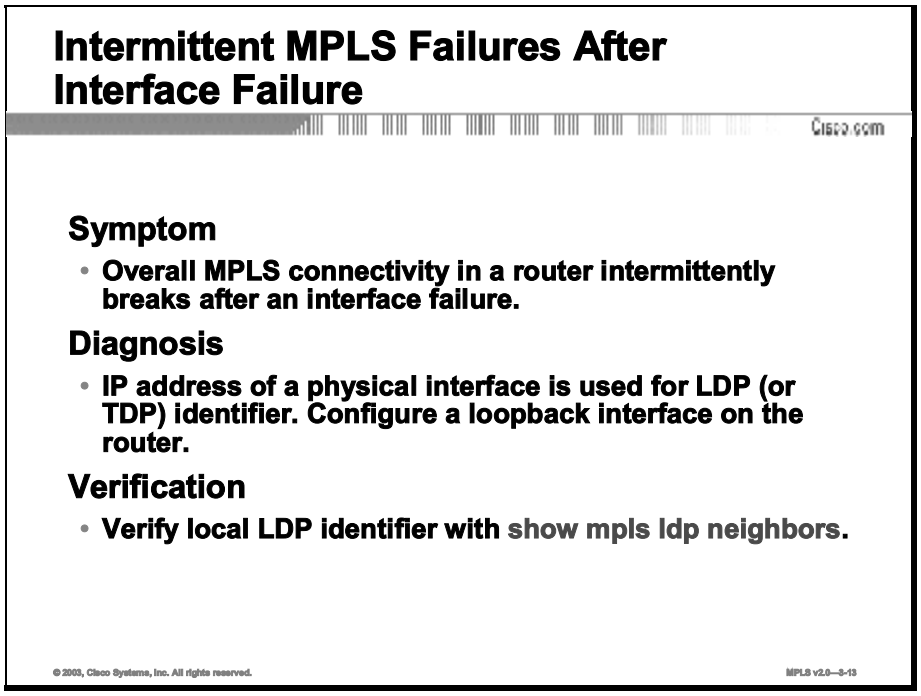

Symptom: MPLS connectivity is established, labels are exchanged, and packets are labeled and forwarded. However, an interface failure can sporadically stop an MPLS operation on unrelated interfaces in the same router.

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.

Packet Propagation Issues

This topic describes possible packet propagation issues in an MPLS network.

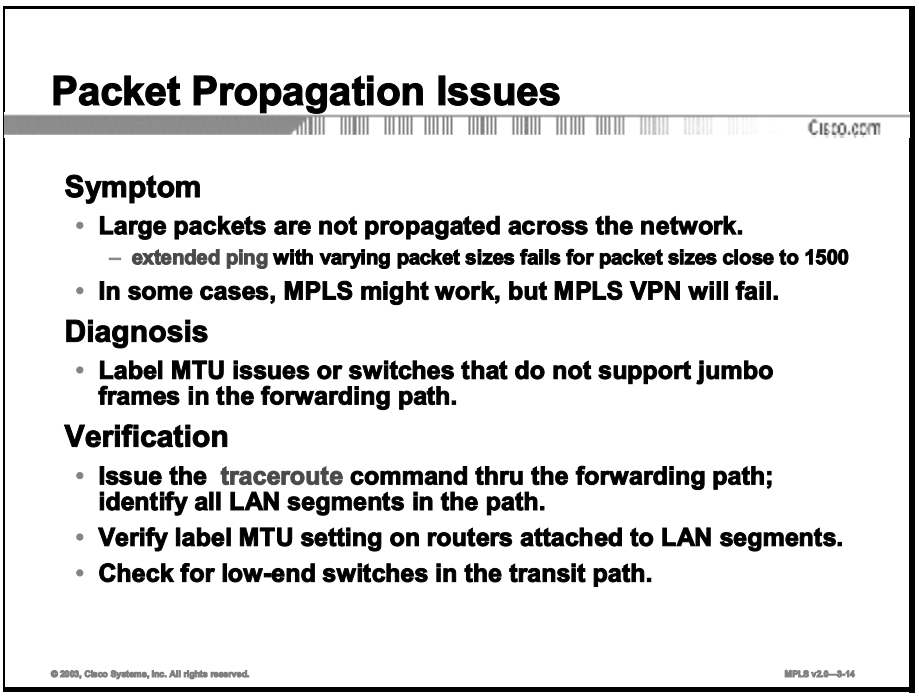

Symptom: Packets are labeled and sent, but they are not received on the neighboring router. A LAN switch between the adjacent MPLS-enabled routers may drop the packets if it does not support jumbo frames.

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.

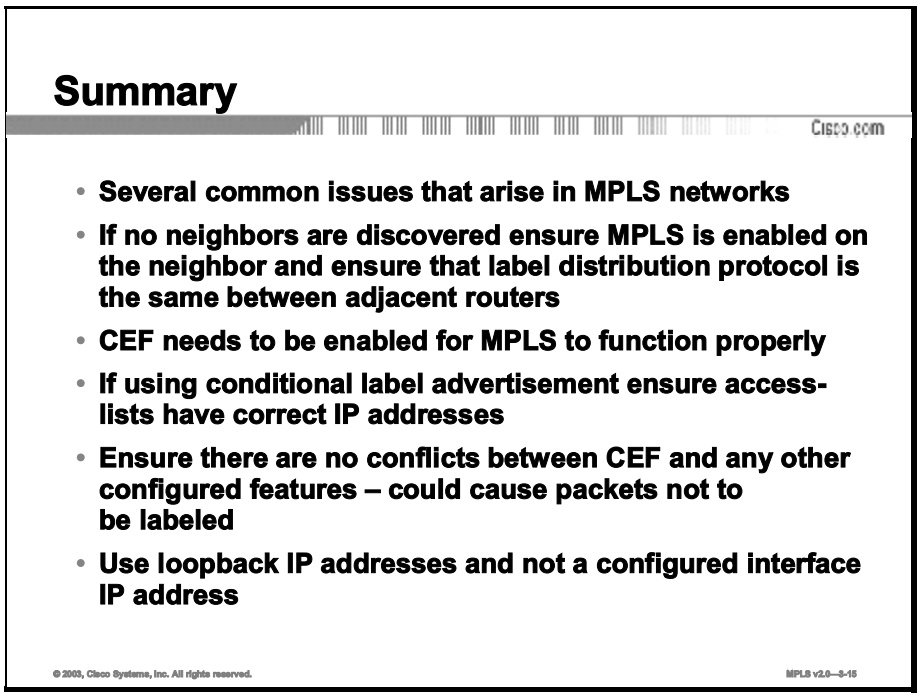

References

For additional information, refer to this resource:

- Search www.cisco.com for additional resources covering the topics discussed in this lesson.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two of the following would cause an LDP (or TDP) session not to be established between two LSRs? (Choose two.)

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

Q2) Which command is issued to troubleshoot label allocation issues?

A) **show cef**

B) **show lfib**

C) **show ip cef**

D) **show mpls lfib**

Q3) Which command is issued to see if labels are being distributed from the local LSR?

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

Q4) Which of the following is correct?

A) router>show cef interface

B) **router#show cef interface**

C) **router(config)#show cef interface**

D) **router(config-router))#show cef interface**

Q5) To reduce the chances of having intermittent MPLS failures because of an interface failing, a _____ address should be configured.

Q6) 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:

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.

Quiz Answer Key

Q1) C, D

Relates to: LDP Session Startup Issues

Q2) C

Relates to: Label Allocation Issues

Q3) D

Relates to: Label Distribution Issues

Q4) B

Relates to: Packet-Labeling Issues

Q5) loopback

Relates to: Intermittent MPLS Failures After Interface Failure

Q6) C

Relates to: Packet Propagation Issues

Configuring LC-ATM MPLS

Overview

This lesson explains how to configure MPLS on router label–controlled ATM (LC-ATM) interfaces and Cisco IOS software-based ATM switches. It presents configuration tasks, syntax definitions, and example configurations.

Relevance

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

This lesson describes how to configure LC-ATM MPLS Cisco IOS platforms.

Upon completing this lesson, you will be able to:

- Describe 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 basic LC-ATM configuration between a router and a Catalyst ATM switch
- Describe what additional LC-ATM parameters can be configured
- Describe how to perform the configuration to disable VC merge

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Troubleshooting Frame-Mode MPLS on Cisco IOS Platforms” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Configuration Tasks for MPLS on LC-ATM Interfaces
- Configuring an LC-ATM Interface on a Router
- Configuring an LC-ATM Interface on a Catalyst ATM Switch
- Basic LC-ATM Configuration
- Configuring Additional LC-ATM Parameters
- Disabling VC Merge
- Summary
- Quiz

Configuration Tasks for MPLS on LC-ATM Interfaces

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

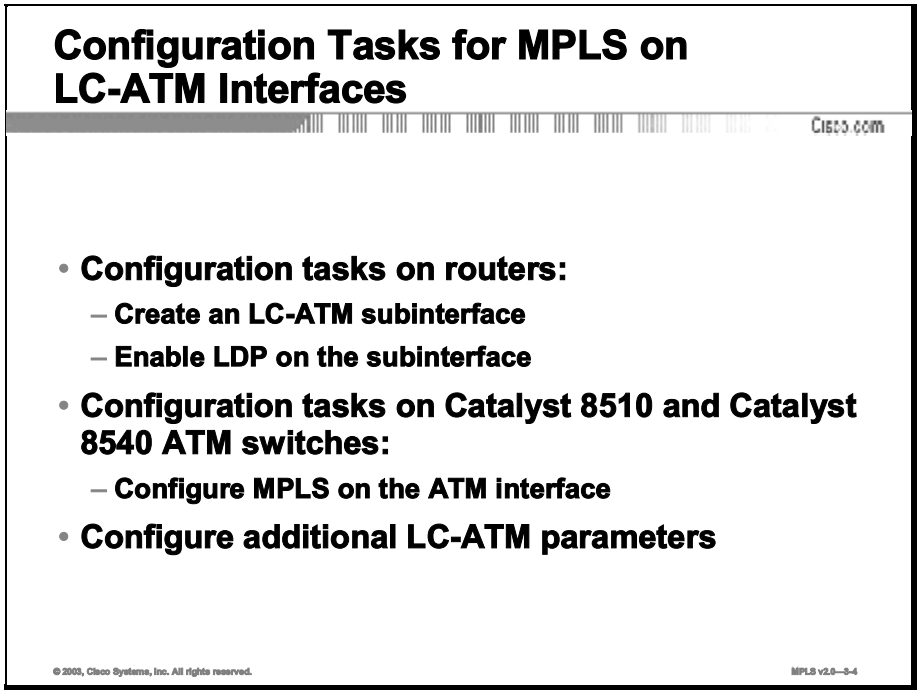

Configuration of cell-mode MPLS differs from configuration of frame-mode MPLS. An additional command specifies the type of subinterface that is to be used.

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).

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.

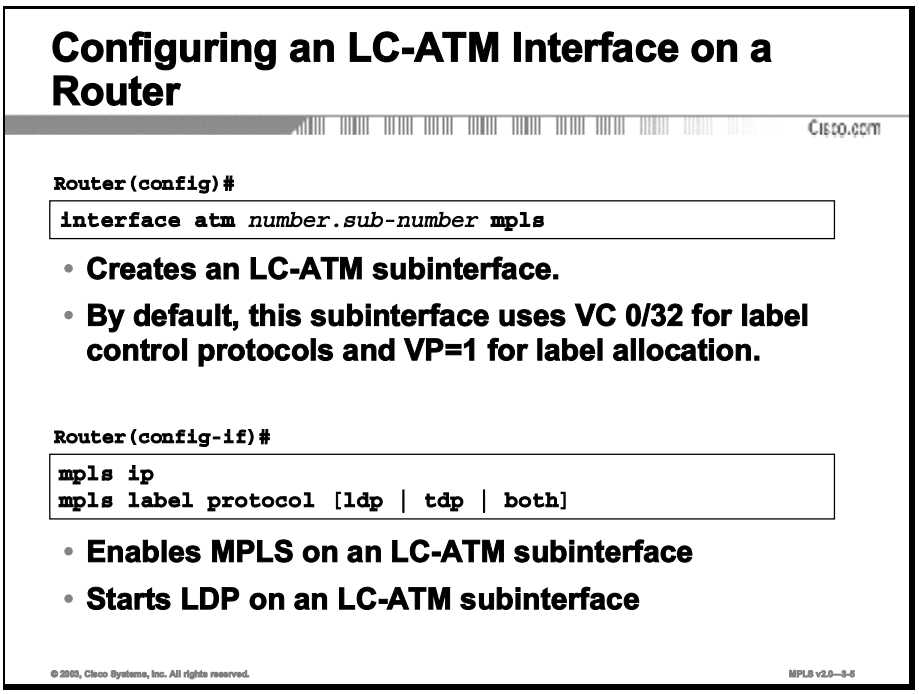

On Cisco IOS platform routers, subinterfaces are typically used. Use the **mpls** keyword to specify the type of subinterface when you are entering interface configuration mode. This command specifies that cell-mode MPLS should be used instead of frame-mode (which is the default).

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:

- **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:

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

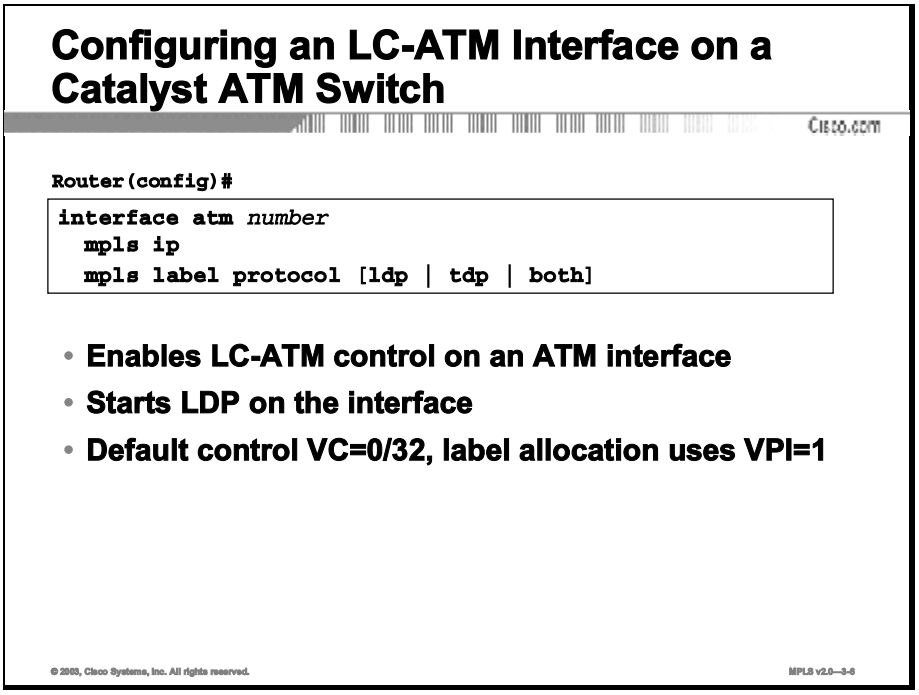

Use these commands to enable MPLS on an interface of a Catalyst ATM switch. Cell-mode MPLS is implied. Enabling both distribution protocols can be useful in a mixed environment when the supported protocol for every device connected to the switch does not need to be determined.

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 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:

- **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:

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

Basic LC-ATM Configuration

This topic describes the basic steps of configuring MPLS between a router and a switch.

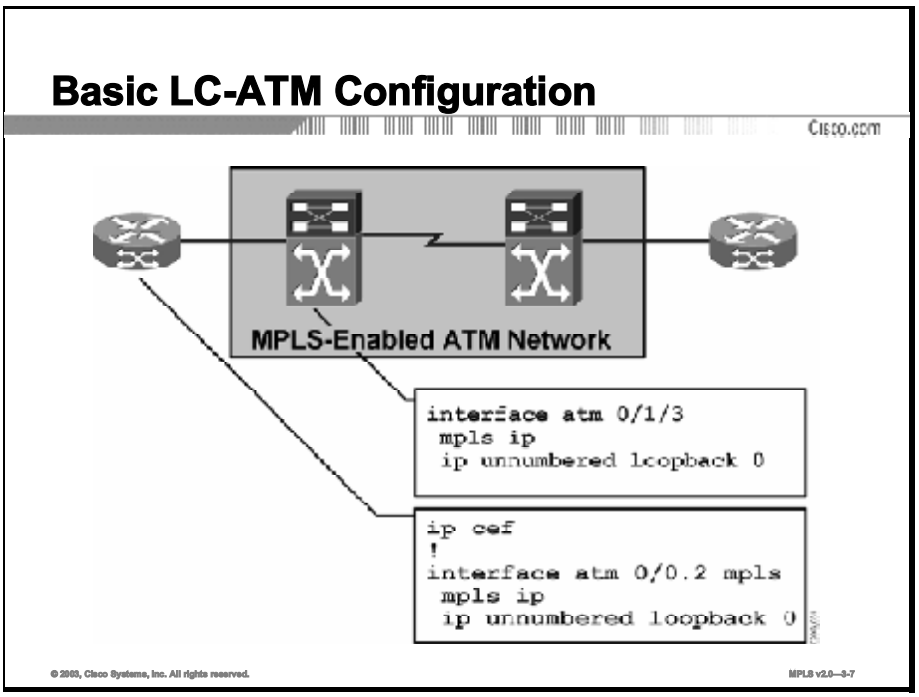

To enable cell-mode MPLS between a router and a switch, ensure that the router uses the MPLS type for the subinterface.

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 additional LC-ATM parameters that can be configured.

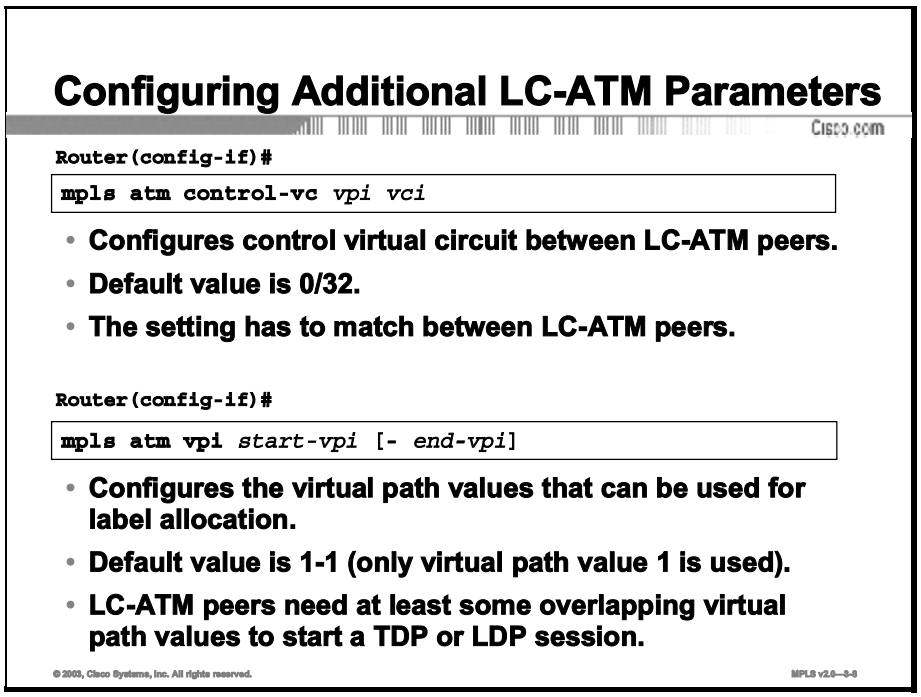

Use the **mpls atm control-vc** command to change the default VPI and VCI numbers used for the control VC. Use the **mpls atm vpi** command to change the default VPI values for the LVCs.

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:

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

Syntax Description

Parameter	Description
vpi	Virtual path identifier.
vci	Virtual channel identifier.

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:

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

Syntax Description

Parameter	Description
vpi	Virtual path identifier (low end of range).
- vpi	(Optional) Virtual path identifier (high end of range).

Defaults

The default is 1-1.

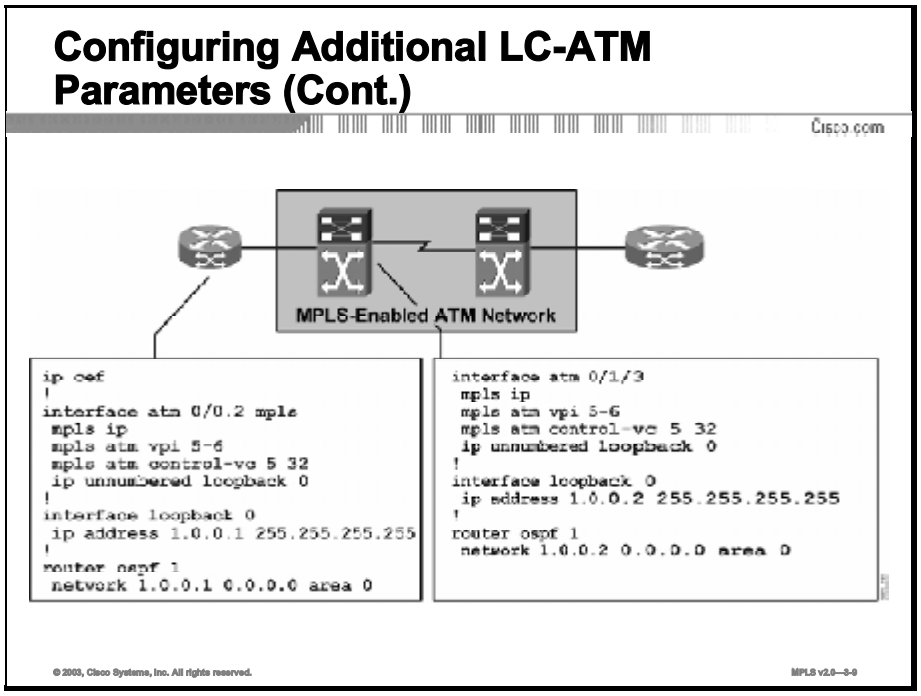

The example shows how to change the default VPI range from 1-1 to 5-6. The control virtual circuit can also use the VPI value used for LVCs.

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

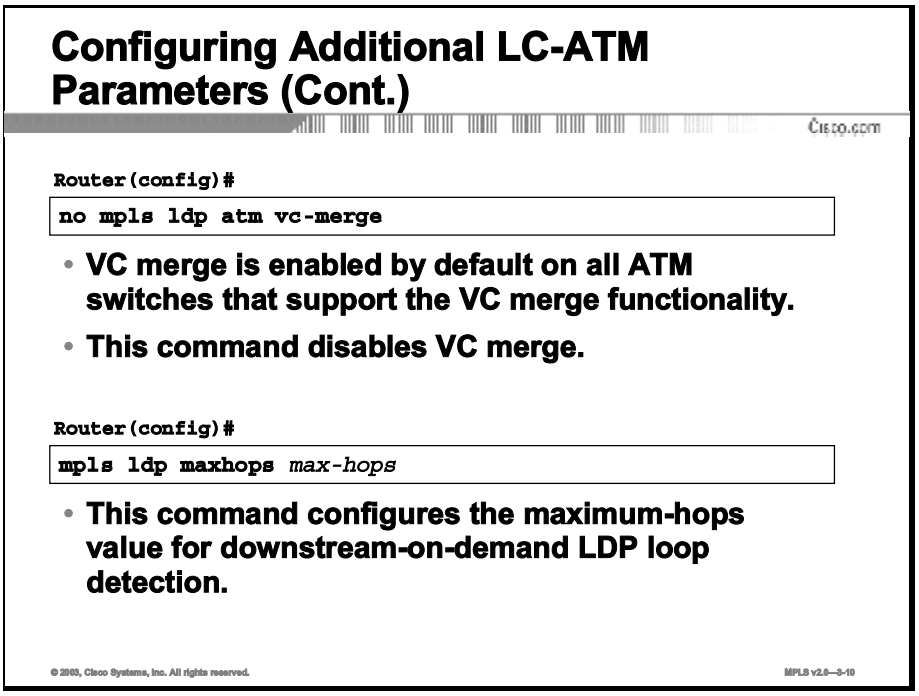

mpls ldp atm vc-merge

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:

- **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 they use the same VPI/VCI value pair. There are 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:

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

Syntax Description

Parameter	Description
number	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, nor does it 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 the configuration needed to disable VC merge.

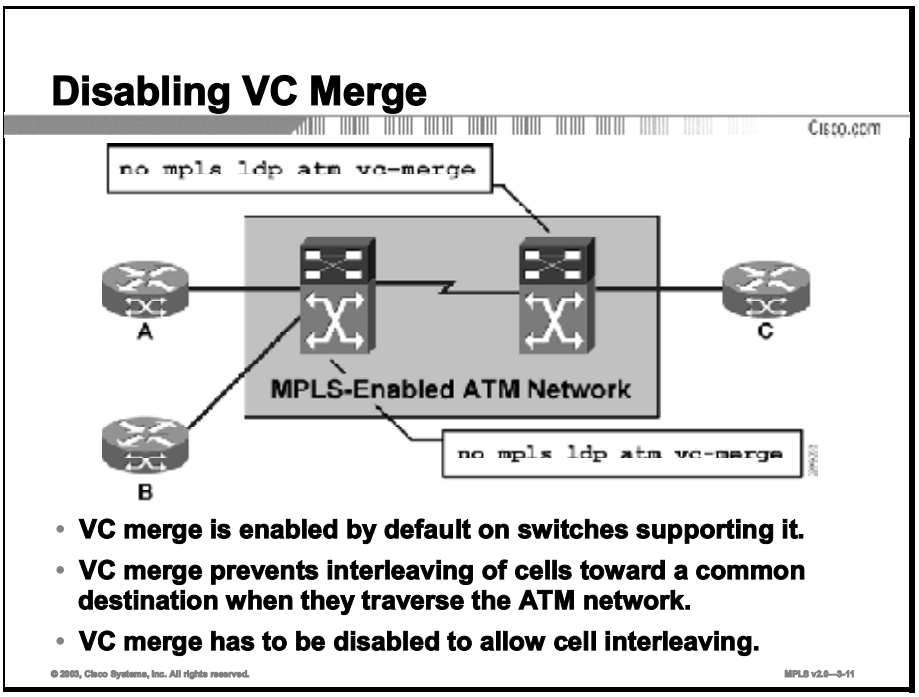

The VC merge feature is enabled by default on all switches that support it. If the feature is not required (that is, small network, different line speeds, buffering not desired), it can be disabled.

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.

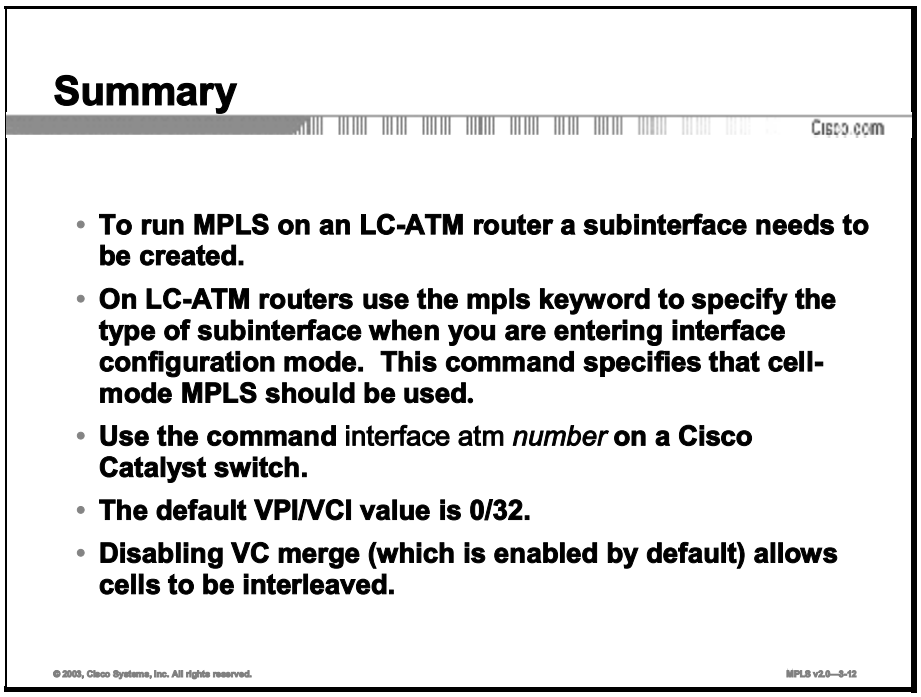

References

For additional information, refer to this resource:

- Search www.cisco.com for additional references for the topics covered in this lesson.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) A _____must be created on an LC-ATM router to support MPLS.

Q2) On Cisco IOS platform routers, _____ mode MPLS is the default.

Q3) By default, all LVCs use which VPI value?

A) 0

B) 1

C) 32

D) 100

Q4) 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?

A) IGP

B) VRI/VDI

C) subinterface number

D) label distribution protocol

Q5) Which command sets the threshold that will prevent forwarding loops in the setting up of label switch paths across an ATM region?

A) **mpls vpi**

B) **mpls atm vpi**

C) **mpls maxhops**

D) **mpls ldp maxhops**

Q6) Which two statements are correct? (Choose two.)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.

Quiz Answer Key

Q1) subinterface

Relates to: Configuration Tasks for MPLS on LC-ATM Interfaces

Q2) frame-

Relates to: Configuring an LC-ATM Interface on a Router

Q3) B

Relates to: Configuring an LC-ATM Interface on a Catalyst ATM Switch

Q4) D

Relates to: Basic LC-ATM Configuration

Q5) D

Relates to: Configuring Additional LC-ATM Parameters

Q6) A, C

Relates to: Disabling VC Merge

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.

Relevance

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

This lesson describes how to configure LC-ATM MPLS over ATM VP.

Upon completing this lesson, you will be able to:

- Describe ATM Virtual Path
- Describe ATM Virtual Path usages
- Describe how to configure MPLS over ATM Virtual Path for switches
- Describe how to configure MPLS over ATM Virtual Path for routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Configuring LC-ATM MPLS” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Introduction to ATM Virtual Path
- ATM Virtual Path Usages
- Configuring MPLS over ATM Virtual Path—Switches
- Configuring MPLS over ATM Virtual Path—Routers
- Summary
- Quiz

Introduction to ATM Virtual Path

This topic introduces ATM VP.

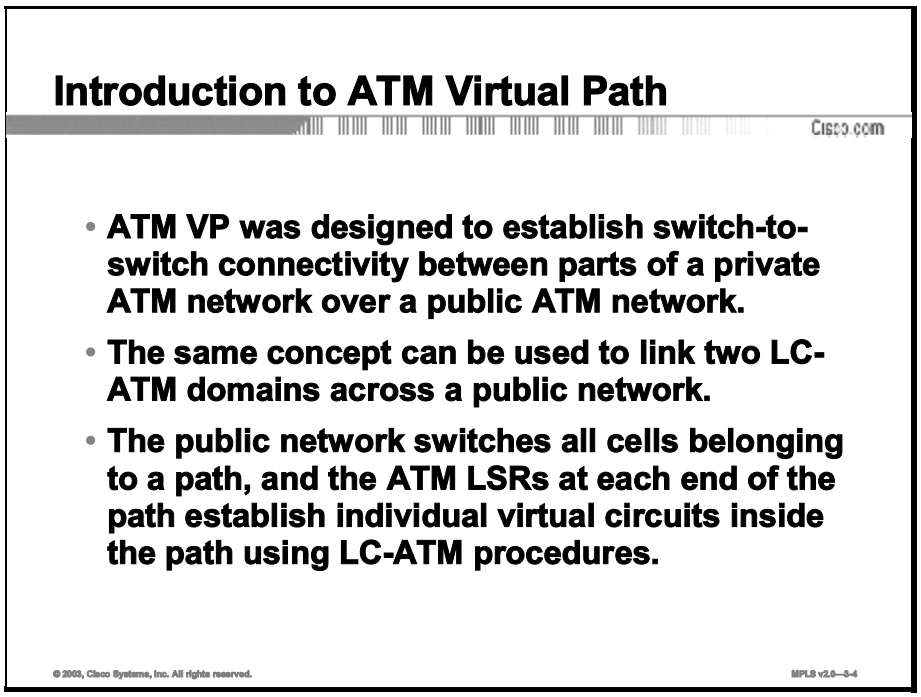

A virtual path is a collection of virtual circuits with a common VPI.

ATM switches forward cells based on the VPI value 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. They 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 Virtual Path can be used.

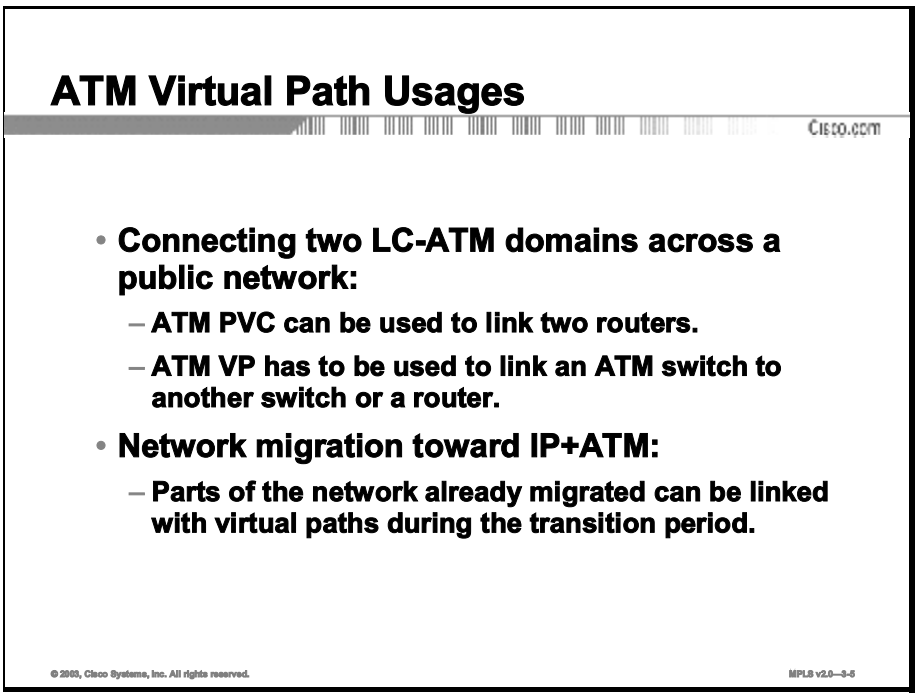

There are two options available to enable two MPLS domains across a public ATM network:

- **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.

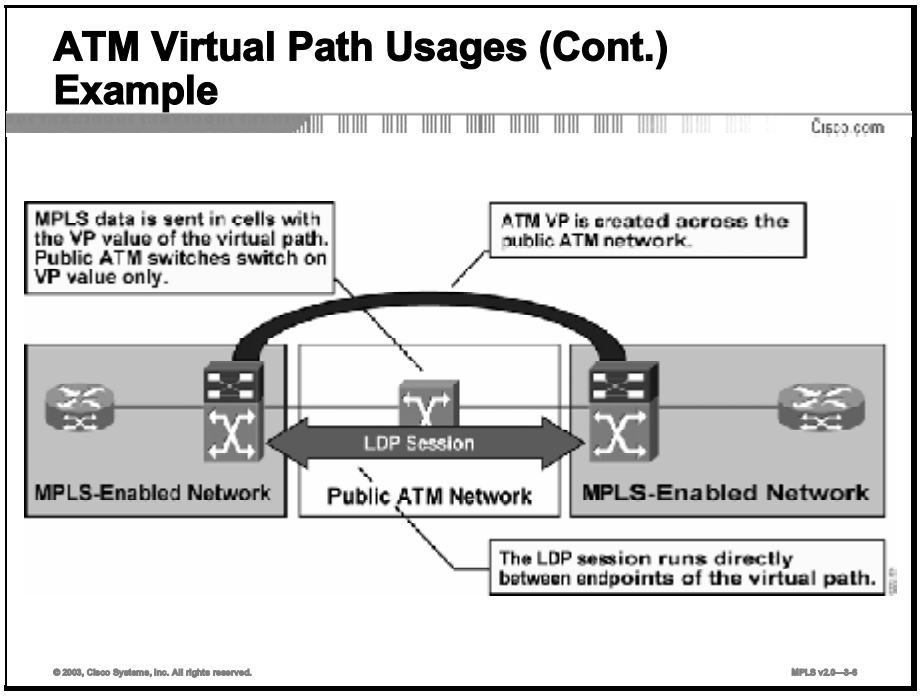

To enable cell-mode MPLS across a virtual path, the control virtual circuit *should use* the VPI of the virtual path.

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 value is part of the LDP virtual path range negotiation.

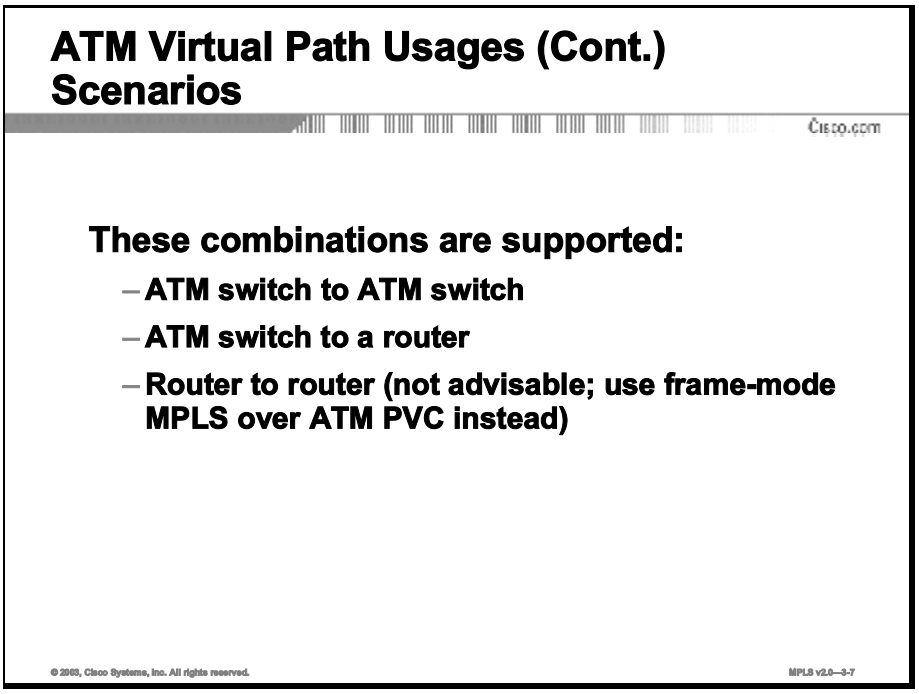

A virtual path can be used to connect any pair of devices across a public ATM network:

- 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 shows how to configure MPLS over ATM VP for switches.

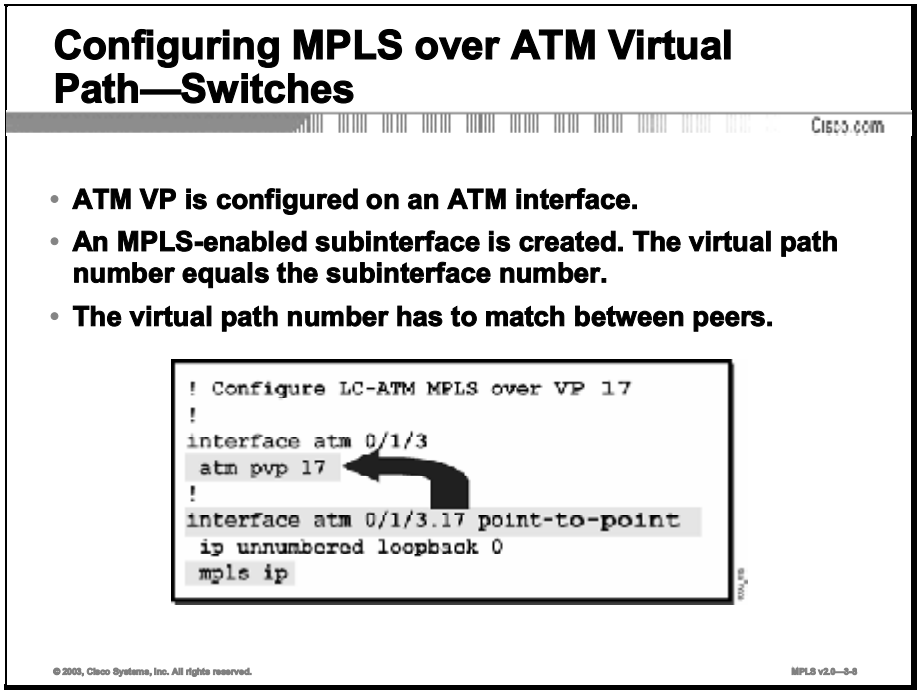

In this example, a virtual path with a VPI of 17 is created.

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

Note	The virtual path number has to match between peers.

Configuring MPLS over ATM Virtual Path—Switches (Cont.)

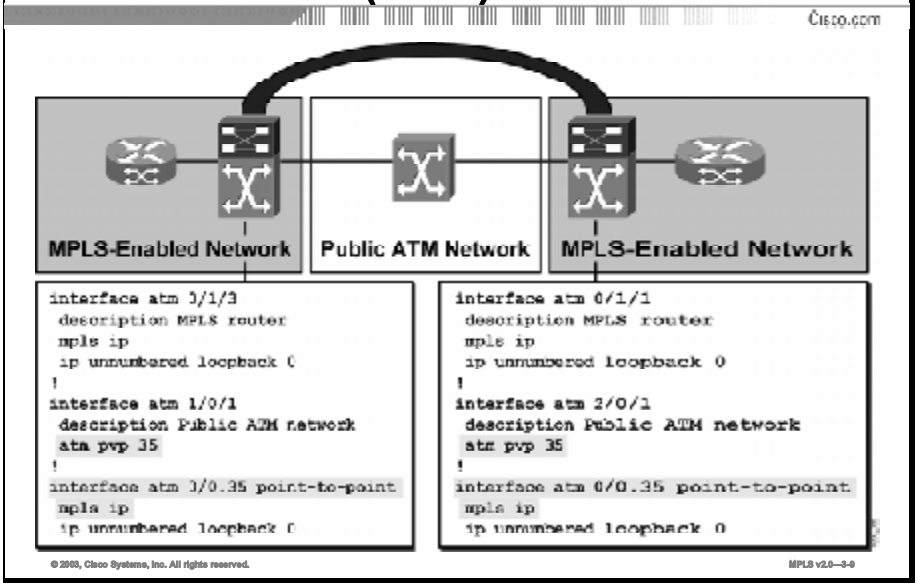

This example shows the configuration of both MPLS-enabled ATM switches connected by a virtual path across a public ATM network.

The VPI value 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 number.

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 shows how to configure MPLS over ATM VP for routers.

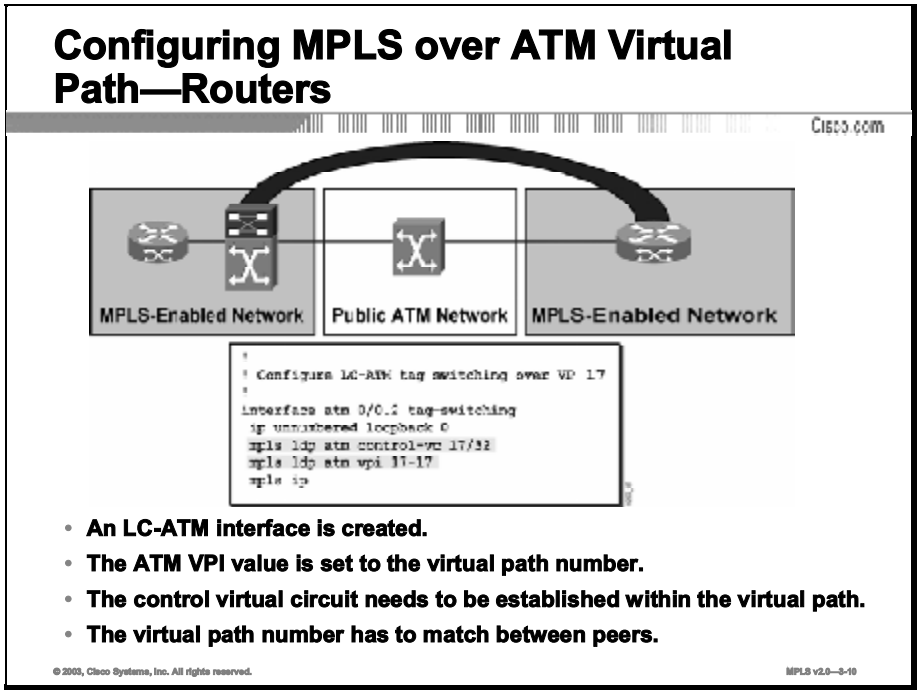

To simplify the provisioning of the connection across a public ATM network, you can also put the control virtual circuit into the virtual path.

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

If the public network is forwarding cells for VPI=17, then 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 value (17-17).

Note	The virtual path number has to match between peers.

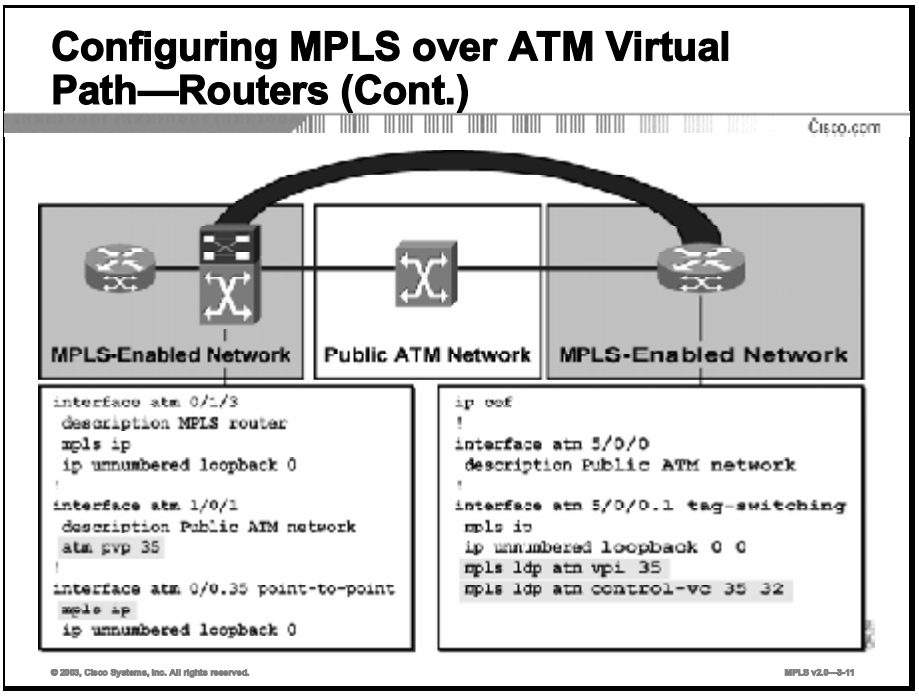

When you connect a router and a switch through a virtual path, you need to set only the parameters for the control virtual circuit and the label range on the router.

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 value 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.

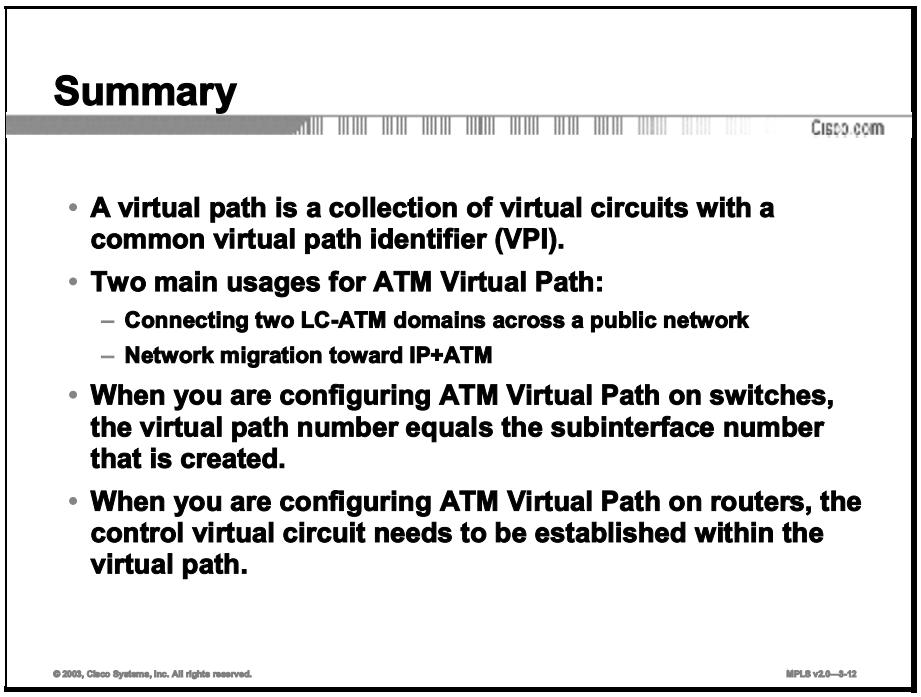

References

For additional information, refer to this resource:

- Search www.cisco.com for additional resources for the items covered in this lesson.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is a 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

Q2) Why is it mandatory that the VPI be used on both ends of the virtual path you are connecting two ATM switches?

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.

Q3) Which two of the following statements are correct when describing the configuration of ATM Virtual Path between two ATM switches? (Choose two.)

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.

Q4) Which two of the following statements are correct when describing the configuration of ATM Virtual Path between two ATM routers? (Choose two.)

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.

Quiz Answer Key

Q1) C

Relates to: Introduction to ATM Virtual Path

Q2) A

Relates to: ATM Virtual Path Usages

Q3) A, C

Relates to: Configuring MPLS over ATM Virtual Path—Switches

Q4) A, D

Relates to: Configuring MPLS over ATM Virtual Path—Routers

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.

Relevance

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

This lesson describes how to monitor LC-ATM MPLS on Cisco IOS platforms.

Upon completing this lesson, you will be able to:

- Use IOS commands to monitor LC-ATM-specific label-switching functions
- Describe and use the **show mpls atm-ldp summary** IOS command
- Describe and use the **show mpls atm-ldp bindings** IOS command
- Describe and use the **show mpls atm-ldp capability** IOS command
- Describe and use IOS debug commands for ATM-specific LDP functions

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Configuring LC-ATM MPLS over ATM Virtual Path” lesson of this module

Outline

This lesson includes these topics:

- Overview
- Monitoring Specific LC-ATM Label-Switching Functions
- **show mpls atm-ldp summary**
- **show mpls atm-ldp bindings**
- **show mpls atm-ldp capability**
- Debugging Specific ATM LDP Functions
- Summary
- Quiz

Monitoring Specific LC-ATM Label-Switching Functions

This topic introduces the Cisco IOS commands that you can use to monitor specific LC-ATM switching functions.

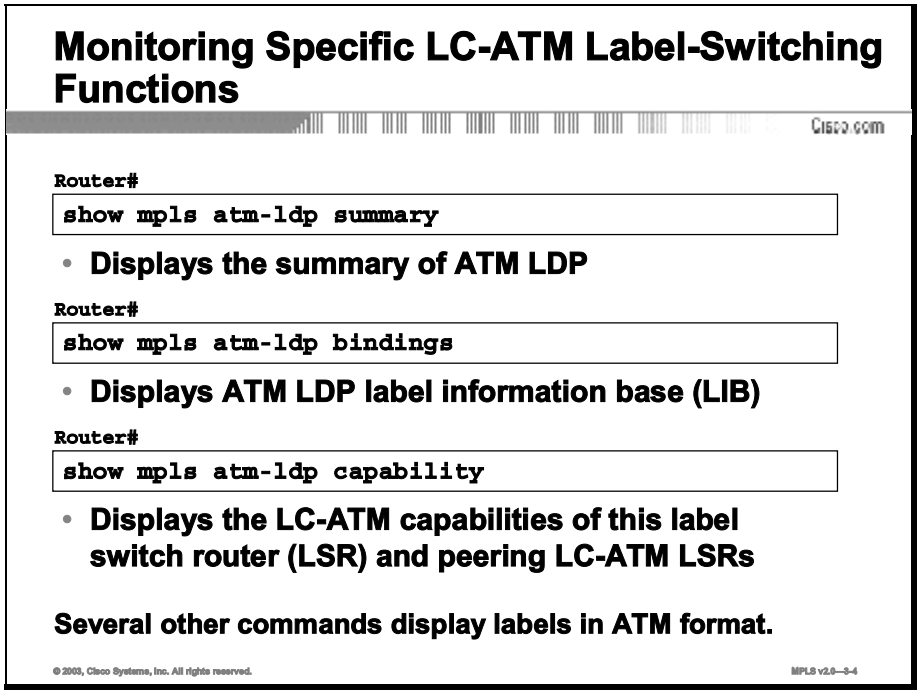

These commands, while similar to the **show mpls ldp** commands, display ATM-specific parameters. Use a question mark to see all of the subcommands or use the **show mpls atm ldp** command.

show mpls atm-ldp summary

To display summary information about all the entries in the ATM label-binding database, use the **show mpls atm-ldp summary** 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:

- **show mpls atm-ldp bindings** [*network* {*mask* | *length*}] [**local-label** *vpi vci*] [**remote-label** *vpi vci*] [**neighbor** *interface*]

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 **show mpls atm-ldp capability** command in privileged EXEC mode:

- **show mpls atm-ldp capability**

show mpls atm-ldp summary

This topic shows the results of issuing the **show mpls atm-ldp summary** IOS command.

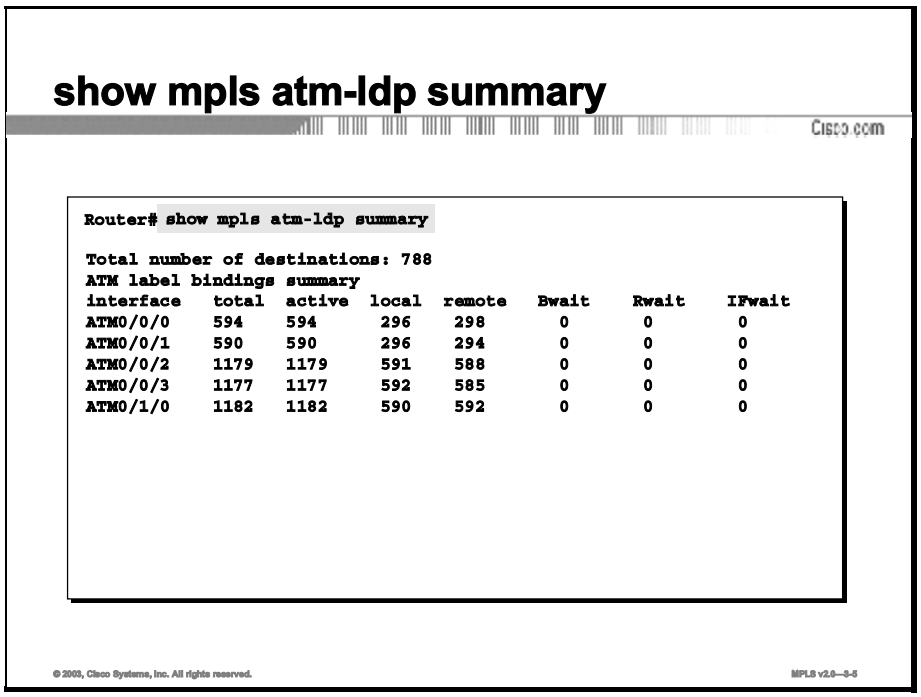

To display summary information about all the entries in the ATM label-binding database, use the **show mpls atm-ldp summary** command in privileged EXEC mode.

The following table describes the significant fields in the display.

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.

show mpls atm-ldp bindings

This topic shows the results of issuing the **show mpls atm-ldp bindings** IOS command.

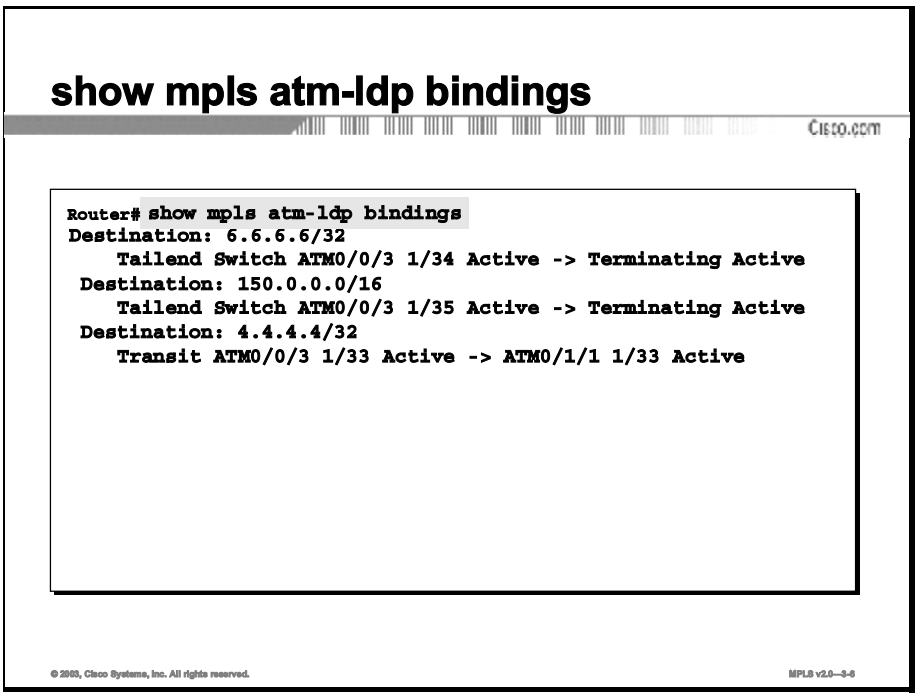

To display current label bindings, use the **show mpls atm-ldp bindings** command in privileged EXEC mode.

The following table describes the significant fields in the display.

Field	Description
Destination	Destination (network/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.
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	Virtual circuit descriptor number.

show mpls atm-ldp capability

This topic shows the results of issuing the **show mpls atm-ldp capability** IOS command.

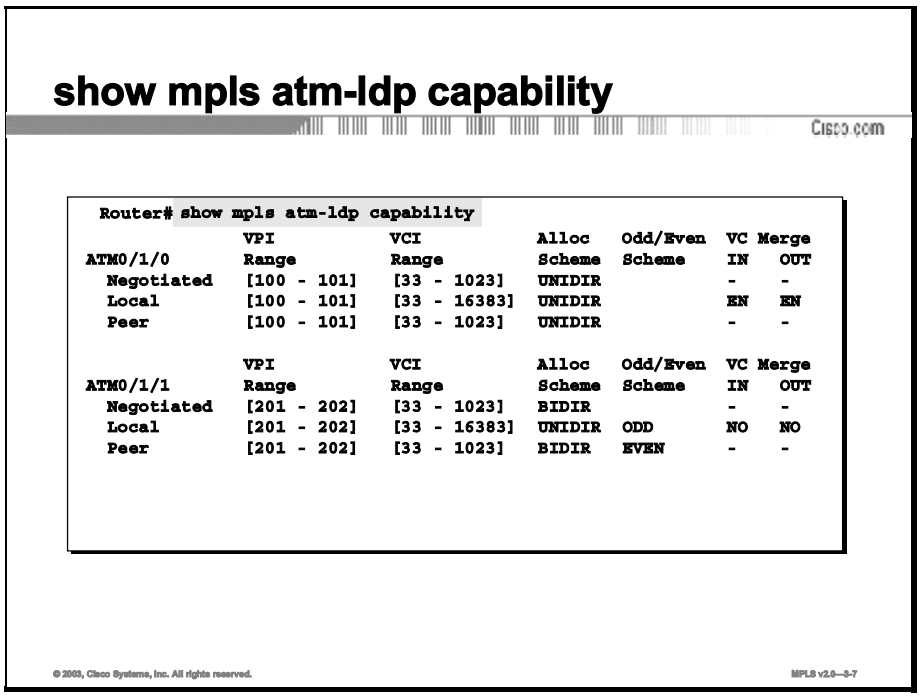

When two LSRs establish an LDP session, they negotiate parameters for the session, that is, a range of VPIs and VCIs that will be used as labels.

The **show mpls atm-ldp capability** command displays the MPLS ATM capabilities negotiated by LDP:

- The 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.

The following table describes the significant fields in the display.

Parameter	Description
VPI Range	Minimum and maximum number of VPIs supported on this interface.
VCI Range	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. It 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 IOS debug commands that you can use when troubleshooting ATM LDP issues.

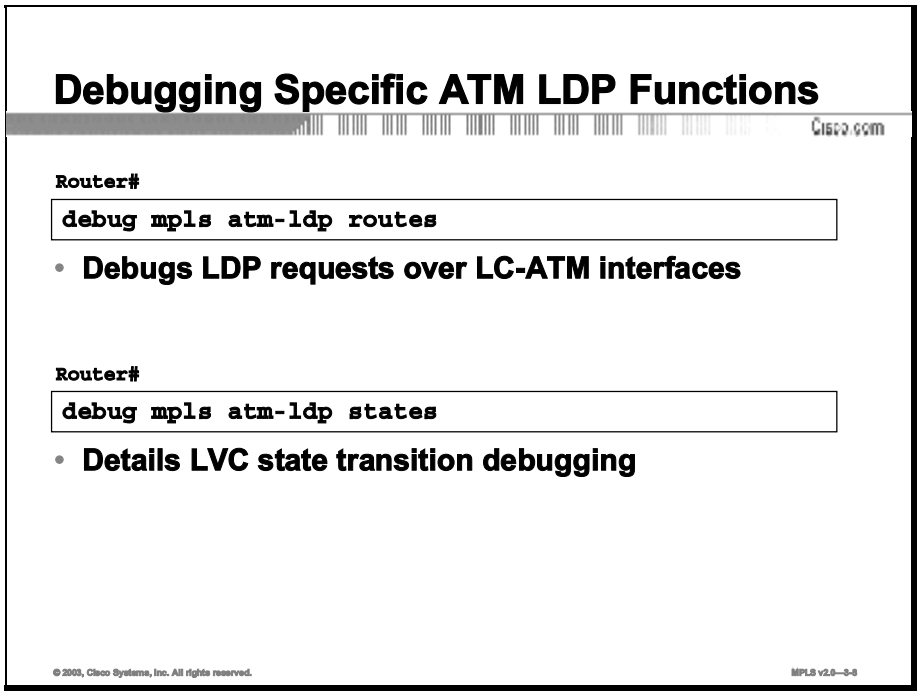

debug mpls atm-ldp routes

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.

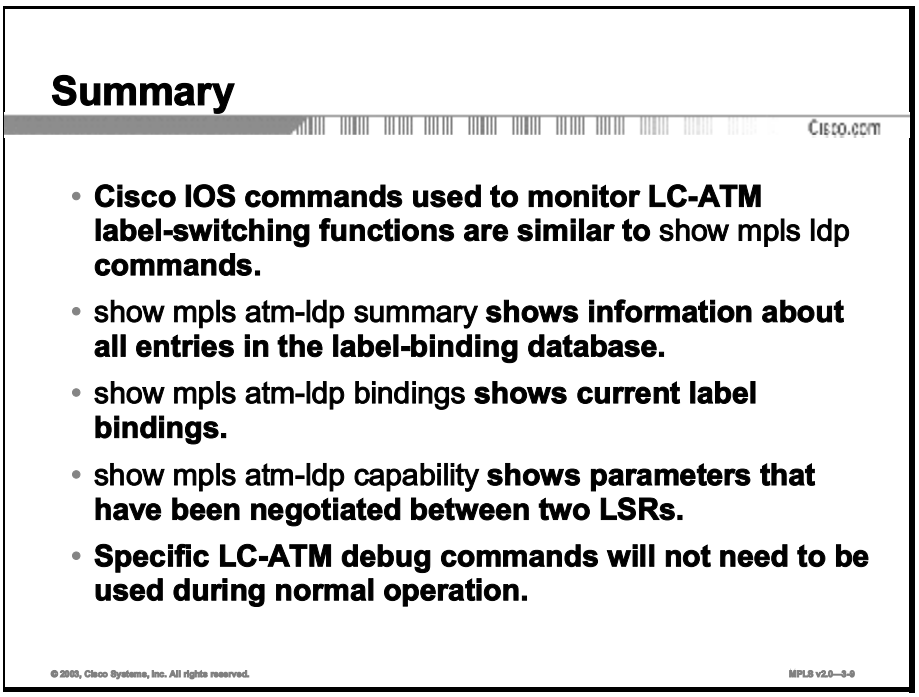

References

For additional information, refer to this resource:

- Search www.cisco.com for additional information about the topics covered in this lesson.

Next Steps

For the associated lab exercise, refer to these sections of the course Lab Guide:

- Lab Exercise 3-1: Establishing the Service Provider IGP Routing Environment
- Lab Exercise 3-2: Establishing the Core MPLS Environment

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which of the following is NOT used to monitor LC-ATM label-switching functions?

A) **show mpls atm-ldp labels**

B) **show mpls atm-ldp bindings**

C) **show mpls atm-ldp summary**

D) **show mpls atm-ldp capability**

Q2) Which of the following provides summary information about all entries in the label-binding database?

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**

Q3) Which of the following shows the current label bindings?

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**

Q4) Which of the following shows the negotiated parameters between LSRs?

A) **show mpls atm-ldp lsr**

B) **show mpls atm-ldp bindings**

C) **show mpls atm-ldp summary**

D) **show mpls atm-ldp capability**

Q5) Which of the following is used to debug an LC-ATM issue?

A) **debug mpls atm-ldp lsrs**

B) **debug mpls atm-ldp routes**

C) **debug mpls atm-ldp nodes**

D) **debug mpls atm-ldp switches**

Quiz Answer Key

Q1) A

Relates to: Monitoring Specific LC-ATM Label-Switching Functions

Q2) B

Relates to: **show mpls atm-ldp summary**

Q3) A

Relates to: **show mpls atm-ldp bindings**

Q4) D

Relates to: **show mpls atm-ldp capability**

Q5) B

Relates to: Debugging Specific ATM LDP Functions

Module 4

MPLS Virtual Private Networks 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 then describes Multiprotocol Label Switching (MPLS) VPN architecture and operations. It 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 do the following:

- Identify the major VPN topologies, their characteristics, and usage scenarios
- Describe the differences between an overlay VPN and peer-to-peer VPN, identifying the implementation, benefits, and drawbacks
- Describe the major VPN topology categories and their implementation
- Describe the major architectural blocks of MPLS VPNs, identifying the functions of route information propagation, route distinguisher, route target, and virtual routing tables
- Identify the routing requirements for MPLS VPNs by describing, from the customer and provider perspective, how routing tables appear on provider edge routers, how customer routes are propagated, and how end-to-end information flows
- Describe how packets are forwarded in an MPLS VPN environment, identifying how VPN labels get propagated, and explaining label imposition and the effect of summarization in the core

Module Outline

The module contains these lessons:

- Introduction to Virtual Private Networks
- Overlay and Peer-to-Peer VPNs
- VPN Categorization
- MPLS VPN Architecture
- MPLS VPN Routing Model
- MPLS VPN Packet Forwarding

Introduction to Virtual Private Networks

Overview

This lesson explains the concept of Virtual Private Networks (VPNs), including the terminology. The lesson also looks at why VPNs were first introduced.

Relevance

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

This lesson identifies the major VPN topologies, their characteristics, and usage scenarios.

Upon completing this lesson, you will be able to:

- 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

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of routing principles and concepts

Outline

This lesson includes these topics:

- Overview
- Traditional Router-Based Networks
- Virtual Private Networks
- VPN Terminology
- Switched WANs VPN Terminology
- Summary
- Quiz

Traditional Router-Based Networks

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

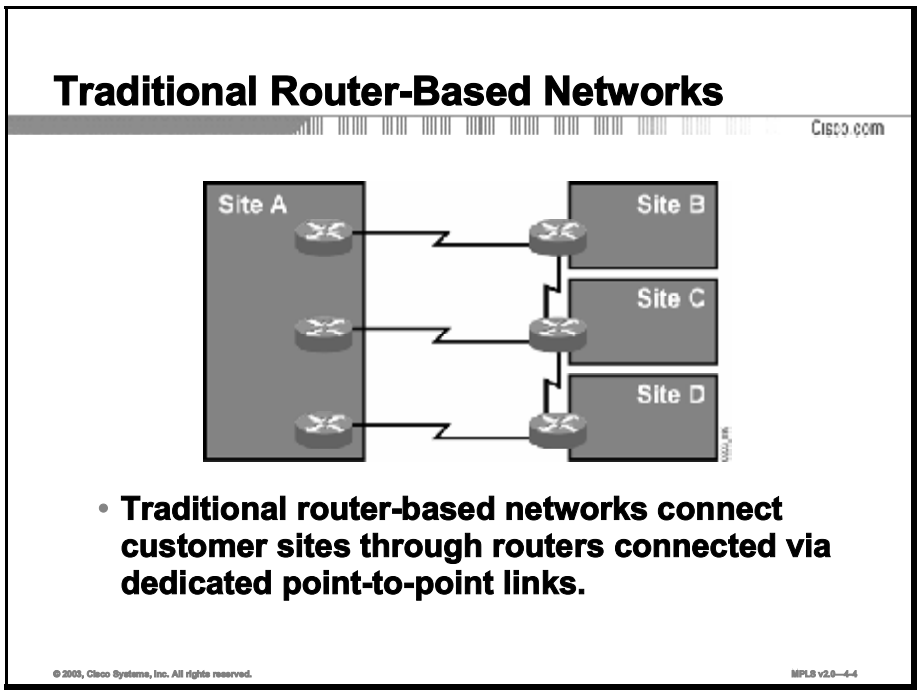

Traditional router-based networks were implemented with dedicated point-to-point links connecting customer sites. The cost of this approach was comparatively high for these reasons:

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

Virtual Private Networks

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

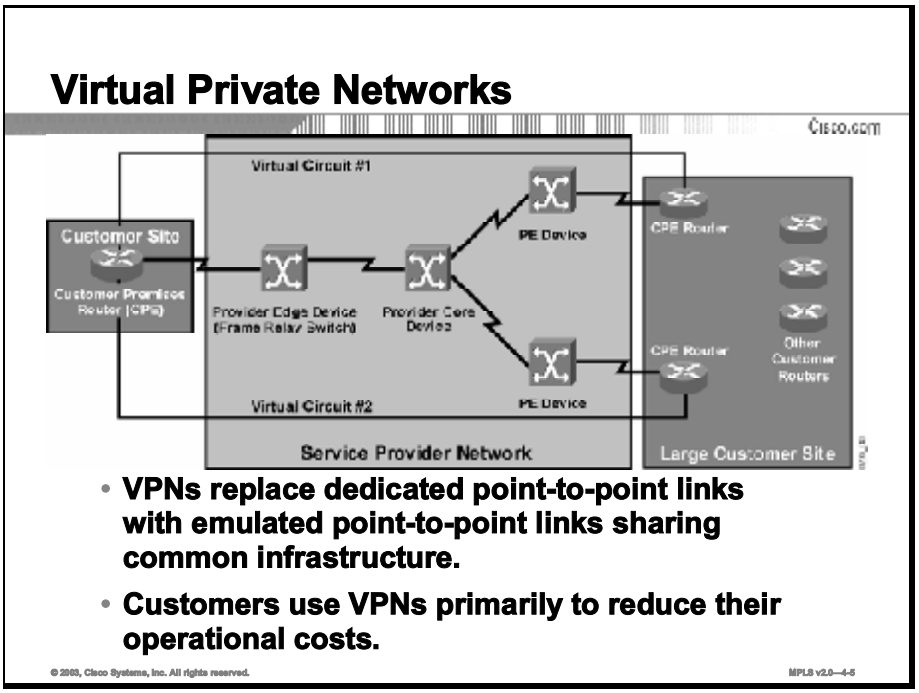

VPNs were introduced very early in the history of data communications with technologies such as X.25 and Frame Relay, which use virtual circuits to establish the end-to-end connection over a shared service provider infrastructure. These technologies, although sometimes considered legacy technologies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

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.

VPN Terminology

This topic describes VPN devices and terminology.

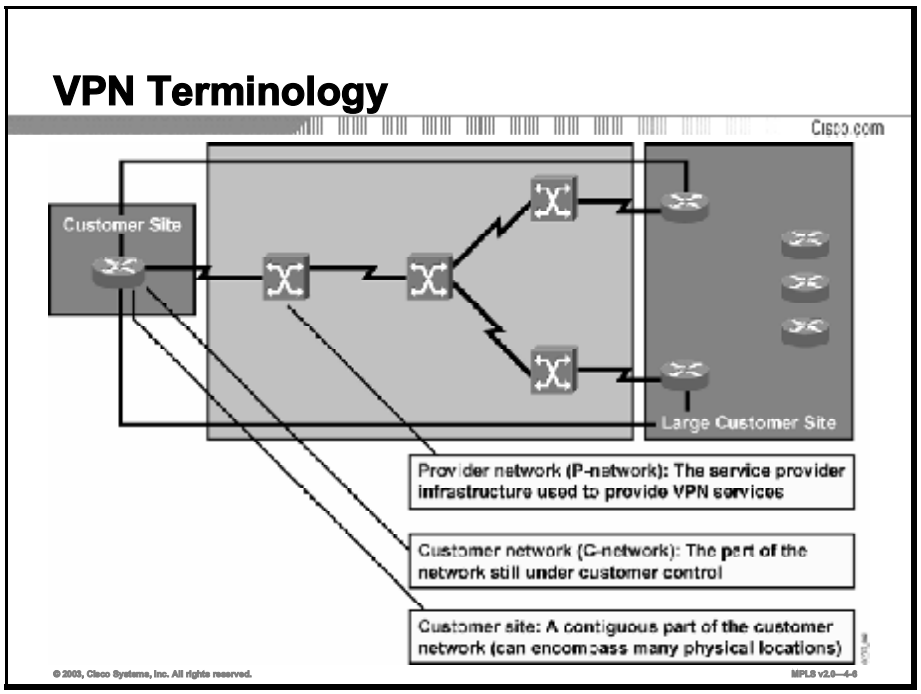

There are many conceptual models and terminologies describing various VPN technologies and implementations. This lesson focuses on the terminology introduced by MPLS VPN architecture. The terminology is generic enough to cover nearly any VPN technology or implementation and is thus extremely versatile.

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).

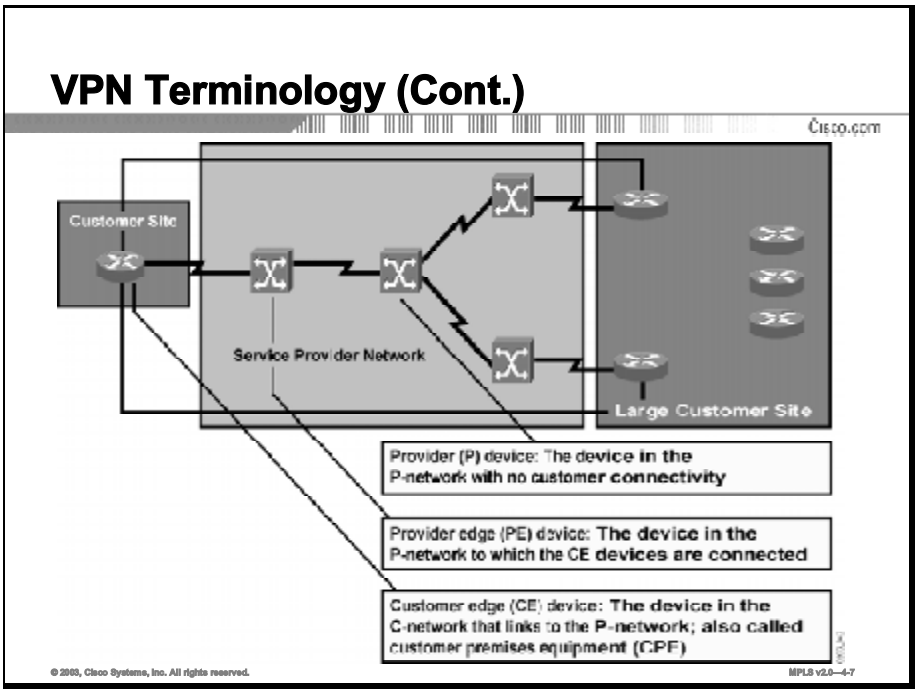

The devices that enable the overall VPN solution are named based on their position in the network:

- 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 customer premises equipment (CPE).

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

- 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 would be core (or transit) switches.

Switched WANs VPN Terminology

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

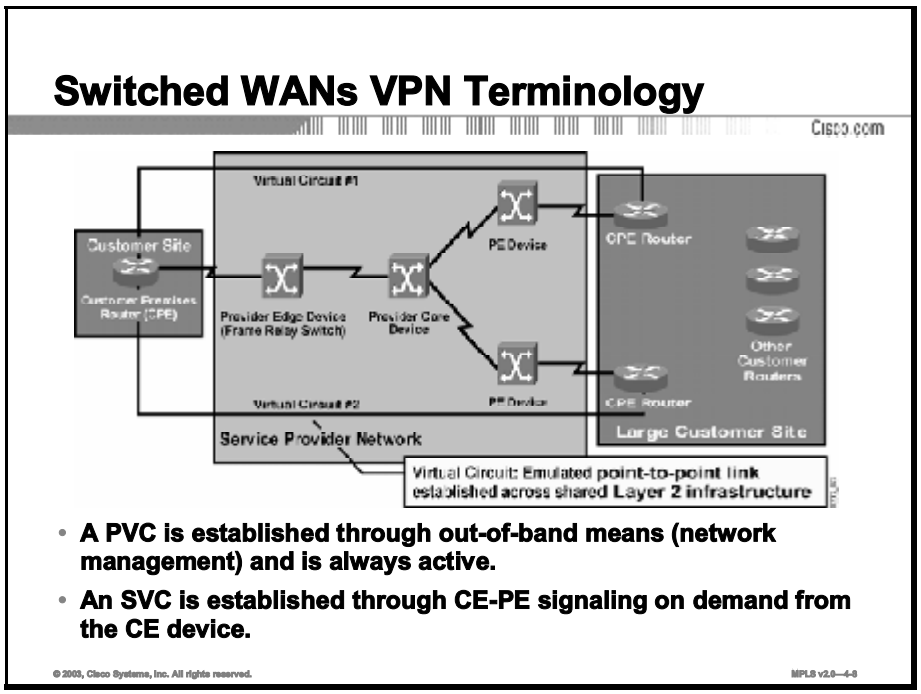

Switched WAN technologies introduced the virtual circuit, an emulated point-to-point link established across the Layer 2 infrastructure (for example, a Frame Relay network). Virtual circuits are further differentiated into permanent virtual circuits (PVCs), which are pre-established by means of network 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.

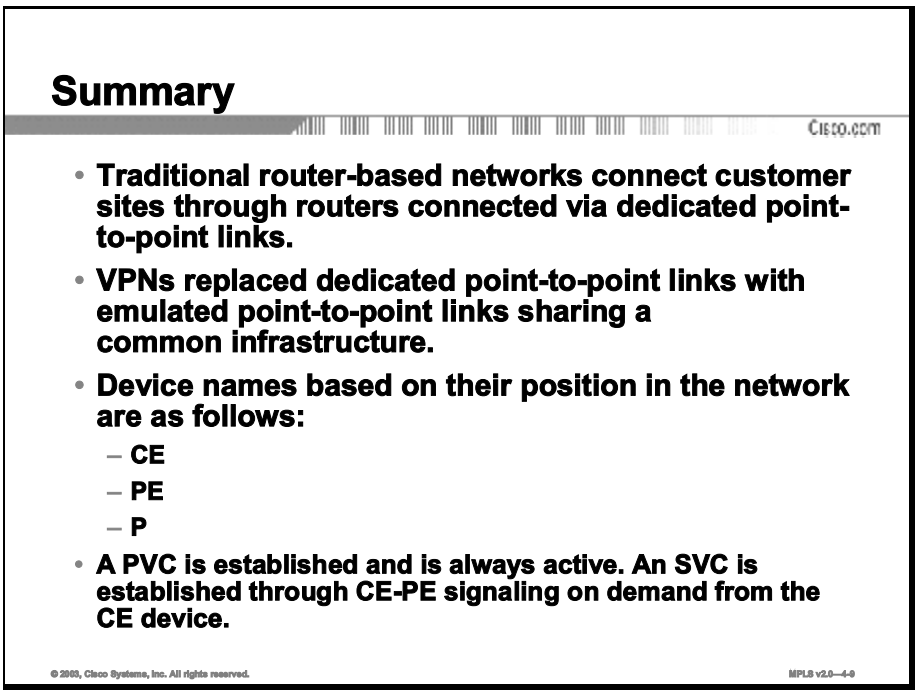

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Traditional router-based networks were implemented using what type of links?

A) PVC

B) dedicated point-to-point

C) SVC

D) emulated point-to-point

Q2) VPNs are implemented using what type of links?

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.)

A) P device

B) CE device

C) PE device

D) CPE device

Q4) What are the two types of virtual circuit supported by switched WAN technologies?

Quiz Answer Key

Q1) B

Relates to: Traditional Router-Based Networks

Q2) A

Relates to: Virtual Private Networks

Q3) A, C

Relates to: VPN Terminology

Q4) switched virtual circuits (SVCs), permanent virtual circuits (PVCs)

Relates to: Switched WANs VPN Terminology

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.

Relevance

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

This lesson describes the differences between overlay VPNs and peer-to-peer VPNs, explaining their implementation, benefits, and drawbacks.

Upon completing this lesson, you will be able to:

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

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components as well as basic routing and ATM principles and concepts

Outline

This lesson includes these topics:

- Overview
- VPN Implementation Technologies
- Overlay VPNs
- Peer-to-Peer VPNs
- Benefits of VPN Implementations
- Drawbacks of VPN Implementations
- Drawbacks of Traditional Peer-to-Peer VPNs
- Summary
- Quiz

VPN Implementation Technologies

This topic describes the two major VPN models.

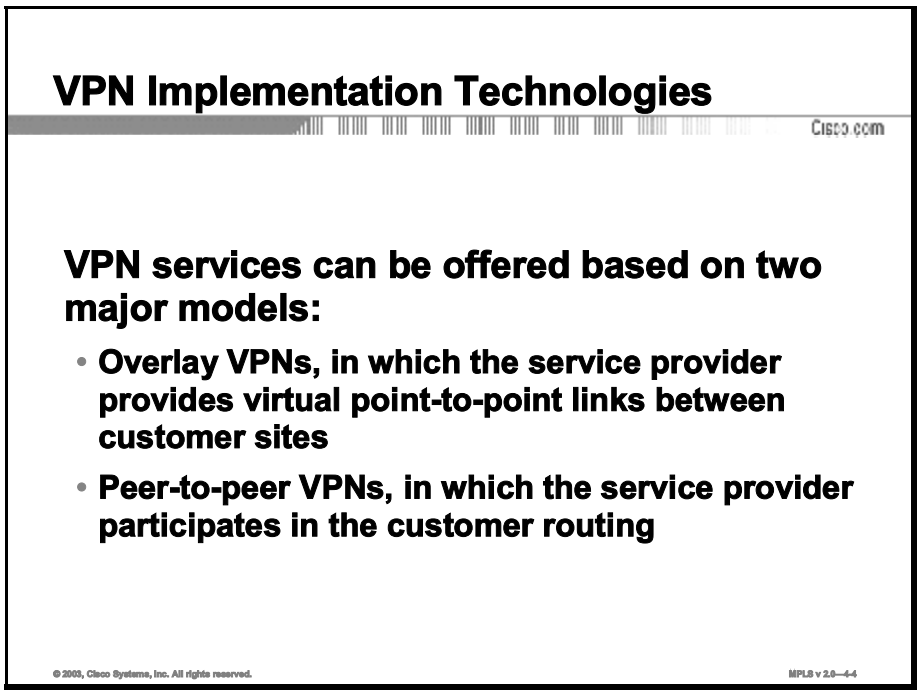

Traditional VPN implementations were all based on the overlay model, in which the service provider sold virtual circuits between customer sites as a replacement for dedicated point-to-point links. The overlay model had a number of drawbacks, which are identified in this lesson. To overcome these drawbacks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

Overlay VPNs

This topic describes the implementation of overlay VPNs.

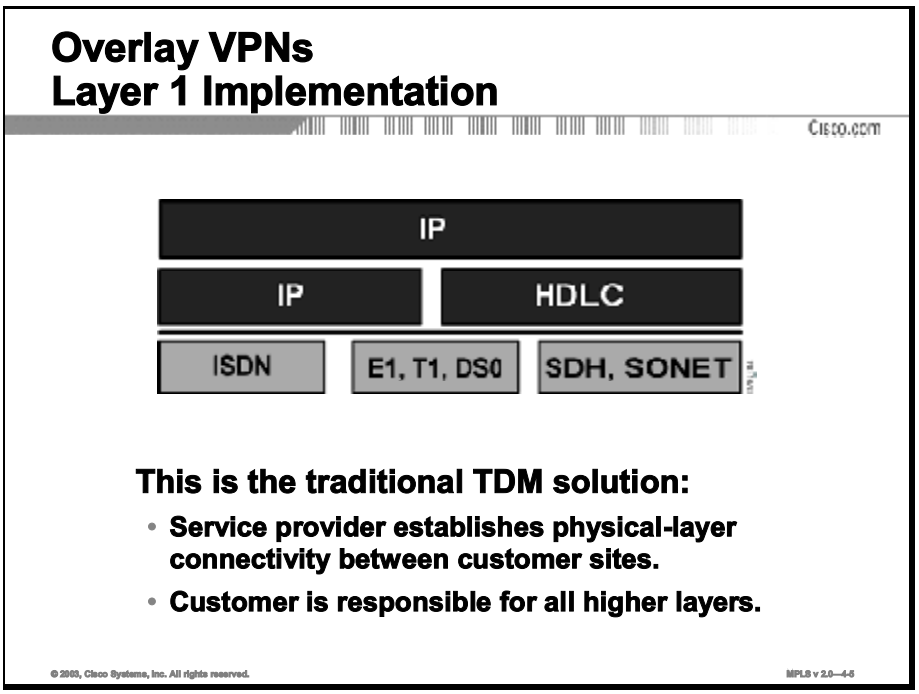

In the Layer 1 overlay VPN implementation, the service provider sells Layer 1 circuits (bit pipes) implemented with technologies such as ISDN, digital service zero (DS0), E1, T1, Synchronous Digital Hierarchy (SDH), or SONET. The customer is responsible for Layer 2 encapsulation between customer devices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

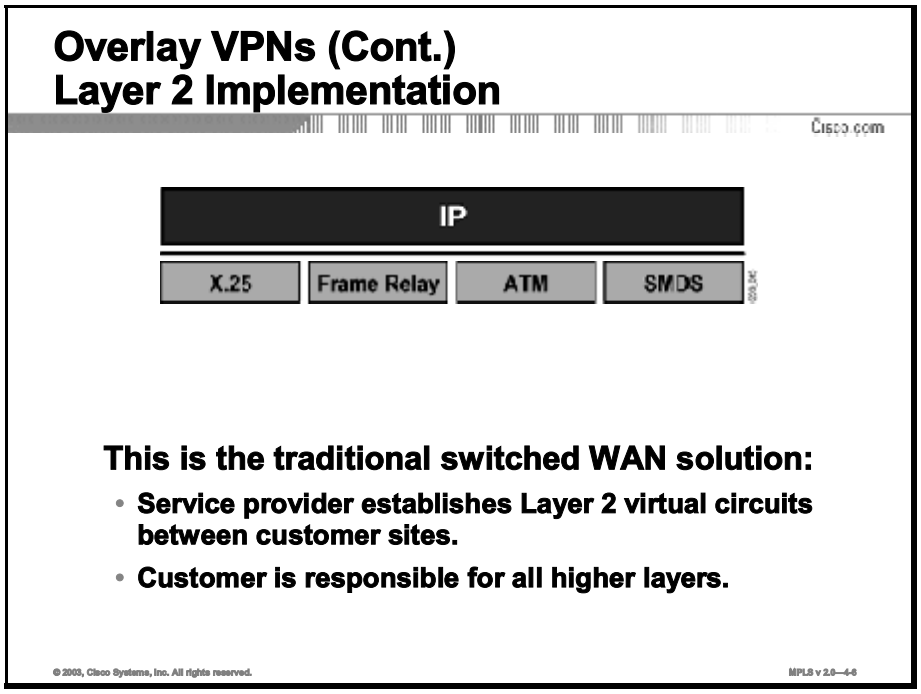

A Layer 2 VPN implementation is the traditional switched WAN model, implemented with technologies such as X.25, Frame Relay, ATM, and Switched Multimegabit Data Service (SMDS). The service provider is responsible for transport of Layer 2 frames between customer sites, and the customer is responsible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

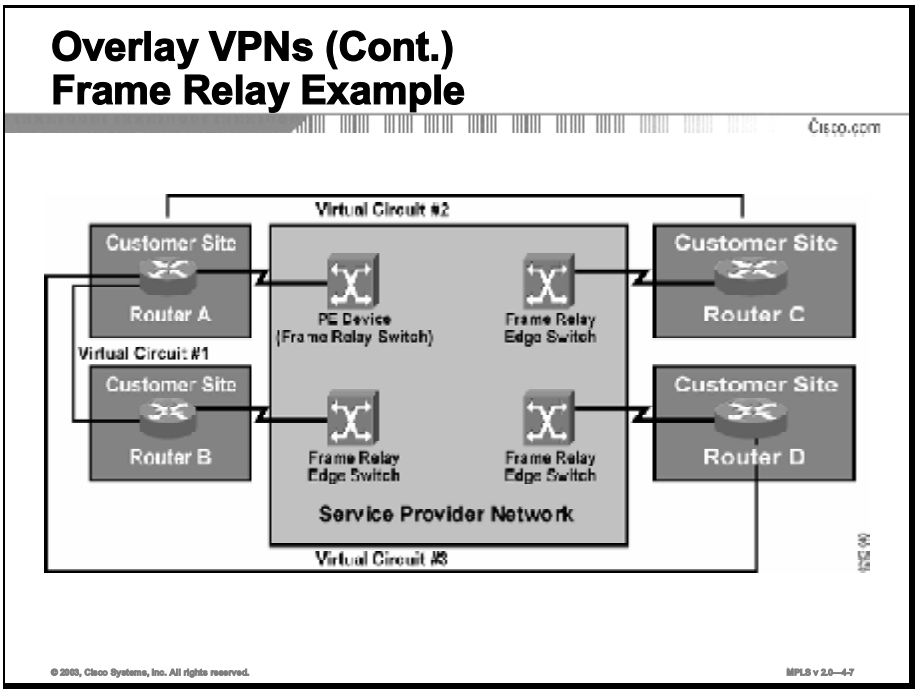

The figure shows a typical overlay VPN implemented by a Frame Relay network. The customer needs to connect three sites to site A (central site, or hub) and orders connectivity between site A (hub) and site B (spoke), between site A and site C (spoke), and between site A and site D (spoke). The service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

It should be noted that this implementation does not provide full connectivity. Data flow between spoke sites is through the hub.

Overlay VPNs (Cont.)
IP Tunneling

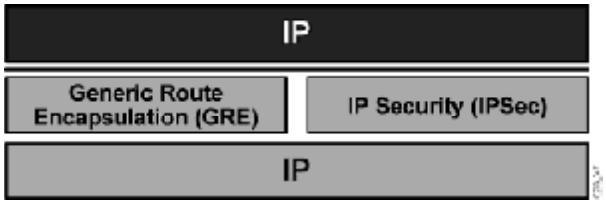

- **VPN is implemented with IP-over-IP tunnels:**
  - **Tunnels are established with GRE or IPSec.**
  - **GRE is simpler (and quicker); IPSec provides authentication and security.**

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

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

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

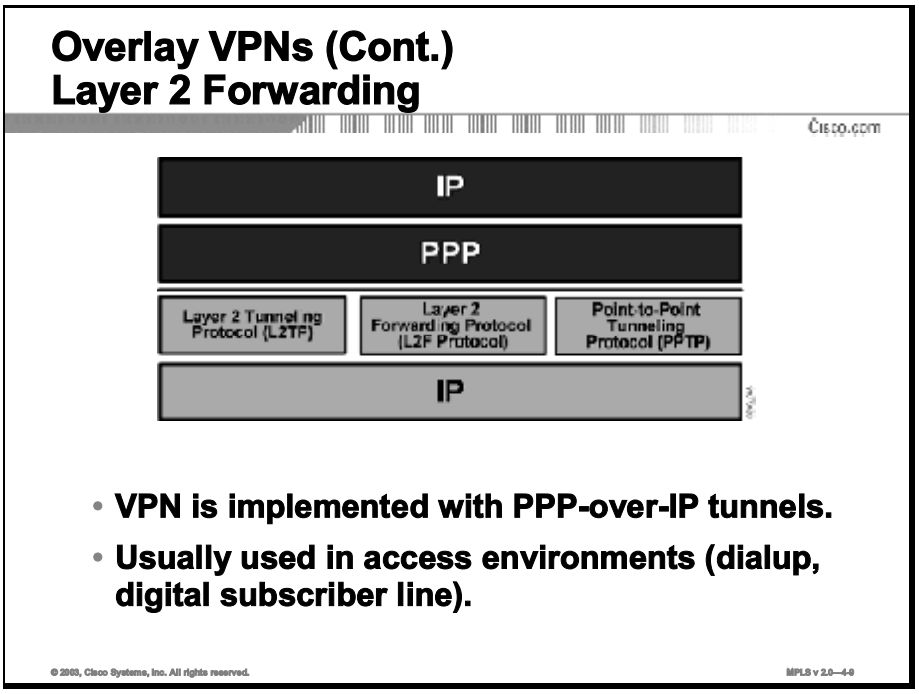

Yet another tunneling technique was first implemented in dialup networks, where service providers wanted to tunnel customer dialup data encapsulated in PPP frames over an IP backbone to the customer central site. To make the service provider transport transparent to the customer, PPP frames are exchanged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

There 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 (Cont.)
Layer 3 Routing

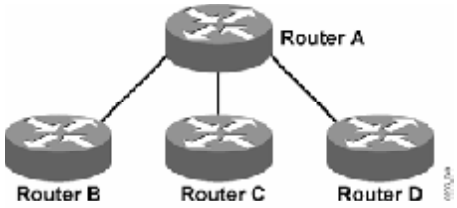

- **Service provider infrastructure appears as point-to-point links to customer routes.**
- **Routing protocols run directly between customer routers.**
- **Service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.**

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.

Peer-to-Peer VPNs

This topic describes the peer-to-peer VPN model.

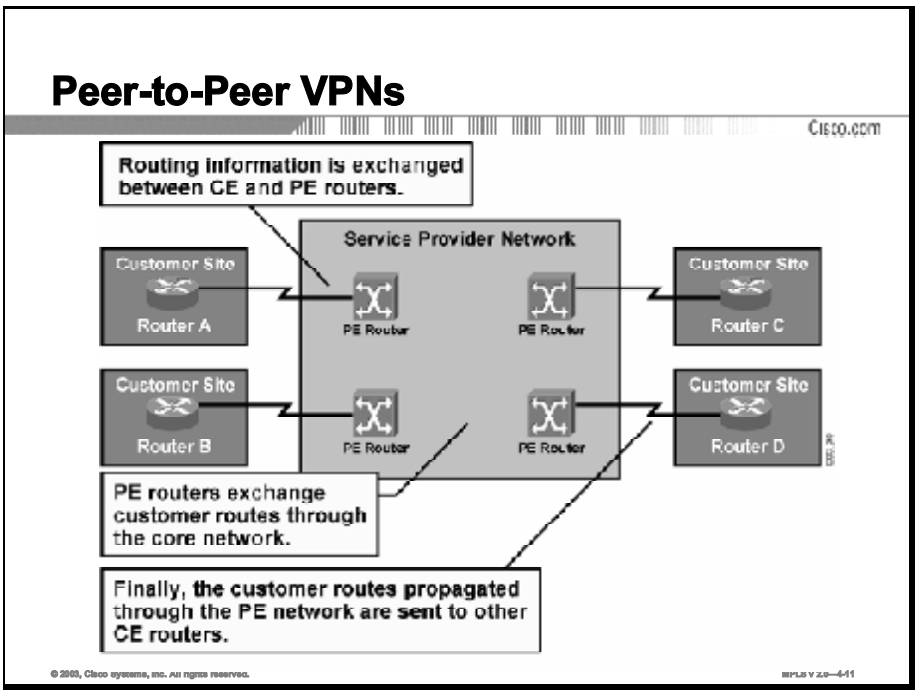

The overlay VPN model has a number of drawbacks, most significantly the need for customers to establish point-to-point links or virtual circuits between sites. The formula to calculate how many point-to-point links or virtual circuits are needed in the worst case is ([*n*][*n*-1])/2, where *n* is the 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 them across the service provider backbone, and finally propagating them to other customer sites.

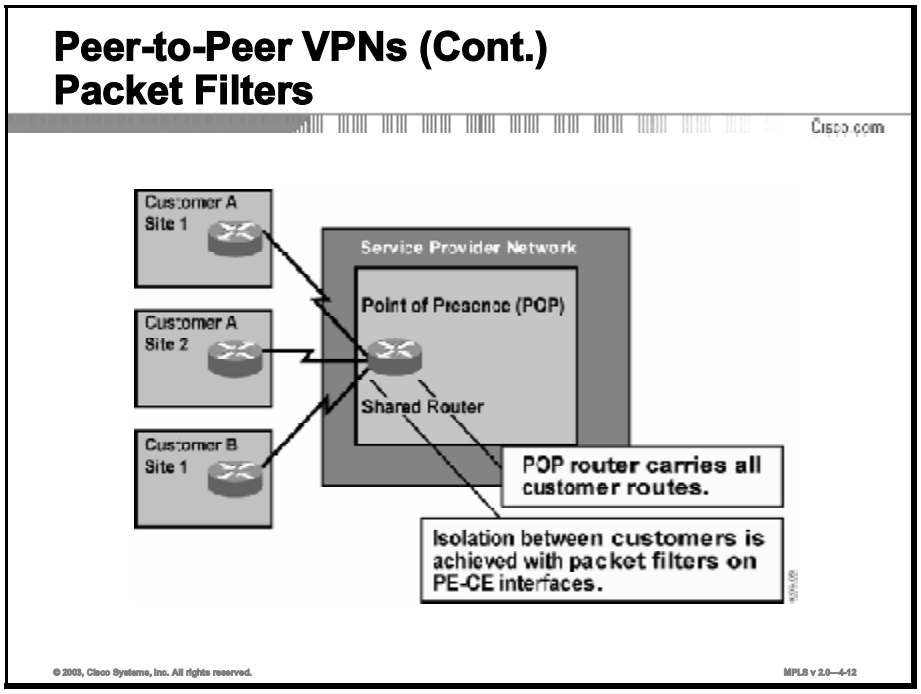

The first peer-to-peer VPN solutions appeared with the widespread deployment of IP in service provider networks. Architectures similar to that of the Internet were used to build them. Special provisions were taken into account to transform the architecture, which was targeted toward public backbones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

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

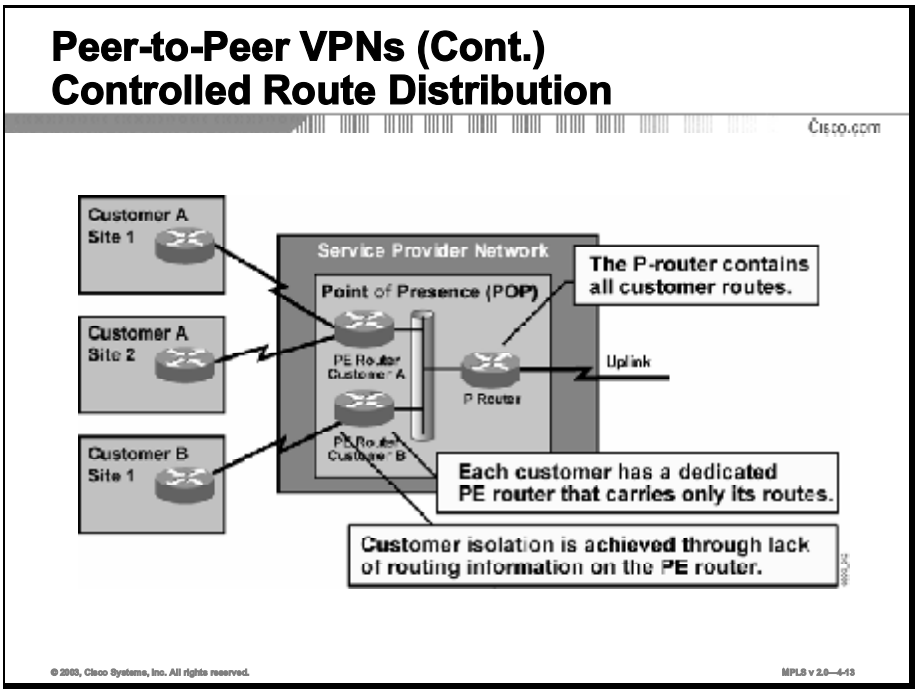

Maintaining packet filters is a mundane and error-prone task. Some service providers have thus implemented more innovative solutions based on controlled route distribution. In this approach, the customer has a dedicated PE router. The core service provider routers (P routers) contain all customer routes,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

In the example here, 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.

Benefits of VPN Implementations

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

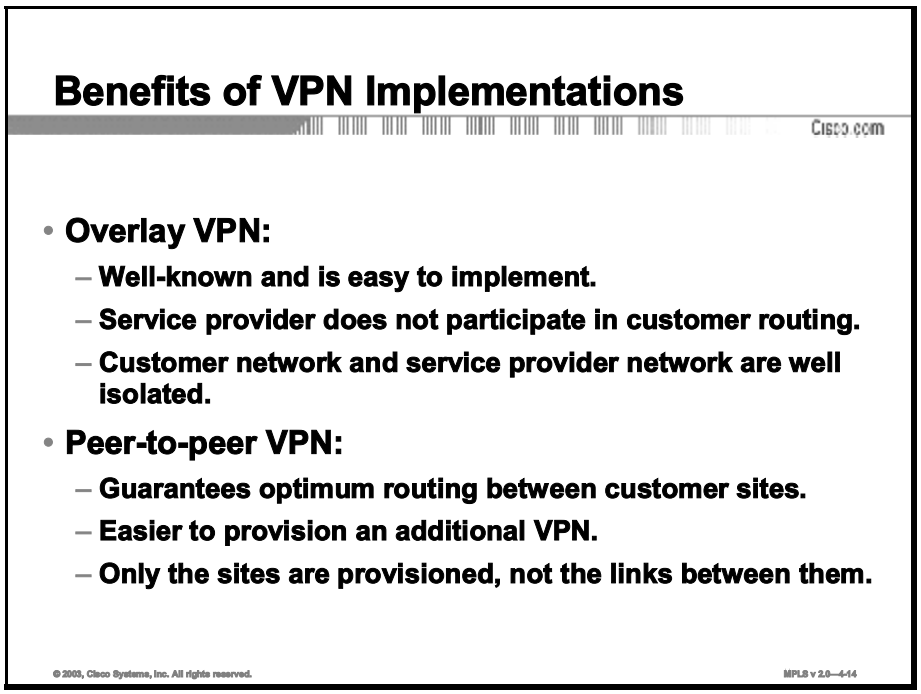

Each VPN model has a number of benefits. For example, overlay VPNs have these advantages:

- They 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

Drawbacks of VPN Implementations

This topic describes the drawbacks of each VPN implementation model.

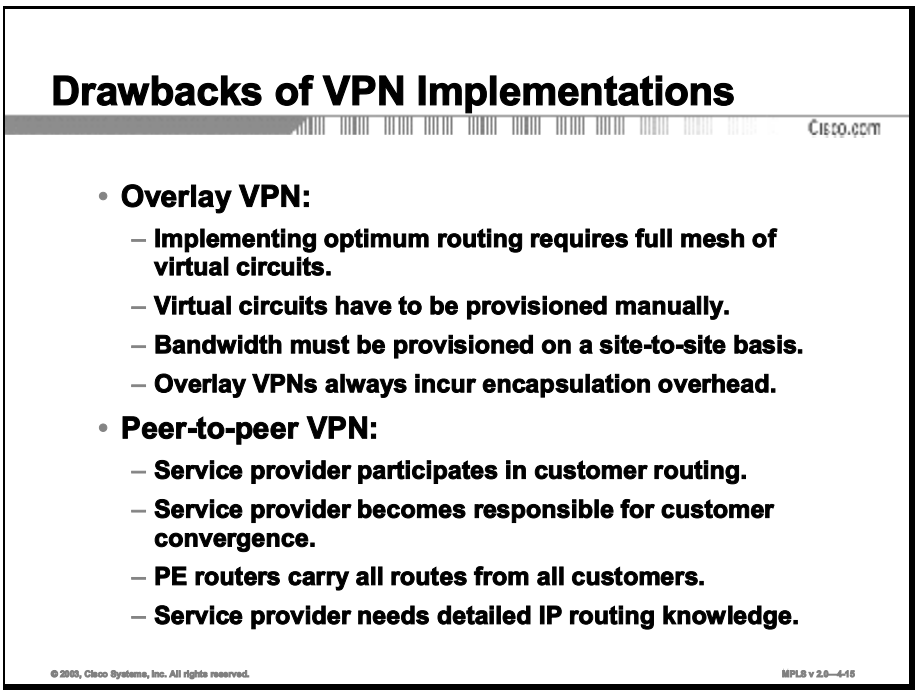

Each VPN model also has a number of drawbacks. Overlay VPNs have these disadvantages:

- They 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 to 80 bytes per transported datagram).

The major drawbacks of peer-to-peer VPNs arise from service provider involvement in customer routing:

- 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.

Drawbacks of Traditional Peer-to-Peer VPNs

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

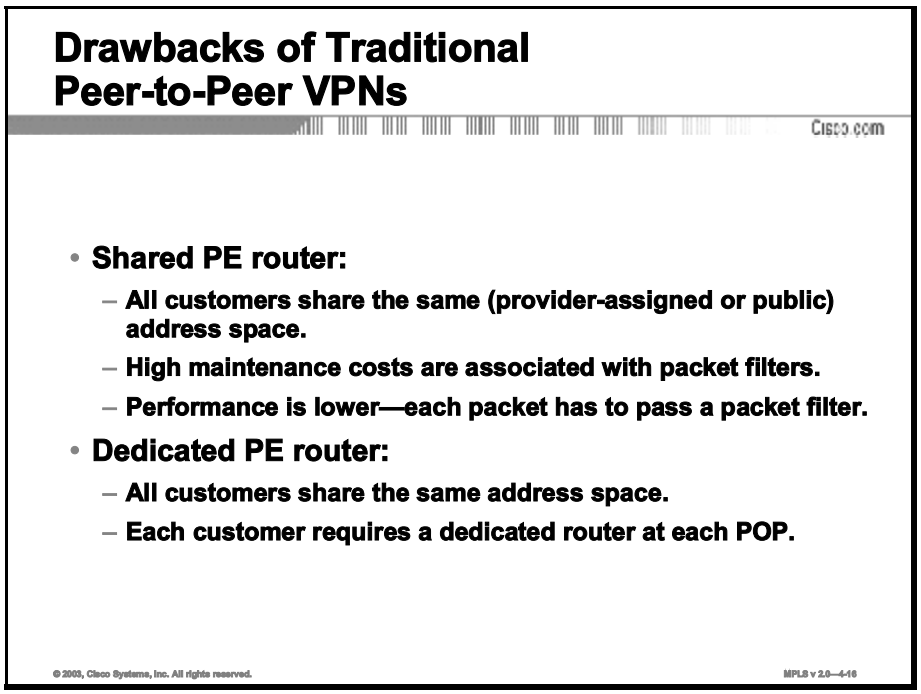

Pre-MPLS VPN implementations or peer-to-peer VPNs all share a common drawback. Customers have to share the same global address space, either using their own public IP addresses or relying on provider-assigned IP addresses. In both cases, connecting a new customer to a peer-to-peer VPN service usually 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 as well as 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 where the service provider has only a small number of large customers.

Summary

This topic summarizes the key points discussed in this lesson.

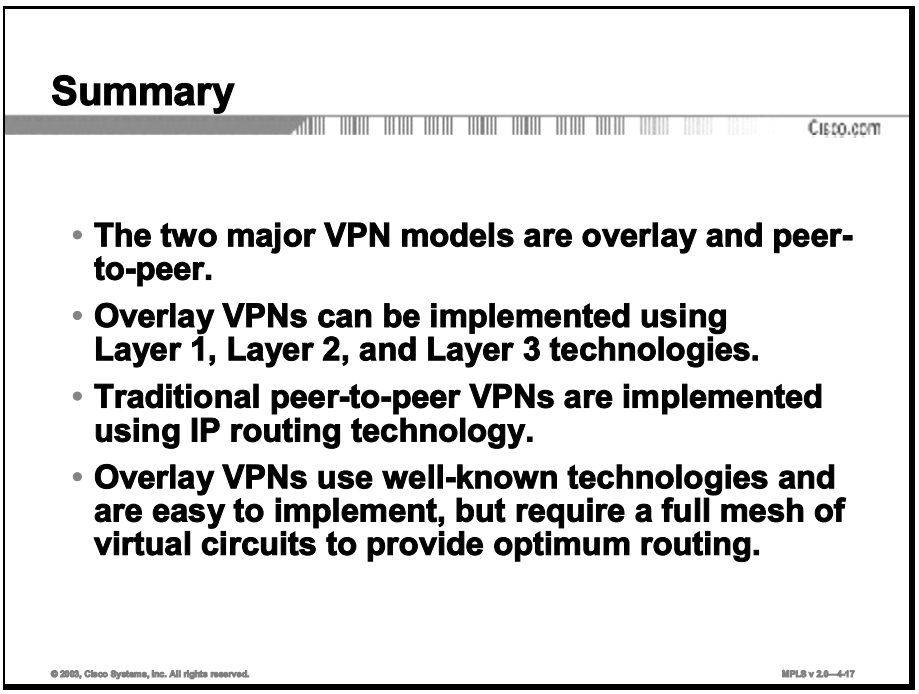

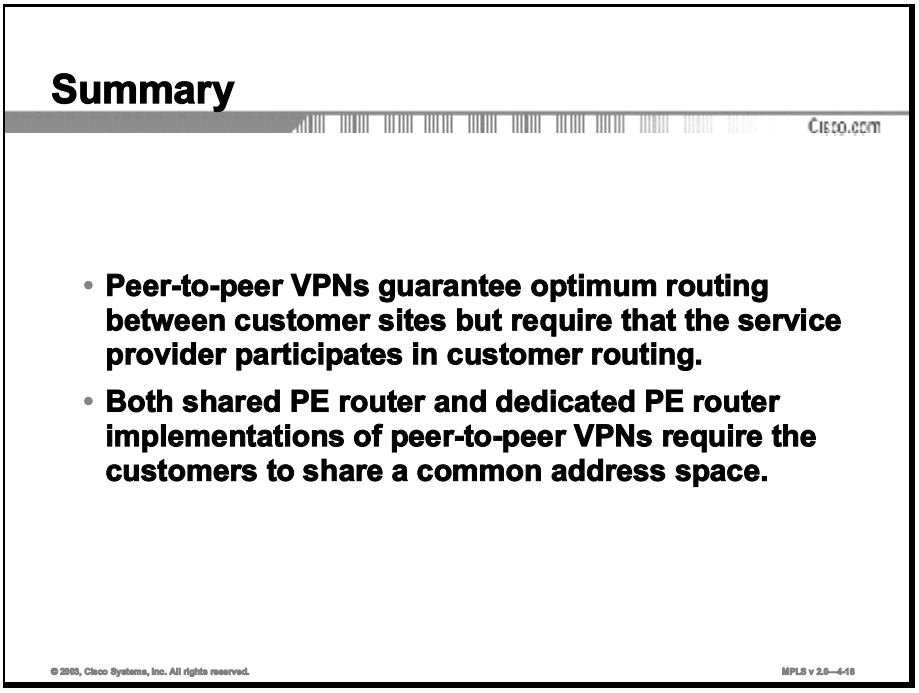

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) In the traditional switched WAN model for Layer 2 VPN implementation, what are the service provider and customer responsible for?

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?

_____

_____

_____

Q2) How is a peer-to-peer VPN implemented using packet filters?

_____

_____

_____

Q3) How do you implement a peer-to-peer VPN based on controlled route distribution?

_____

_____

_____

Q4) Which VPN type does not require the service provider to participate in customer routing?

A) overlay

B) peer-to-peer

Q5) For which VPN type is it easier to provision an additional VPN?

A) overlay

B) peer-to-peer

Q6) Which VPN type requires the PE router to carry all routes from all customers?

A) overlay

B) peer-to-peer

Q7) Which VPN type requires the service provider to participate in customer routing?

A) overlay

B) peer-to-peer

Q8) Describe the use of address space and packet routing in each of the following peer-to-peer implementations:

- Shared PE router

_____

_____

- Dedicated PE router

_____

_____

Quiz Answer Key

Q1) providing end-to-end connectivity, routing updates
The need for customers to establish point-to-point links or virtual circuits between sites.

Relates to: Overlay VPNs; Peer-to-Peer VPNs

Q2) 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.

Relates to: Peer-to-Peer VPNs

Q3) The core service provider routers (P routers) contain all customer routes, and the PE routers contain only routes of a single customer.

Relates to: Peer-to-Peer VPNs

Q4) A

Relates to: Benefits of VPN Implementations

Q5) D

Relates to: Benefits of VPN Implementations

Q6) B

Relates to: Drawbacks of VPN Implementations

Q7) B

Relates to: Drawbacks of VPN Implementations

Q8) 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.

Relates to: Drawbacks of Traditional Peer-to-Peer VPNs

VPN Categorization

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.

Relevance

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

Objectives

This lesson describes the characteristics of the different VPN topology categories.

Upon completing this lesson, you will be able to:

- Identify the major components of the overlay VPN topology category
- 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

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components as well as basic routing and ATM principles and concepts

Outline

This lesson includes these topics:

- Overview
- Overlay VPN Category
- Hub-and-Spoke Overlay VPN Topology
- Partial Mesh Overlay VPN Topology
- VPN Business Category
- Extranet VPNs
- VPN Connectivity Category
- Central Services Extranet
- Managed Network Implementation
- Summary
- Quiz

Overlay VPN Category

This topic identifies the major components of the overlay VPN topology category.

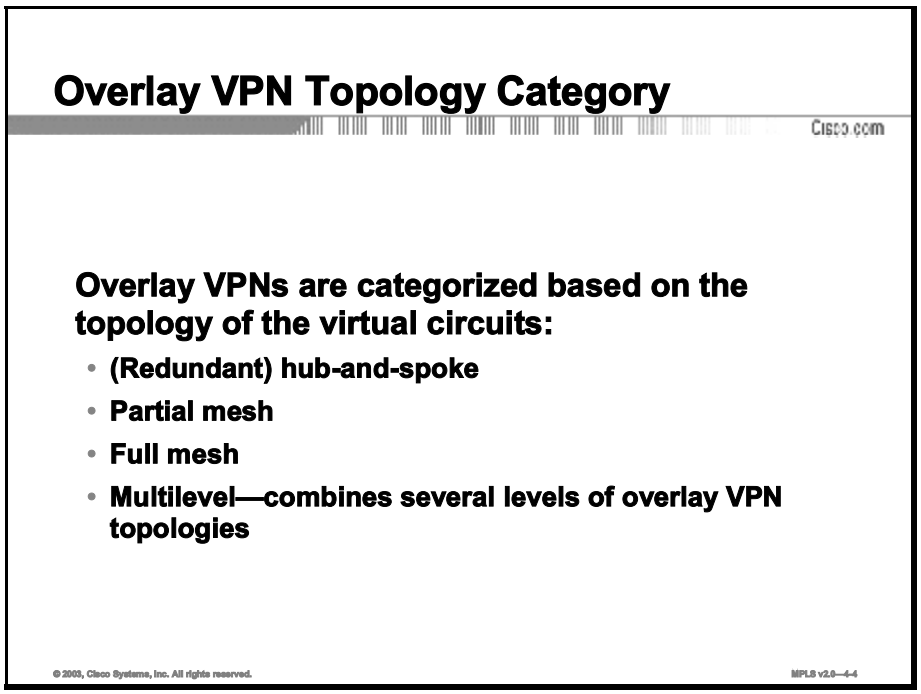

The oldest VPN category is based on the topology of point-to-point links in an overlay VPN implementation:

- **Hub-and-spoke:** Hub-and-spoke topology is the ultimate reduction of partial mesh topology: many sites (spokes) are connected only with the central site(s), or hub(s), with no direct connectivity between the spokes. To prevent single points of failure, hub-and-spoke topology is sometimes extended to *redundant* hub-and-spoke topology.
- **Full mesh:** Full mesh topology provides a dedicated virtual circuit between any two CE routers in the network.
- **Partial mesh:** 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)

Hub-and-Spoke Overlay VPN Topology

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

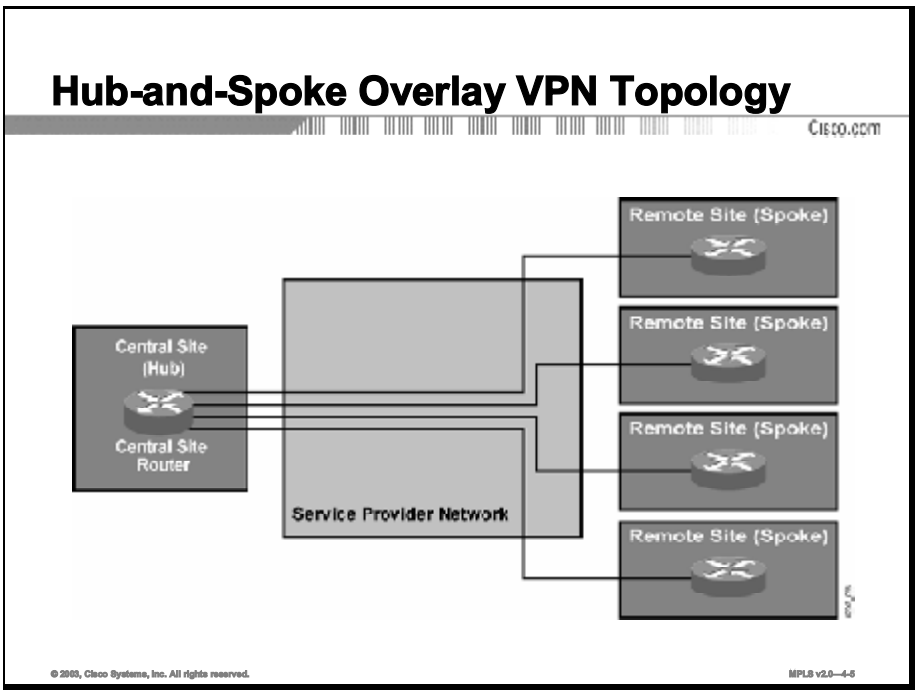

The hub-and-spoke topology is the simplest overlay VPN topology—all remote sites are linked with a single virtual circuit to a central CE router. The routing is also extremely simple—static routing or a distance vector protocol such as Routing Information Protocol (RIP) is more than adequate. If a dynamic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

Hub-and-Spoke Overlay VPN Topology (Cont.)
Redundant Hub-and-Spoke Topology

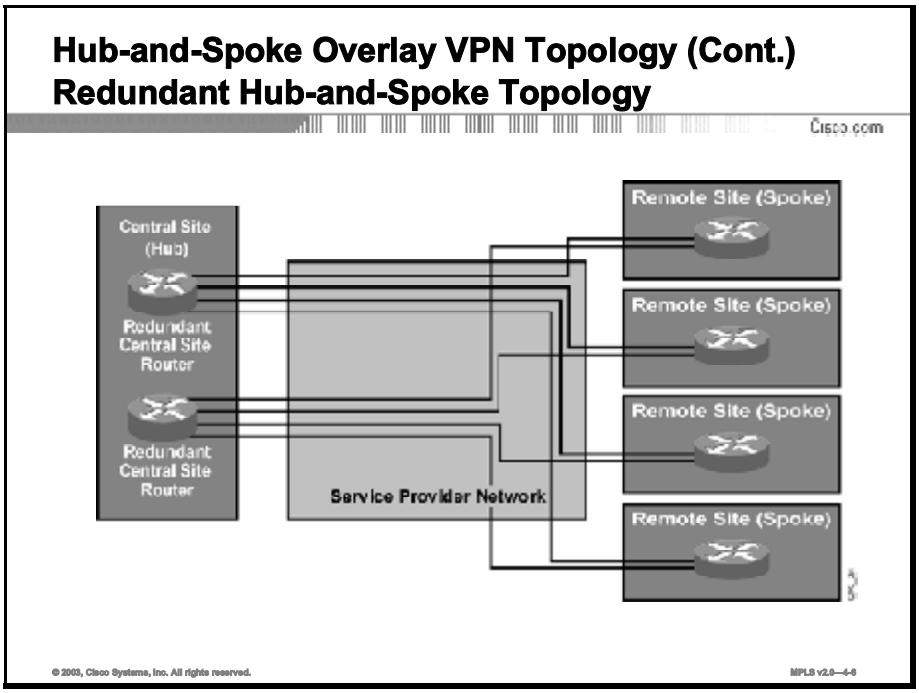

A typical redundant hub-and-spoke topology introduces central site redundancy (more complex topologies might also introduce router redundancy at spokes).

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/backup configuration.

Partial Mesh Overlay VPN Topology

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

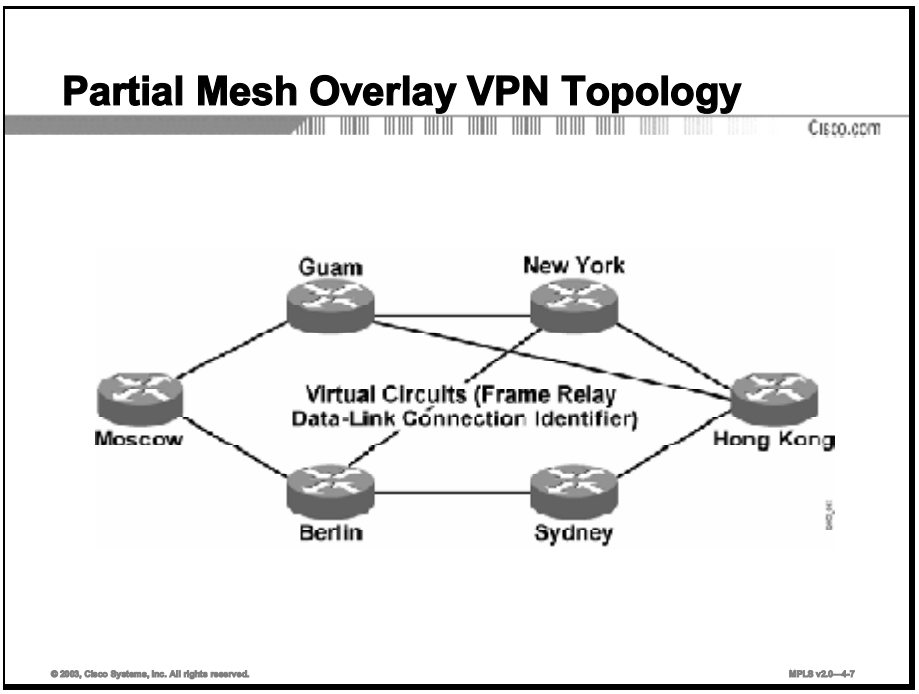

Partial mesh topology is used in environments where cost or complexity factors prevent a full mesh between customer sites. The virtual circuits in a partial mesh can be established based on a wide range of criteria:

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

VPN Business Category

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

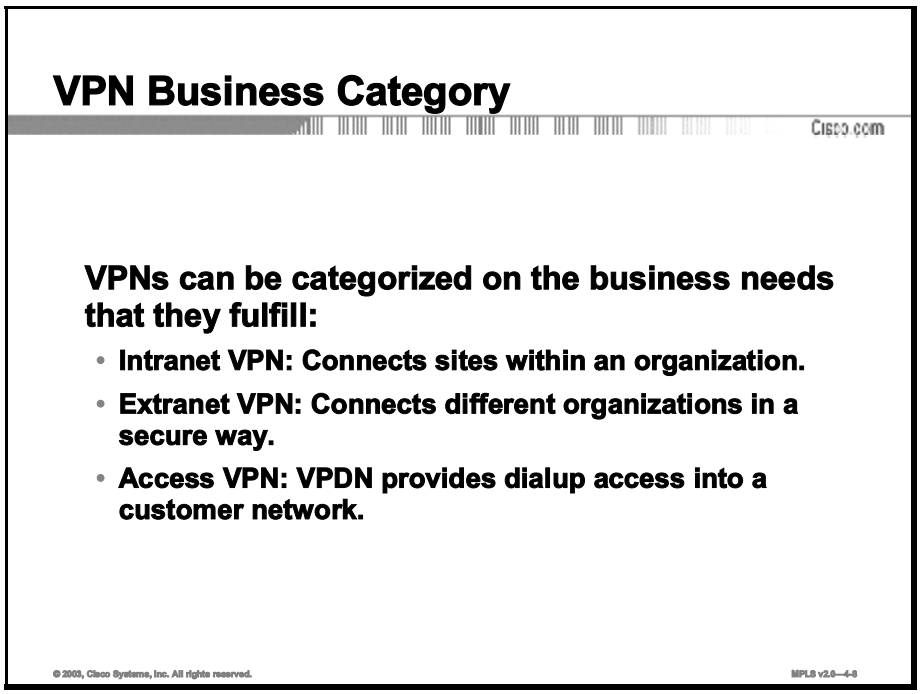

Another very popular VPN category classifies VPNs based on the business needs that they fulfill:

- **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.

The following two figures compare an overlay VPN implementation of an extranet with a peer-to-peer implementation. Similar comparisons can be made for intranets as well.

Extranet VPNs

This topic describes the characteristics of the extranet component of the VPN business category.

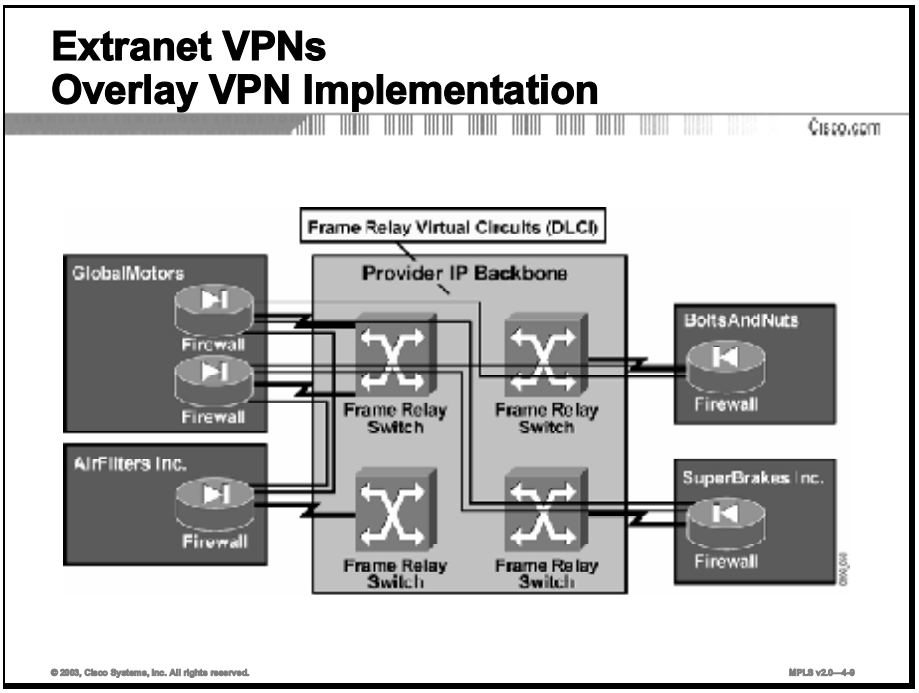

In an overlay implementation of an extranet, organizations are linked with dedicated virtual circuits. Traffic between two organizations can flow only if one of the following conditions is met:

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

Extranet VPNs (Cont.)
Peer-to-Peer VPN Implementation

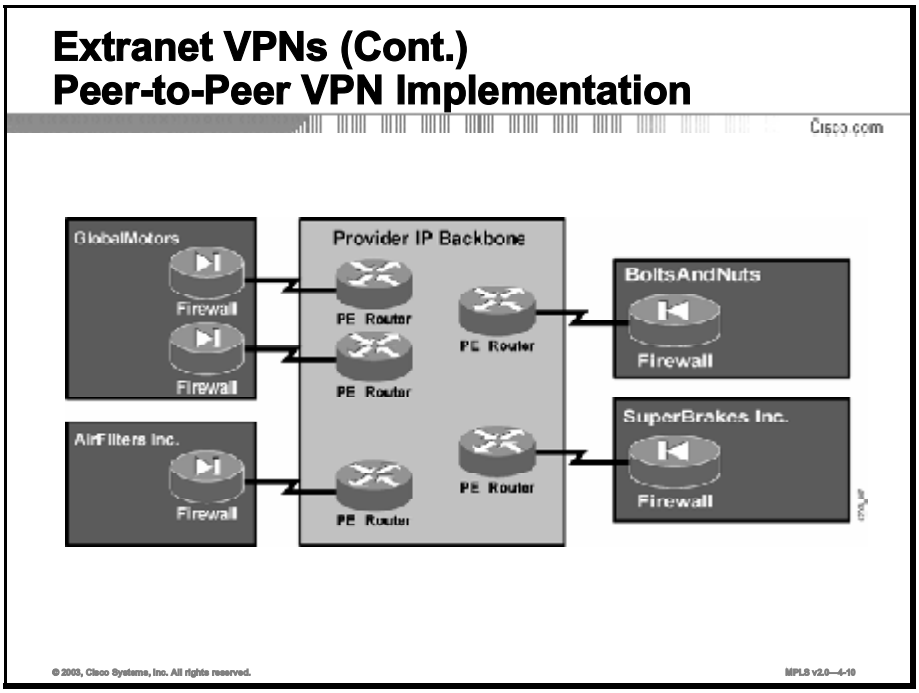

Peer-to-peer VPN implementation of an extranet VPN is very simple compared with overlay VPN implementation—all sites are connected to the P-network, and optimum routing between sites is enabled by default.

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.

VPN Connectivity Category

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

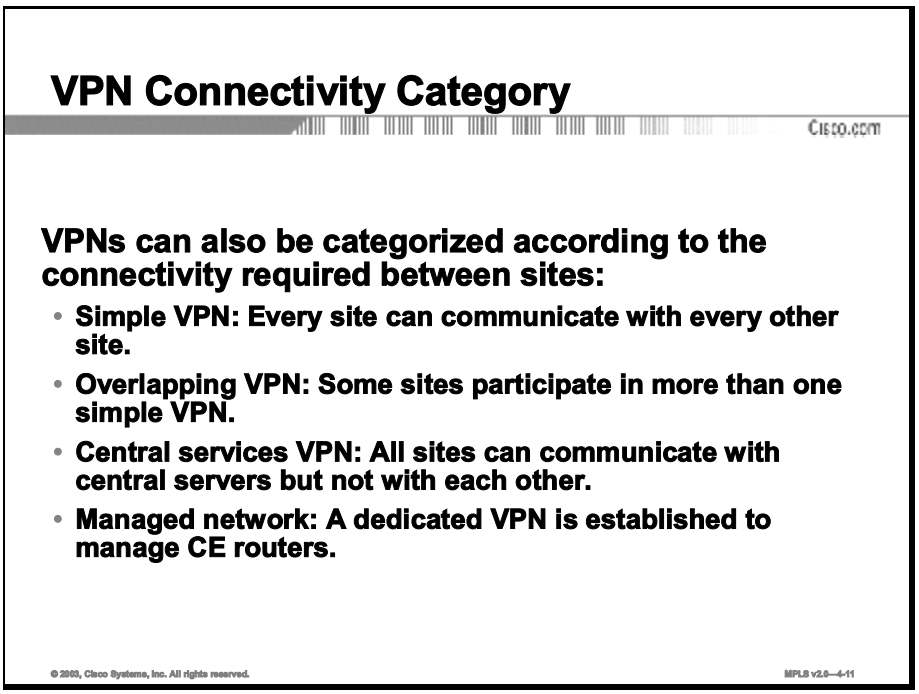

The VPNs discussed so far have usually been very simple in terms of connectivity:

- 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.

There are, however, 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

Central Services Extranet

This topic describes the characteristics of the central services extranet component of the VPN connectivity category.

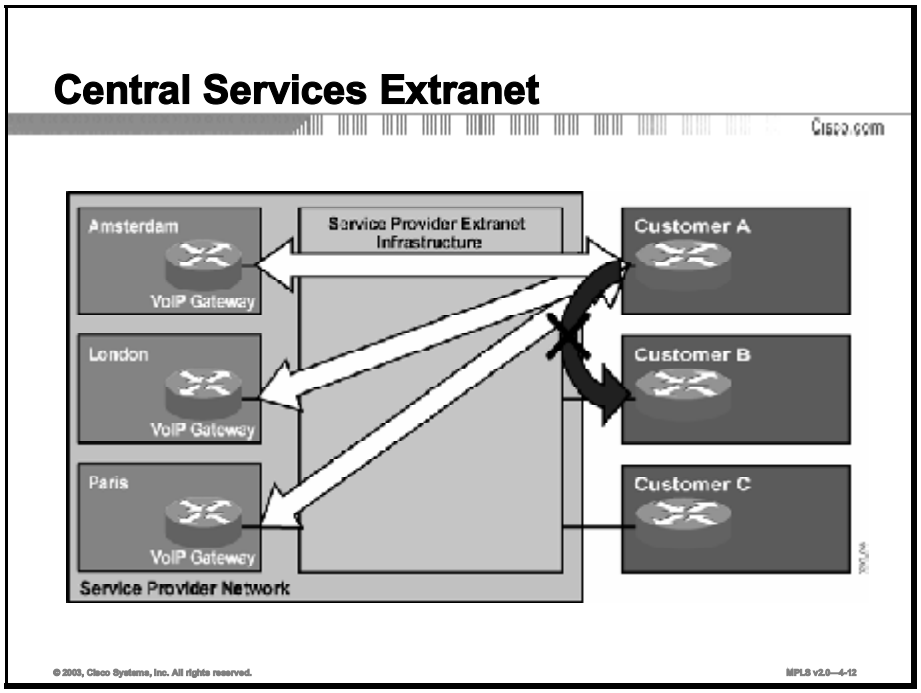

The figure shows a central services extranet implementing international Voice over IP (VoIP) service. Every customer of this service can access voice gateways in various countries but cannot access other customers using the same service.

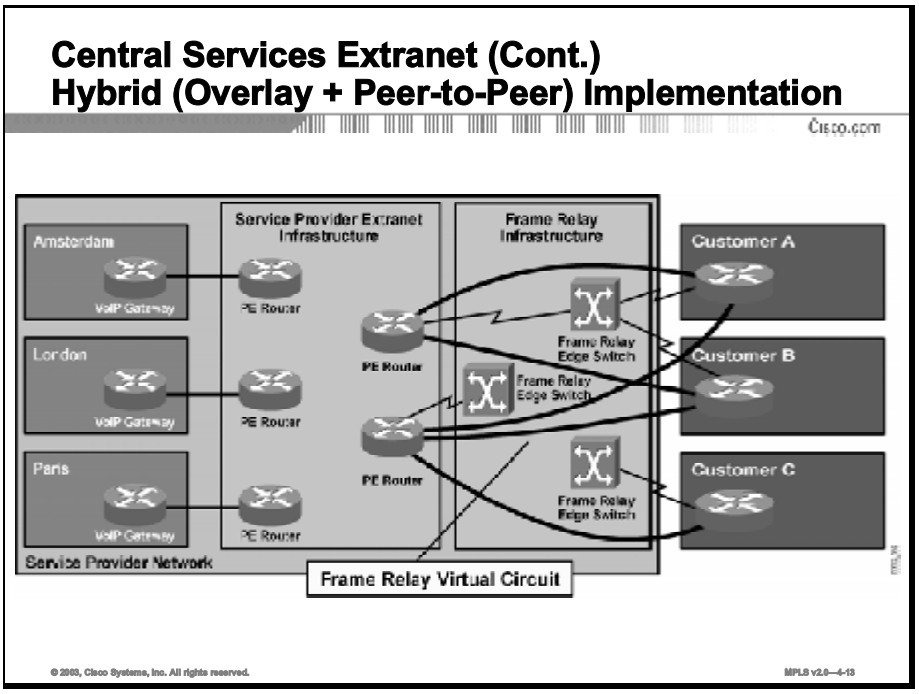

The network diagram shows an interesting scenario where peer-to-peer VPN and overlay VPN implementation can be used together to provide end-to-end service to the customer.

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.

Managed Network Implementation

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

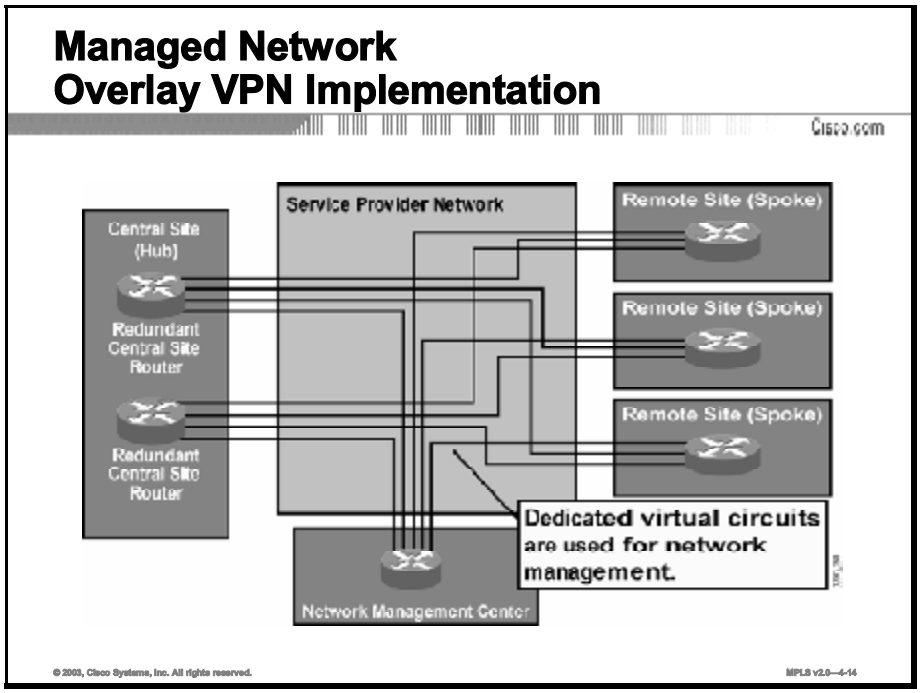

A managed network VPN is traditionally implemented in combination with overlay VPN services. Dedicated virtual circuits are deployed between any managed CE router and the central network management system (NMS) router to which the NMS is connected.

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.

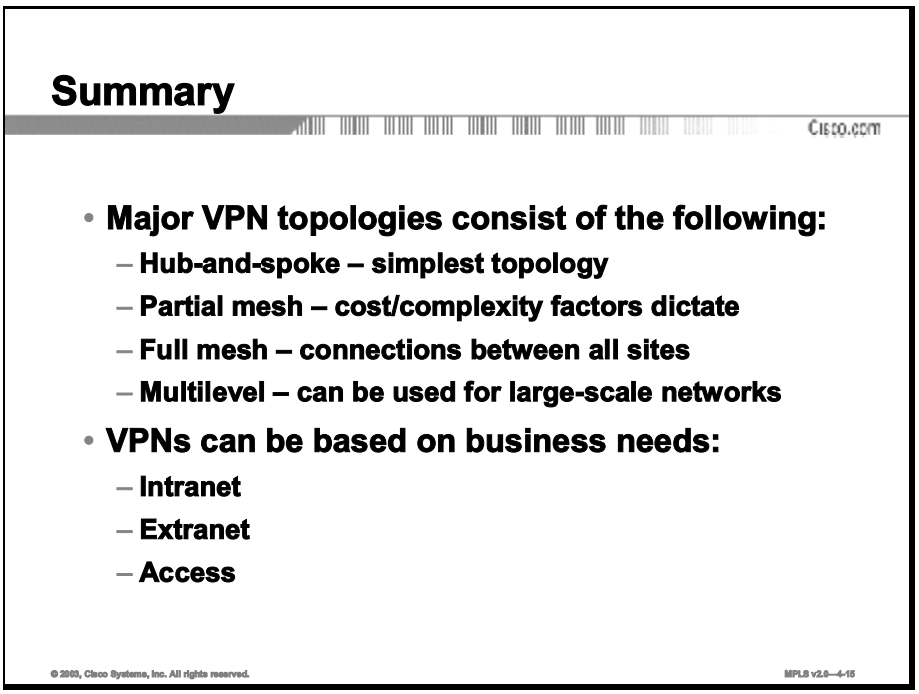

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Name the VPN topology that has many sites connecting to a central site.

_____

_____

Q2) When you are using a dynamic routing protocol such as RIP in a redundant hub-and-spoke topology, which of the following is true?

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.

Q3) Identify the criteria that a customer should consider when determining where virtual circuits are established in a partial mesh topology.

_____

_____

_____

Q4) Which component of the VPN business category is used to connect different organizations?

A) intranet VPNs

B) Internet VPNs

C) access VPNs

D) extranet VPNs

Q5) Which component of the VPN business category relies on security mechanisms to ensure protection of participating individual organizations?

A) intranet VPNs

B) Internet VPNs

C) access VPNs

D) extranet VPNs

Q6) Which implementation of the VPN business category provides the most cost-effective model?

A) overlay

B) peer-to-peer

Q7) Which component of the VPN connectivity category provides full connectivity between sites?

A) simple

B) overlapping

C) central services

D) managed services

Q8) Describe the connectivity in a central services extranet.

Q9) Describe the connectivity in a managed network VPN.

Quiz Answer Key

Q1) hub-and-spoke

Relates to: Overlay VPN Category

Q2) C

Relates to: Hub-and-Spoke Overlay VPN Topology

Q3) 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.

Relates to: Partial Mesh Overlay VPN Topology

Q4) D

Relates to: VPN Business Category

Q5) D

Relates to: VPN Business Category

Q6) B

Relates to: Extranet VPNs

Q7) A

Relates to: VPN Connectivity Category

Q8) 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.

Relates to: Central Services Extranet

Q9) Dedicated virtual circuits are deployed between any managed CE router and the central NMS router.

Relates to: Managed Network Implementation

MPLS VPN Architecture

Overview

This lesson explains the MPLS VPN architecture, as well as route information propagation, route distinguishers (RDs), route targets (RTs), and virtual routing tables.

Relevance

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

This lesson describes the major architectural components of an MPLS VPN, identifying the functions of route information propagation, RDs, RTs, and virtual routing tables.

Upon completing this lesson, you will be able to:

- Describe the MPLS VPN architecture, identifying network elements such as the CE, PE, and P routers
- Describe how MPLS VPNs are implemented in the PE router architecture
- Describe different methods of propagating routing information across the P-network, identifying the drawbacks of each method
- Explain the need for route distinguishers and how they are implemented
- Explain the need for route targets and how they are implemented
- Describe how complex VPNs have redefined the meaning of VPNs
- Describe the impact of complex VPN topologies on virtual routing tables

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components as well as basic routing and ATM principles and concepts

Outline

This lesson includes these topics:

- Overview
- MPLS VPN Architecture
- PE Router Architecture
- Propagation of Routing Information Across the P-Network
- Route Distinguishers
- Route Targets
- Virtual Private Networks Redefined
- Impact of Complex VPN Topologies on Virtual Routing Tables
- Summary
- Quiz

MPLS VPN Architecture

This topic describes the MPLS VPN architecture, identifying network elements such as the CE, PE, and P routers.

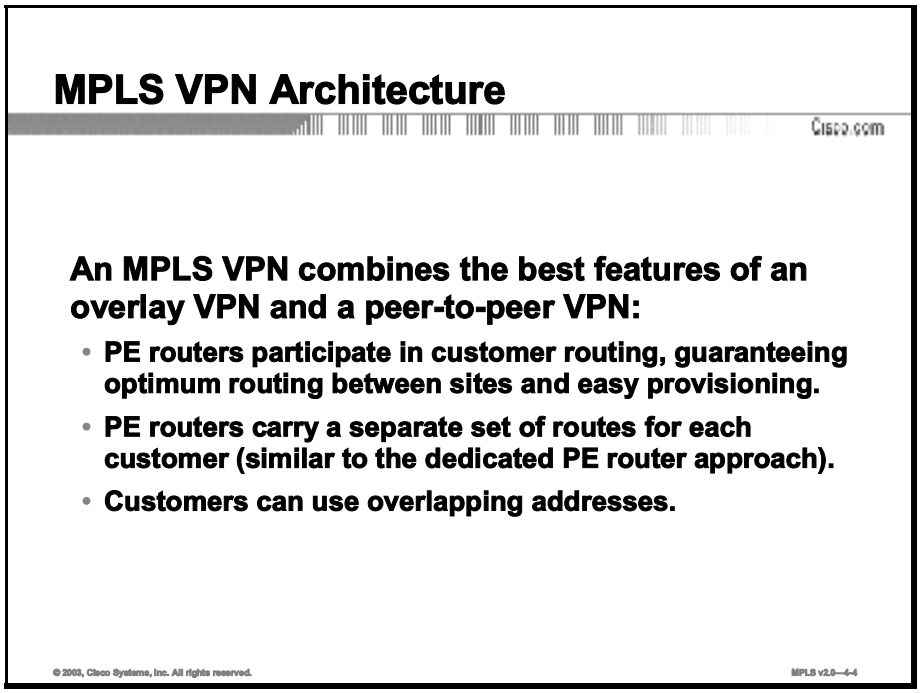

The MPLS VPN architecture offers service providers a peer-to-peer VPN architecture that combines the best features of overlay VPNs (support for overlapping customer address spaces) with the best features of peer-to-peer VPNs:

- 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.

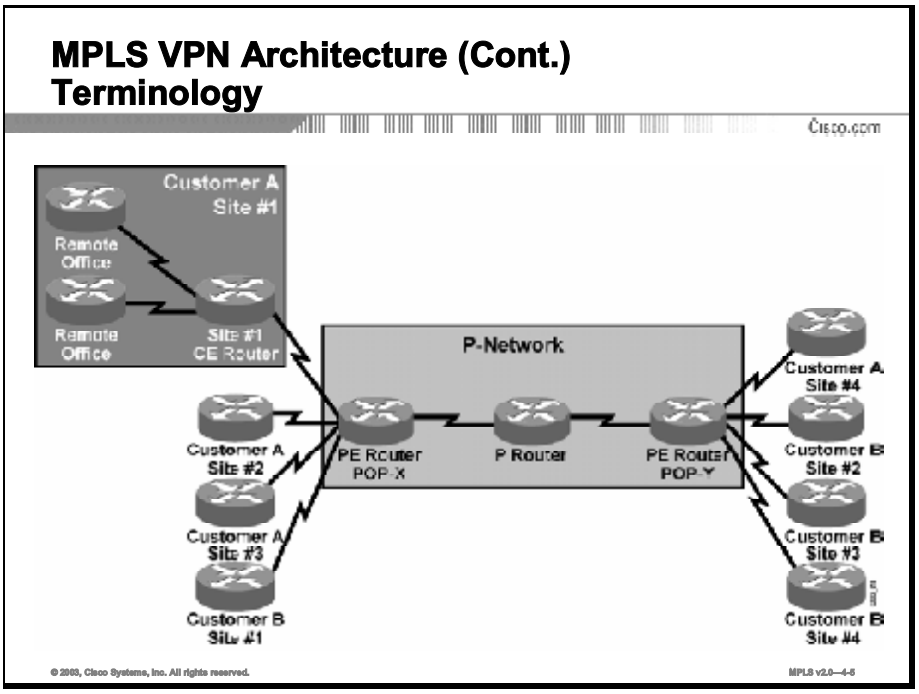

As discussed earlier, MPLS VPN terminology divides the overall network into a customer-controlled part (the customer network, or C-network) and a provider-controlled part (the provider network, or P-network). Contiguous portions of the C-network are called sites and are linked with the P-network via 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.

PE Router Architecture

This topic describes how MPLS VPNs are implemented in the PE router architecture.

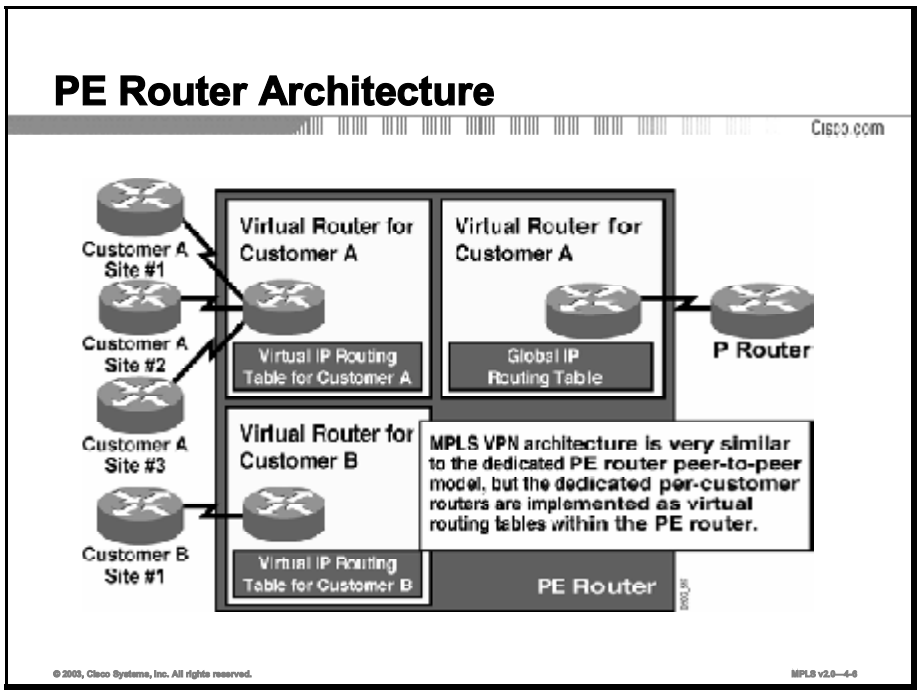

The architecture of a PE router in an MPLS VPN is very similar to the architecture of a POP in the dedicated PE router peer-to-peer model; the only difference is that the whole architecture is condensed into one physical device. Each customer is assigned an independent routing table (virtual routing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

Propagation of Routing Information Across the P-Network

This topic describes the different methods of propagating routing information across the P-network, identifying the drawbacks of each method.

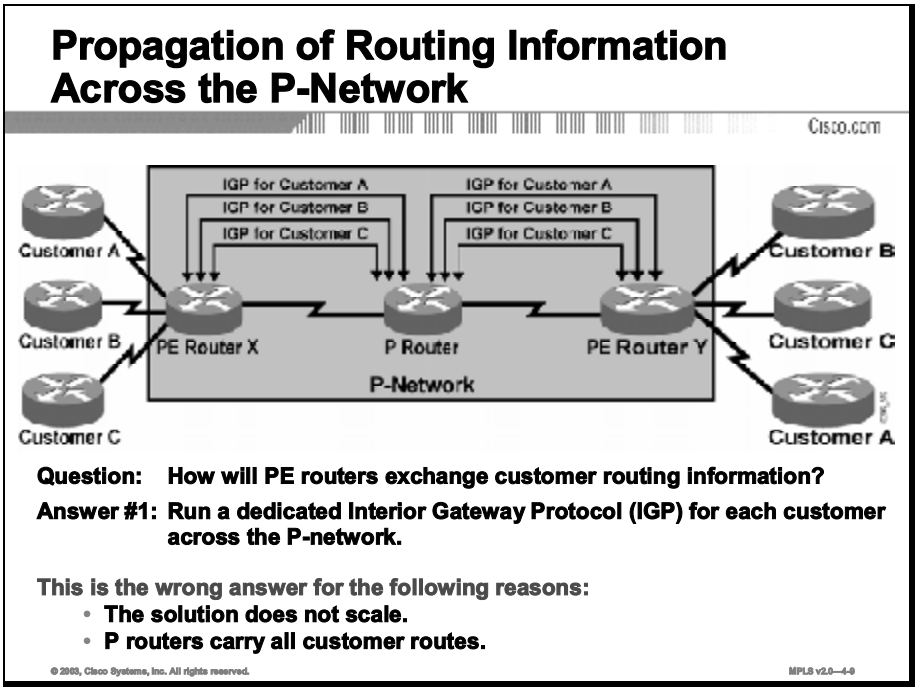

Although virtual routing tables provide isolation between customers, the data from these routing tables still needs to be exchanged between PE routers to enable data transfer between sites attached to different PE routers. Therefore, a routing protocol is needed that will transport all customer routes 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.

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

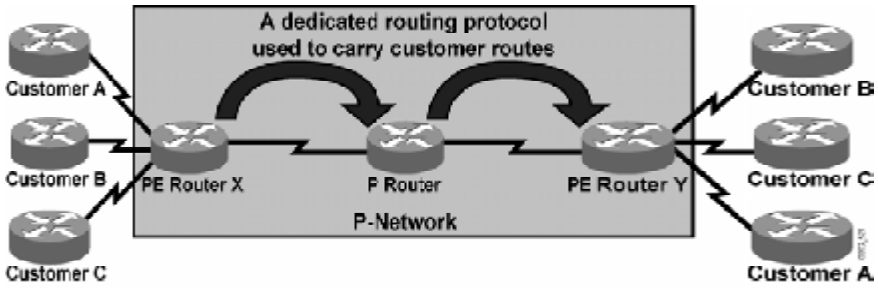

Question: How will PE routers exchange customer routing information?

Answer #2: Run a single routing protocol that will carry all customer routes inside the provider backbone.

Better answer, but still not good enough:

- **P routers carry all customer routes.**

MPLS v2.0—4-12

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, so the proposal retains some of the same scalability issues of the previous one.

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

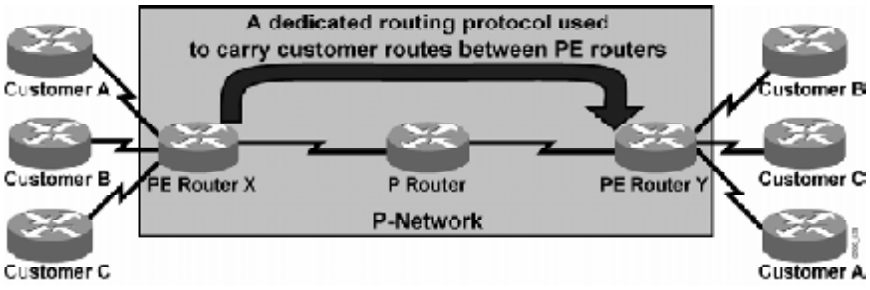

Question: **How will PE routers exchange customer routing information?**

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:

- 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.

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

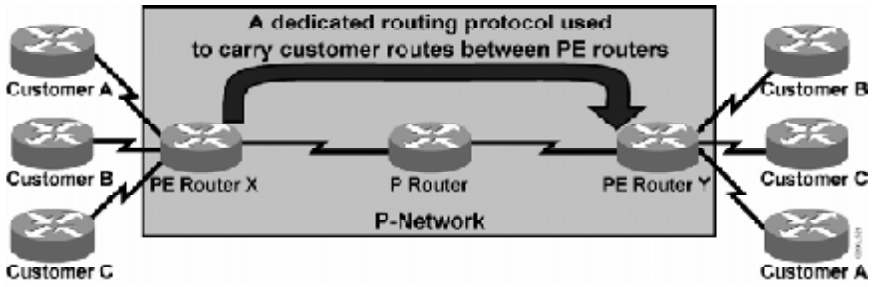

Question: **Which protocol can be used to carry customer routes between PE routers?**

Answer: **The number of customer routes can be very large. BGP is the only routing protocol that can scale to a very large number of routes.**

Conclusion:

BGP is used to exchange customer routes directly between PE routers.

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. And, in fact, BGP is used in MPLS VPN architecture to transport customer routes directly between PE routers.

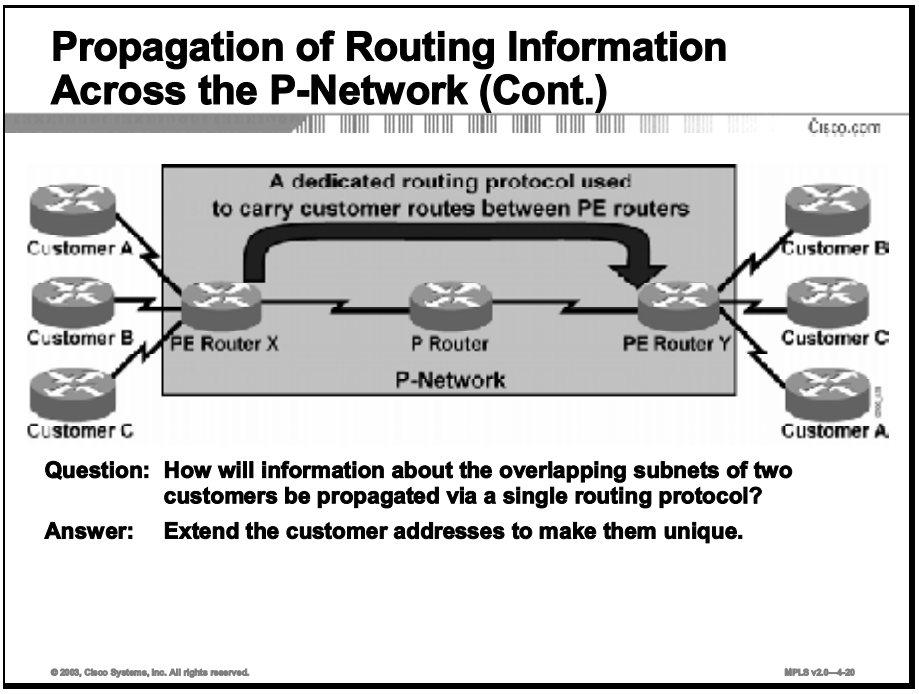

MPLS VPN architecture differs in an important way from traditional peer-to-peer VPN solutions: the support of overlapping customer address spaces.

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 route distinguisher (RD) is used in MPLS VPNs to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE routers.

Route Distinguishers

This topic describes what a route distinguisher (RD) is and how it is used.

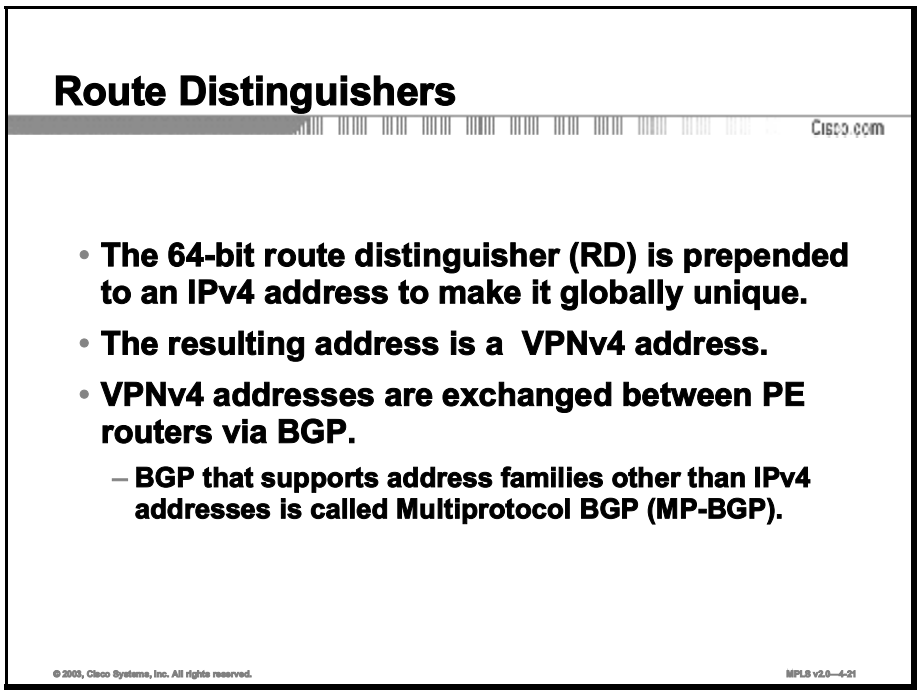

The RD is used only to transform nonunique 32-bit customer IP version 4 (IPv4) addresses into unique 96-bit VPNv4 addresses (also called VPN IPv4 addresses).

VPNv4 addresses are exchanged only between PE routers; they are never used between CE routers. BGP between PE routers must therefore support exchange of traditional IPv4 prefixes as well as 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.

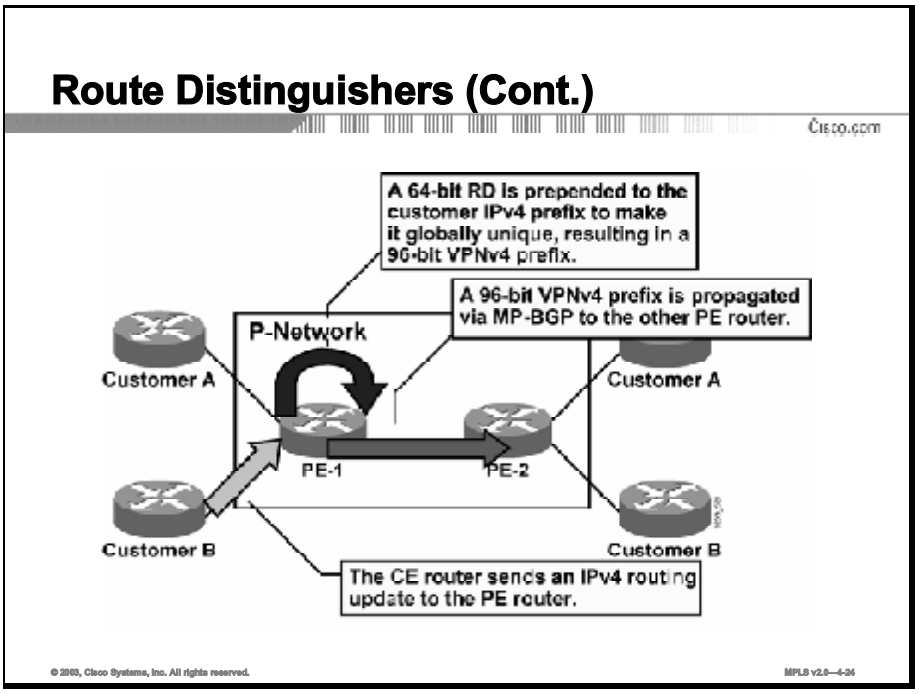

Customer route propagation across an MPLS VPN network is done using the following process:

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.

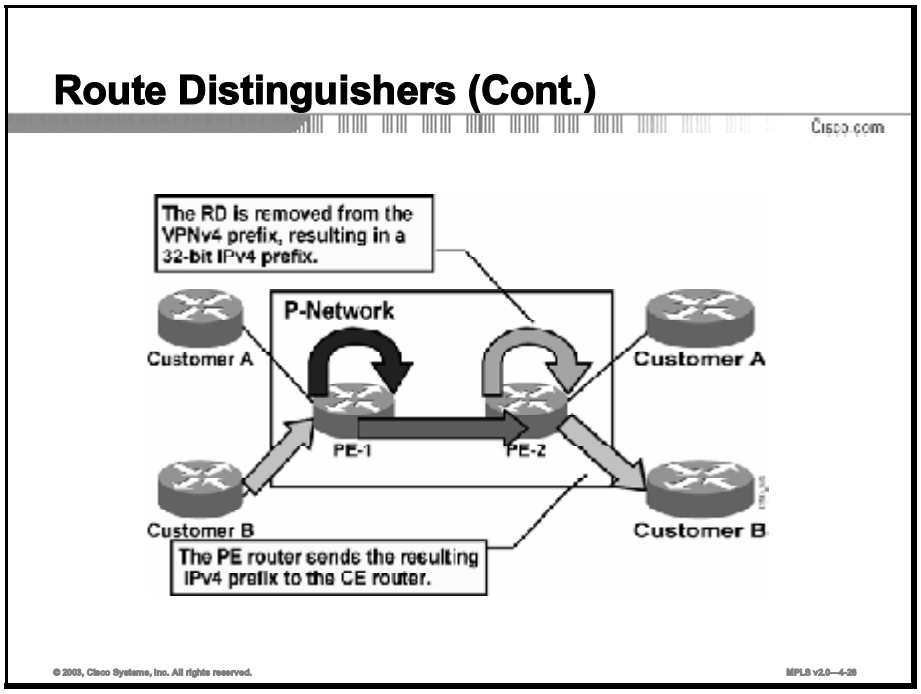

Step 4 The receiving PE routers strip the RD from the VPNv4 prefix, resulting in an IPv4 prefix.

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

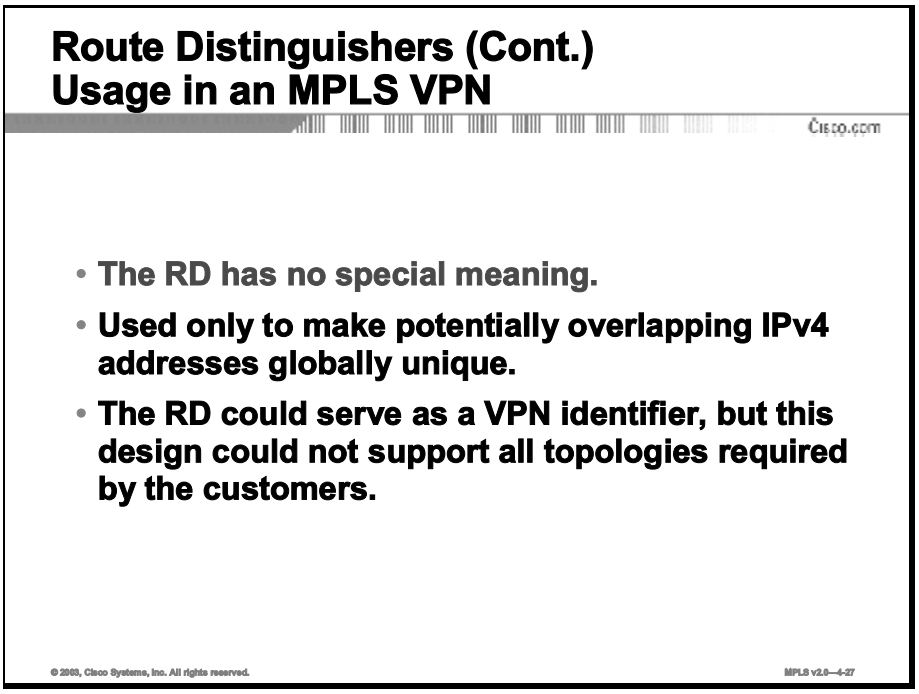

The RD has no special meaning or role in MPLS VPN architecture; its only function is to make overlapping IPv4 addresses globally unique.

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. It 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.

Note	This topic will be discussed in greater detail later in this lesson.

Route Targets

This topic explains the need for route targets (RTs) and how they are implemented.

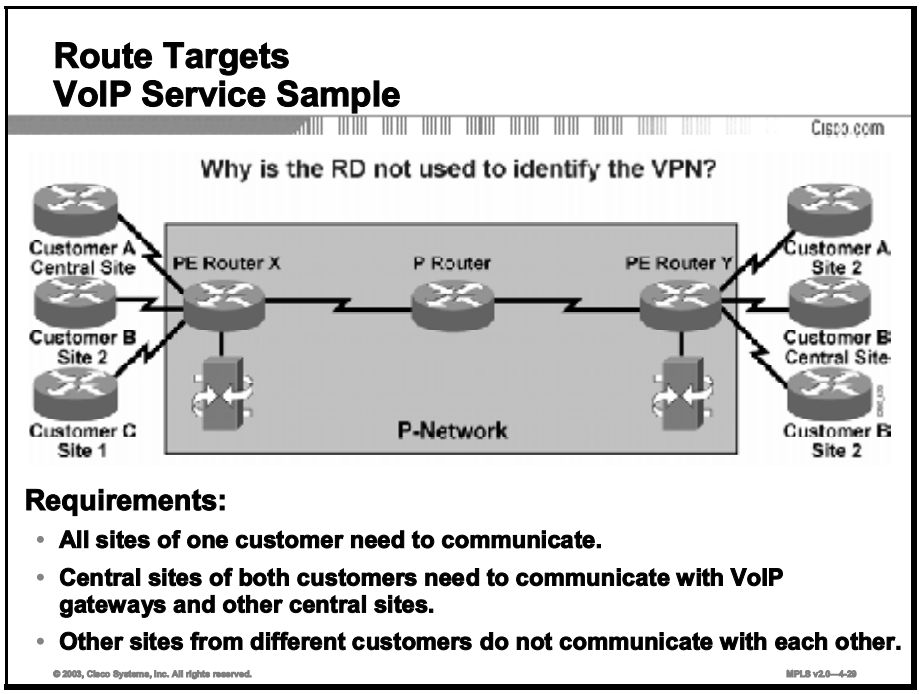

To illustrate the need for a more versatile VPN indicator than the RD, consider the VoIP service in the figure. The connectivity requirements of this service are as follows:

- 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) as well as with other central sites to exchange intercompany voice calls.

Note	Additional security measures would have to be put in place at central sites to make sure 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.

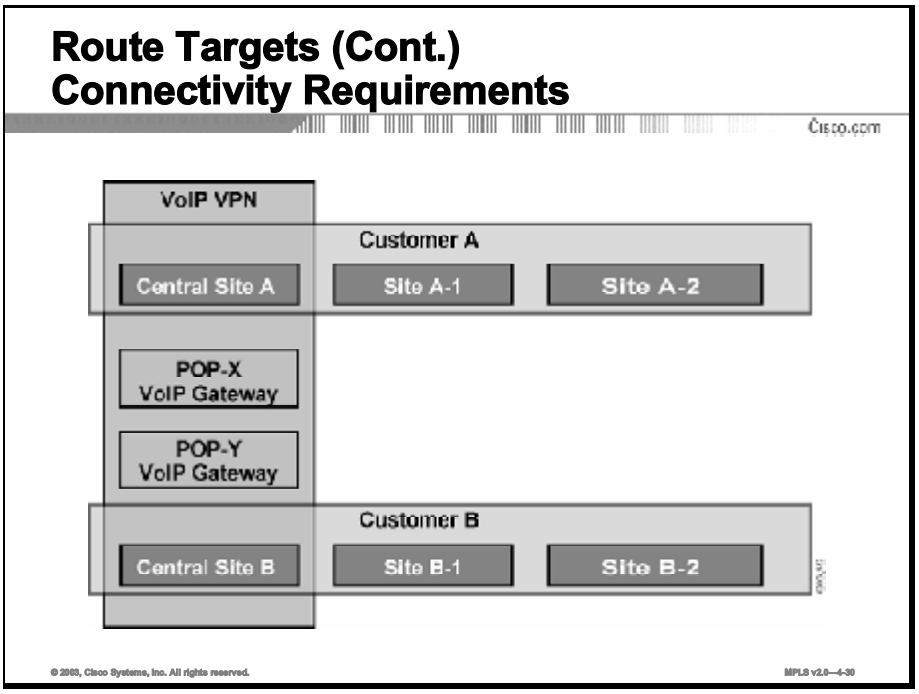

The connectivity requirements of the VoIP service are illustrated in the figure. Three VPNs are needed to implement the desired connectivity: two customer VPNs and a shared VoIP VPN. Central customer sites participate in the customer VPN as well as in the VoIP VPN.

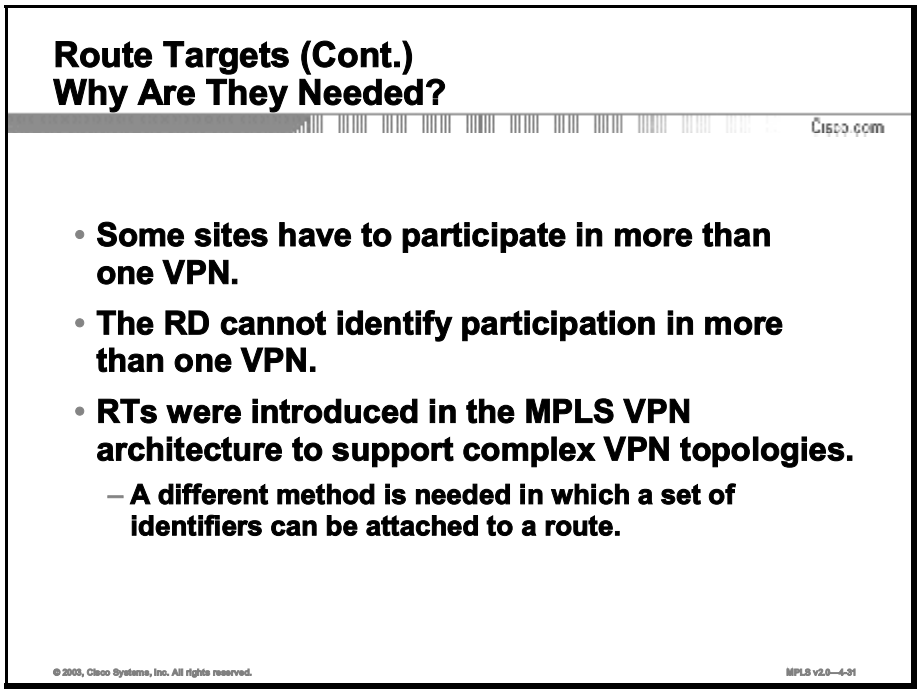

The RD (again, a single entity prepended to an IPv4 route) cannot indicate that a site participates in more than one VPN. A method is needed in which a set of VPN identifiers can be attached to a route to indicate its membership in several VPNs.

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

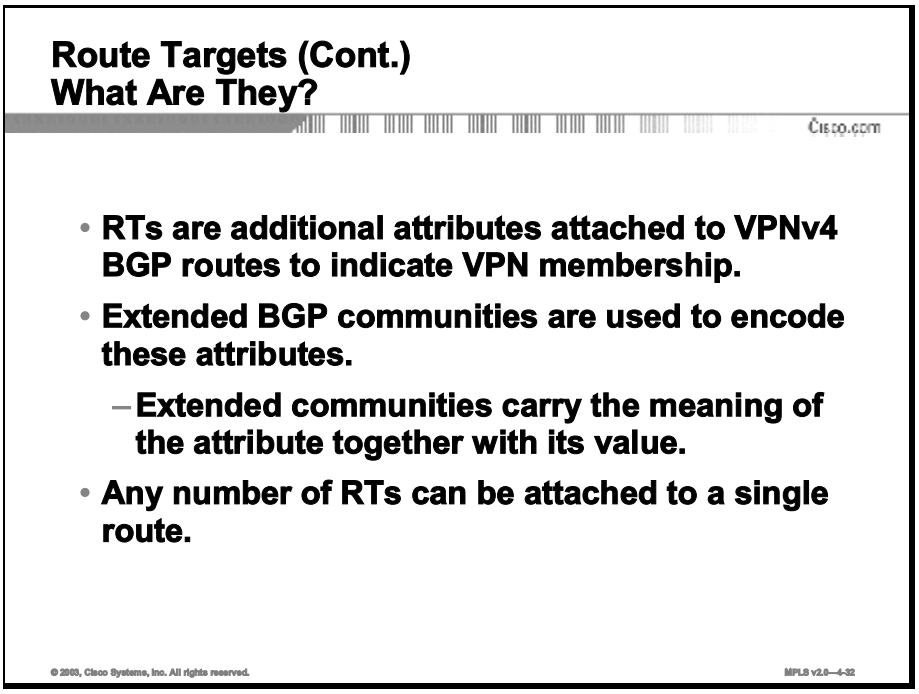

RTs are attributes that are attached to a VPNv4 BGP route to indicate its VPN membership. The extended BGP communities of a routing update is used to carry the RT of that update, thus identifying which VPN the update belongs to.

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 them useful for a number of different applications, such as MPLS VPN RTs.

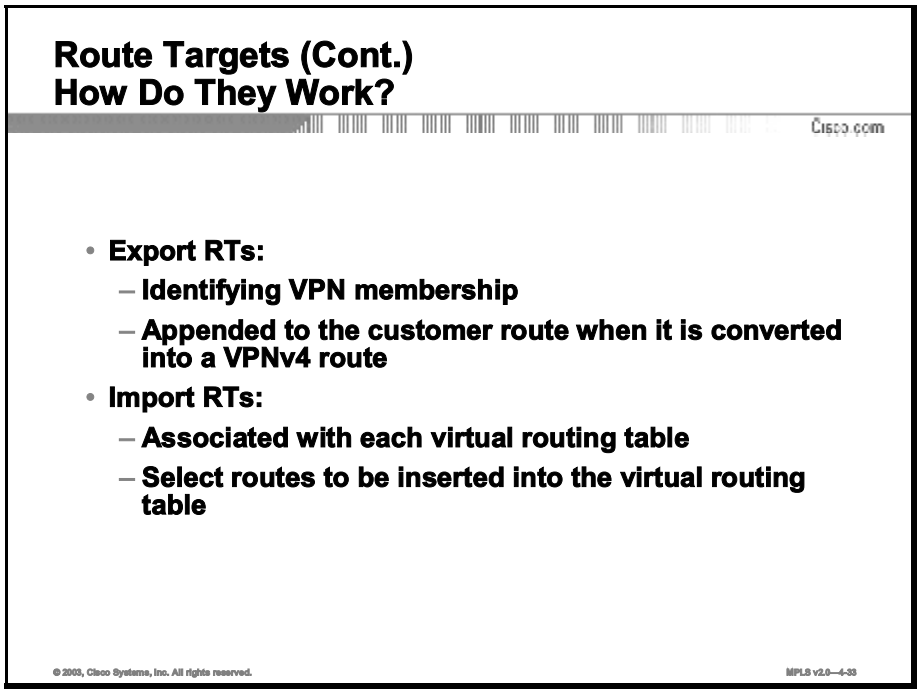

MPLS VPN RTs are attached to a customer route at the moment that it is converted from an IPv4 route to a VPNv4 route by the PE router. The RTs attached to the route are called export RTs and are configured separately for each virtual routing table in a PE router. Export RTs identify a set of VPNs in 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.

Virtual Private Networks Redefined

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

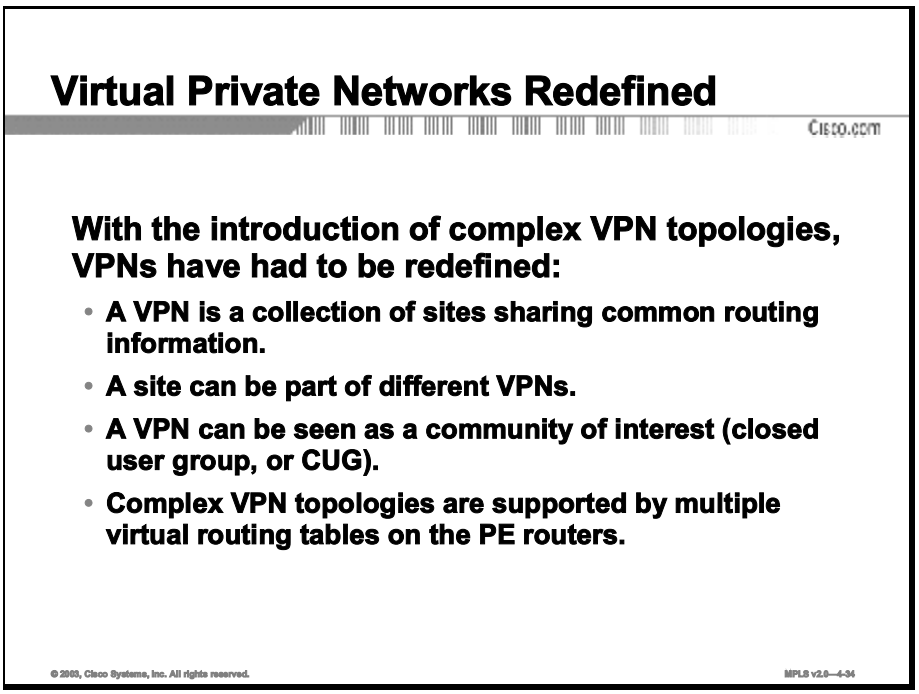

With the introduction of complex VPN topologies, the definition of a VPN has needed to be changed. A VPN is simply a collection of sites sharing common routing information. In traditional switched WAN terms (for example, in X.25 terminology), such a concept would be called a closed user group (CUG).

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.

Impact of Complex VPN Topologies on Virtual Routing Tables

This topic describes the impact of complex VPN topologies on the virtual routing tables.

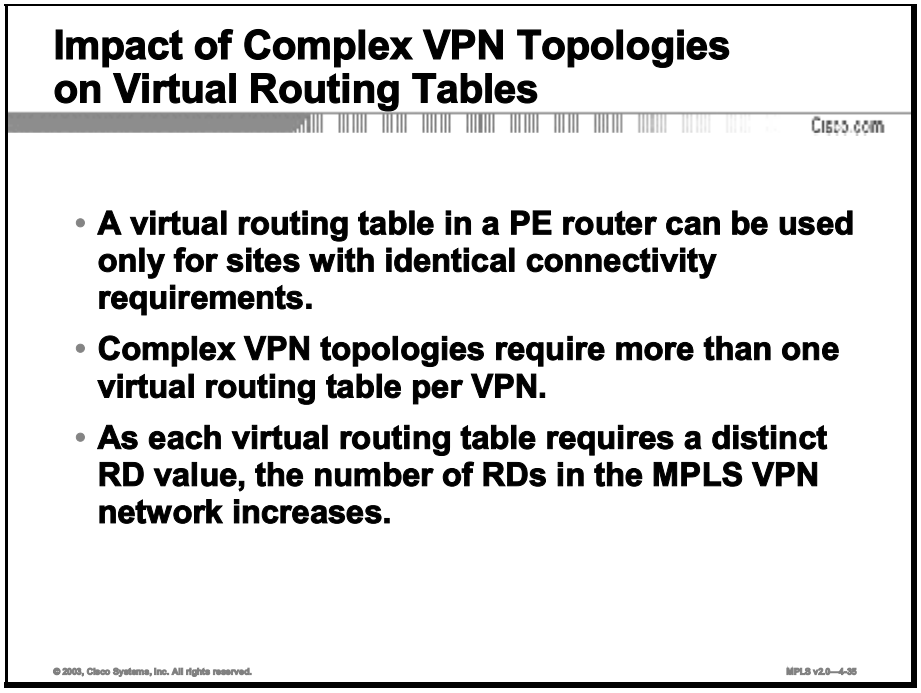

A single virtual routing table can be used only for sites with identical connectivity requirements. Complex VPN topologies, therefore, require more than one virtual routing table per VPN.

Note	If sites with different requirements are associated with the same virtual routing table, some of them 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.

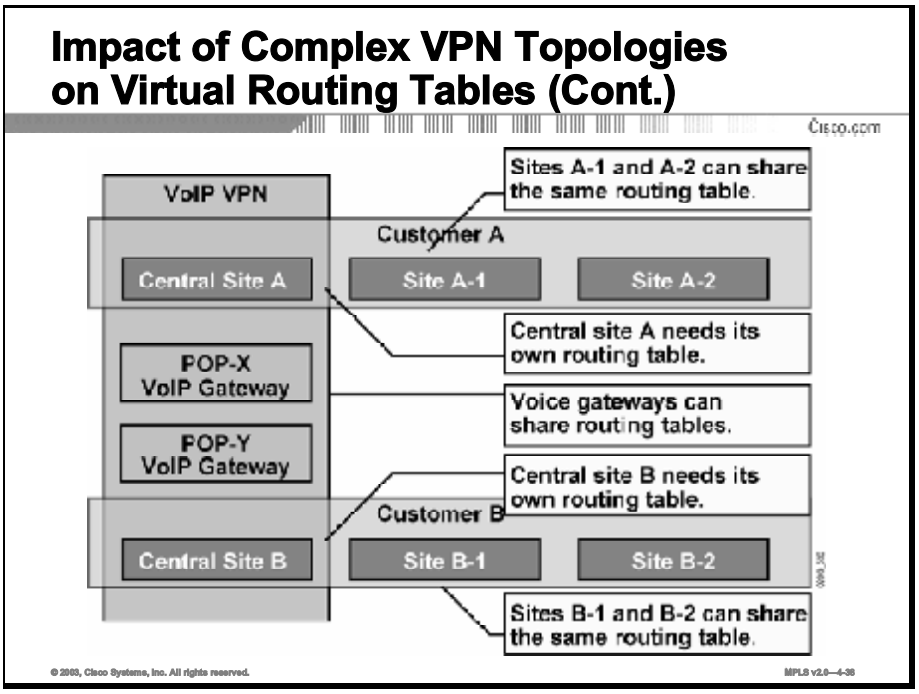

To illustrate the requirements for multiple virtual routing tables, consider a VoIP service with three VPNs (customer A, customer B, and a VoIP VPN). The virtual routing table needs of this service are as follows:

- 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.

So, 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.

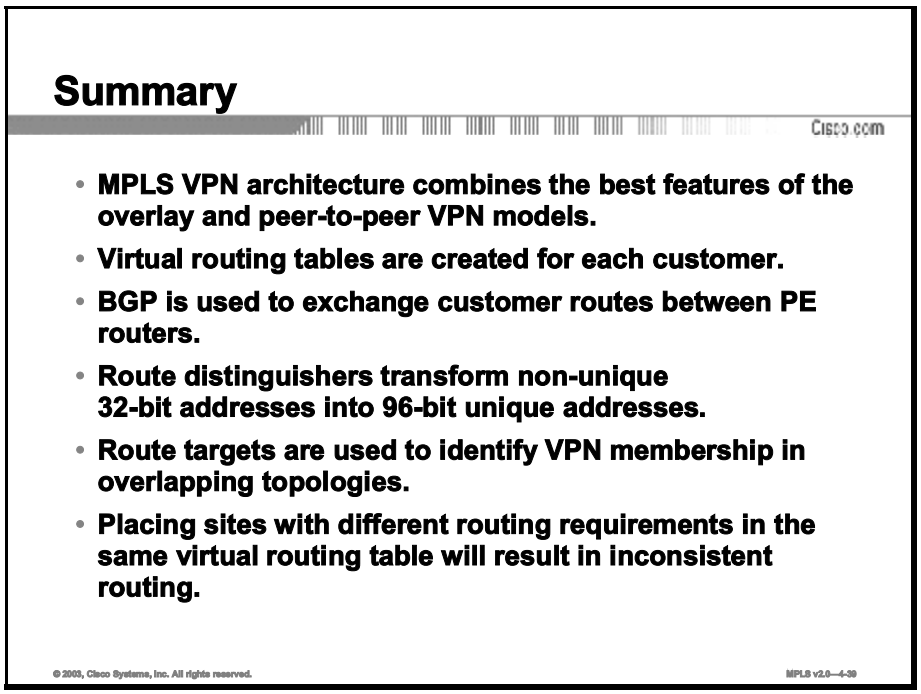

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which routers are MPLS VPN-aware?

Q2) Which traditional VPN module can the architecture of a PE router in an MPLS VPN be compared to?

Q3) Which protocol is used to transport customer routes directly between PE routers?

A) RIP

B) VPN

C) BGP

D) OSPF

Q4) What is the function of the RD in an MPLS VPN?

Q5) What is the function of the RT in MPLS VPNs?

Q6) How has the introduction of complex VPN topologies redefined the meaning of a VPN?

Q7) What could happen if two different sites with different requirements are associated with the same virtual routing table?

Quiz Answer Key

Q1) P routers

Relates to: MPLS VPN Architecture

Q2) the dedicated PE router peer-to-peer model

Relates to: PE Router Architecture

Q3) C

Relates to: Propagation of Routing Information Across the P-Network

Q4) The RD is used to transform the nonunique IP addresses of the customer into unique VPNv4 addresses.

Relates to: Route Distinguishers

Q5) 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.

Relates to: Route Targets

Q6) 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.

Relates to: Virtual Private Networks Redefined

Q7) Some of the sites might be able to access destinations that they should not be able to access.

Relates to: Impact of Complex VPN Topologies on Virtual Routing Tables

MPLS VPN Routing Model

Overview

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

Relevance

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 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

This lesson identifies the routing requirements for MPLS VPNs by describing, from the perspective of the customer and provider, how routing tables appear on PE routers, how customer routes are propagated, and how end-to-end information flows.

Upon completing this lesson, you will be able to:

- Describe the routing requirements of an MPLS VPN
- Describe the MPLS VPN routing model of 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

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components, and an understanding of basic VPN implementations

Outline

This lesson includes these topics:

- Overview
- MPLS VPN Routing Requirements
- MPLS VPN Routing
- Support for Existing Internet Routing
- Routing Tables on PE Routers
- End-to-End Routing Update Flow
- Route Distribution to CE Routers
- Summary
- Quiz

MPLS VPN Routing Requirements

This topic describes the routing requirements for MPLS VPNs.

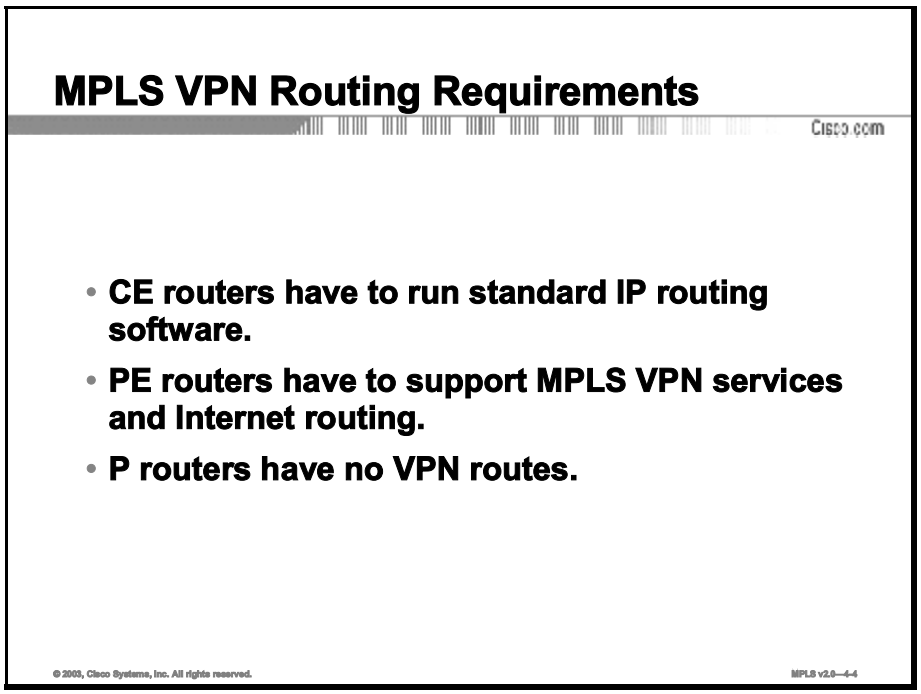

The designers of MPLS VPN technology were faced with the following routing requirements:

- CE routers should not be MPLS VPN-aware. They 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.

MPLS VPN Routing

This topic identifies the MPLS VPN routing model of CE routers, PE routers, and P routers.

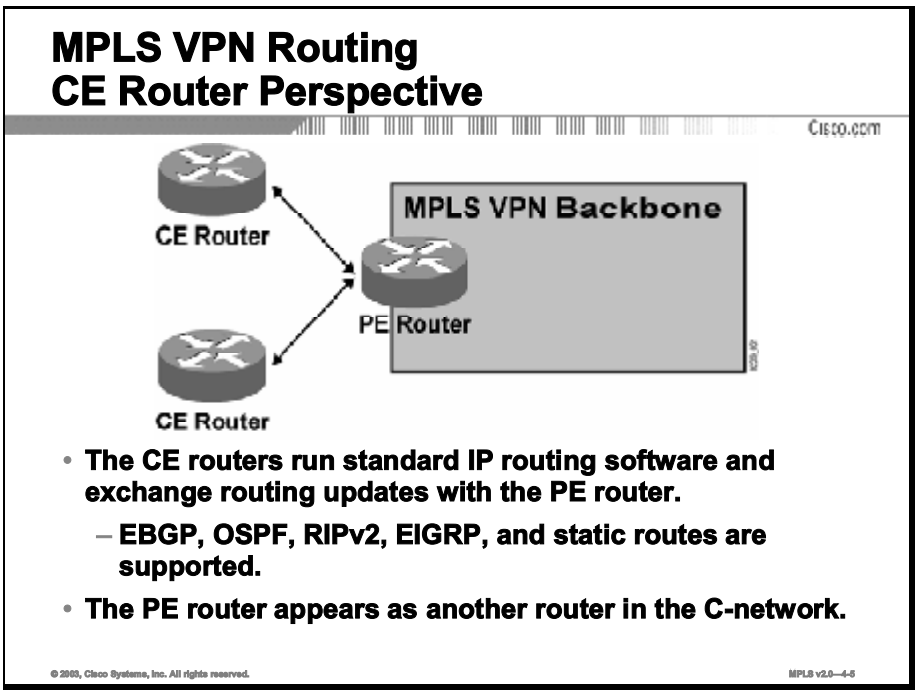

The MPLS VPN backbone should look like a standard corporate backbone to the CE routers. The CE routers run standard IP routing software and exchange routing updates with the PE routers, which appear to them as normal routers in the C-network.

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 (cont.)
Overall Customer Perspective

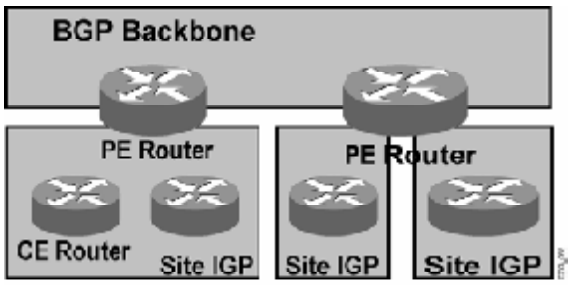

- **To the customer, the PE routers appear as core routers connected via a BGP backbone.**
- **The usual BGP and IGP design rules apply.**
- **The P routers are hidden from the customer.**

From the customer perspective, the MPLS VPN backbone looks like an intracompany 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 (Cont.)
P Router Perspective

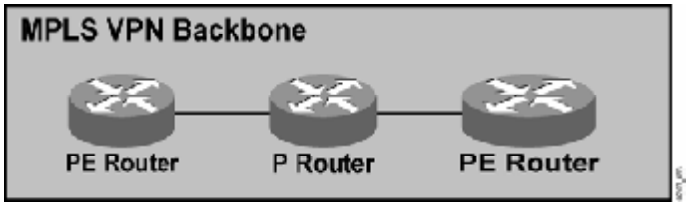

- **P routers do not participate in MPLS VPN routing and do not carry VPN routes.**
- **P routers run backbone IGP with the PE routers and exchange information about global subnets (core links and loopbacks).**

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. They run only a backbone Interior Gateway Protocol (IGP) with other P routers and with PE routers, and exchange information about core subnets. BGP deployment on P routers is not needed for proper MPLS VPN operation; it might be needed, however, to support traditional Internet connectivity that has not yet been migrated to MPLS.

MPLS VPN Routing (Cont.)
PE Router Perspective

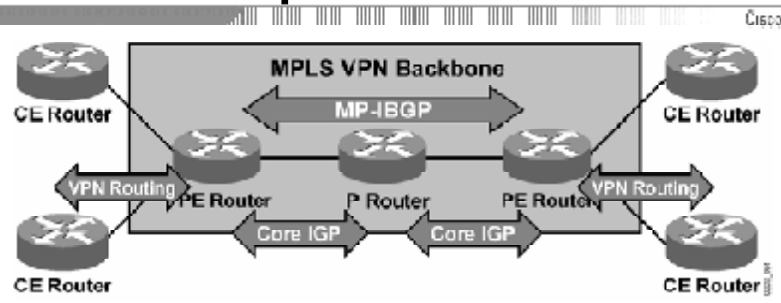

PE routers:

- **Exchange VPN routes with CE routers via per-VPN routing protocols**
- **Exchange core routes with P routers and PE routers via core IGP**
- **Exchange VPNv4 routes with other PE routers via MP-IBGP sessions**

The PE routers are the only routers in MPLS VPN architecture that see all routing aspects of the MPLS VPN:

- They exchange IPv4 VPN routes with CE routers via various routing protocols running in the virtual routing tables.
- They exchange VPNv4 routes via MP-IBGP sessions with other PE routers.
- They exchange core routes with P routers and other PE routers via core IGP.

Support for Existing Internet Routing

This topic describes how IPv4 is used to provide support for existing Internet routing.

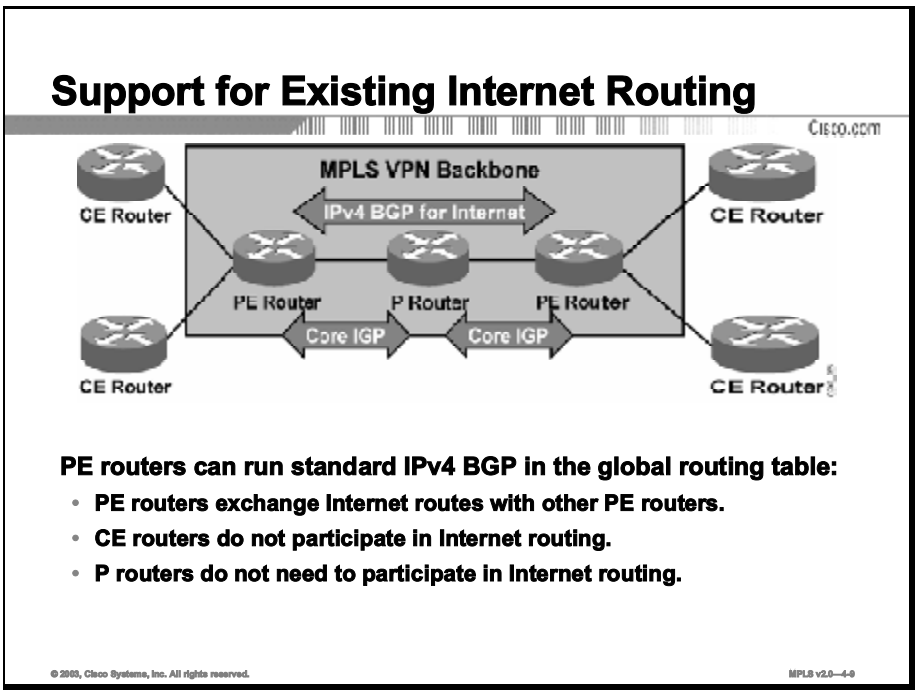

The routing requirements for PE routers also extend to supporting Internet connectivity—PE routers have to exchange Internet routes with other PE routers. The CE routers cannot participate in Internet routing if the Internet routing is performed in global address space. The P routers could participate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

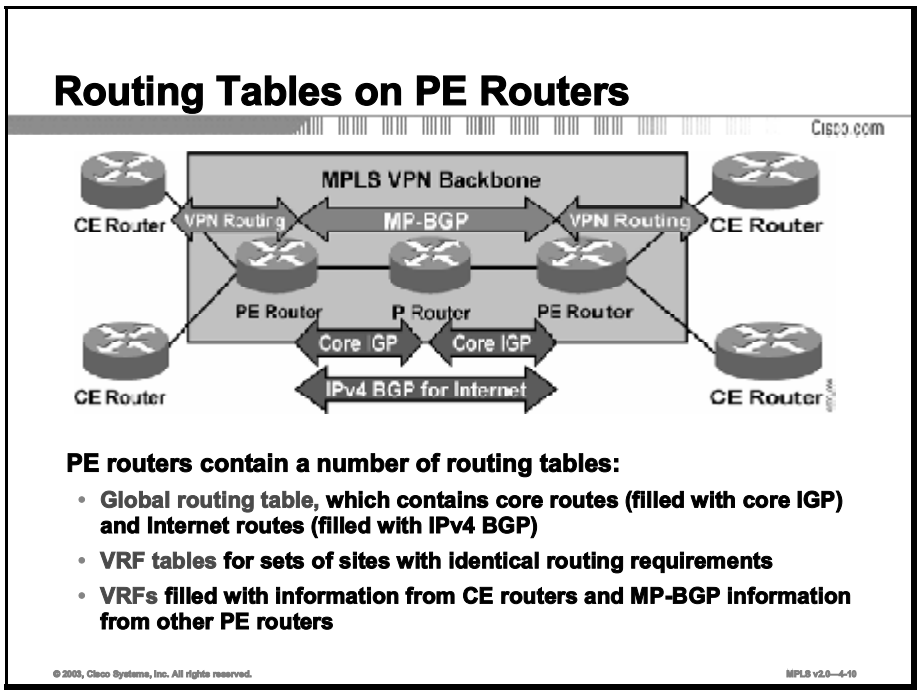

The PE routers fulfill various routing requirements imposed on them by using a number of IP routing tables:

- 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.

End-to-End Routing Update Flow

This topic describes the end-to-end flow of routing updates in an MPLS VPN.

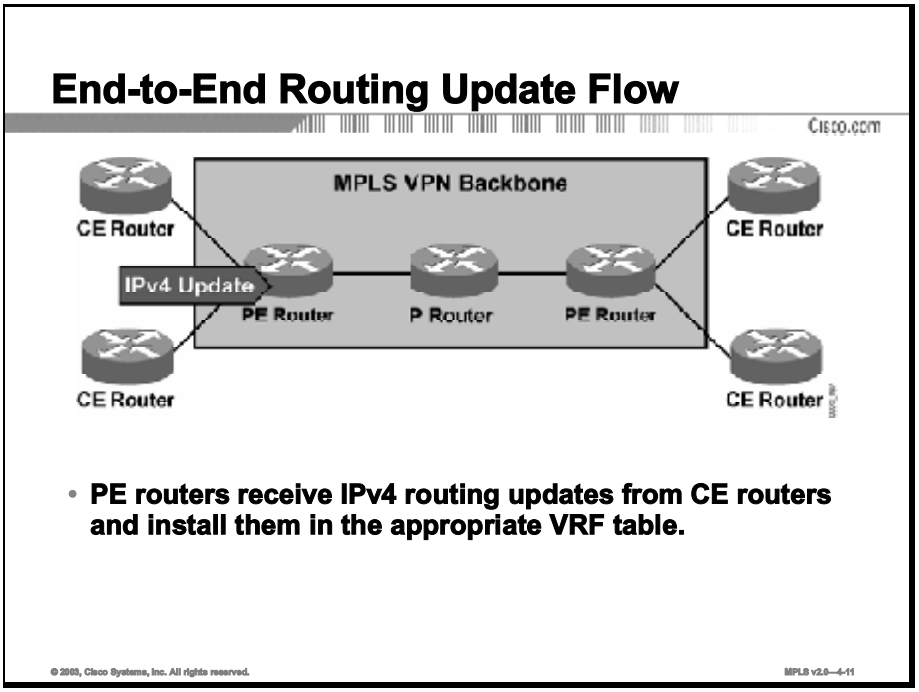

The next few figures provide an overview of end-to-end routing information flow in an MPLS VPN network. The PE routers receive IPv4 routing updates from the CE routers and install them in the appropriate VRF table.

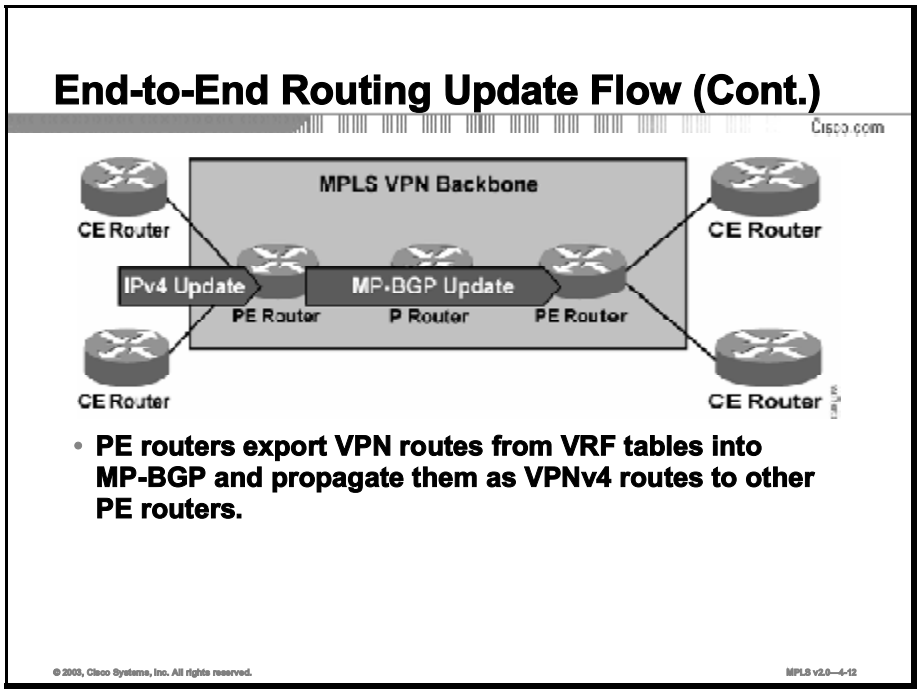

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers.

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.

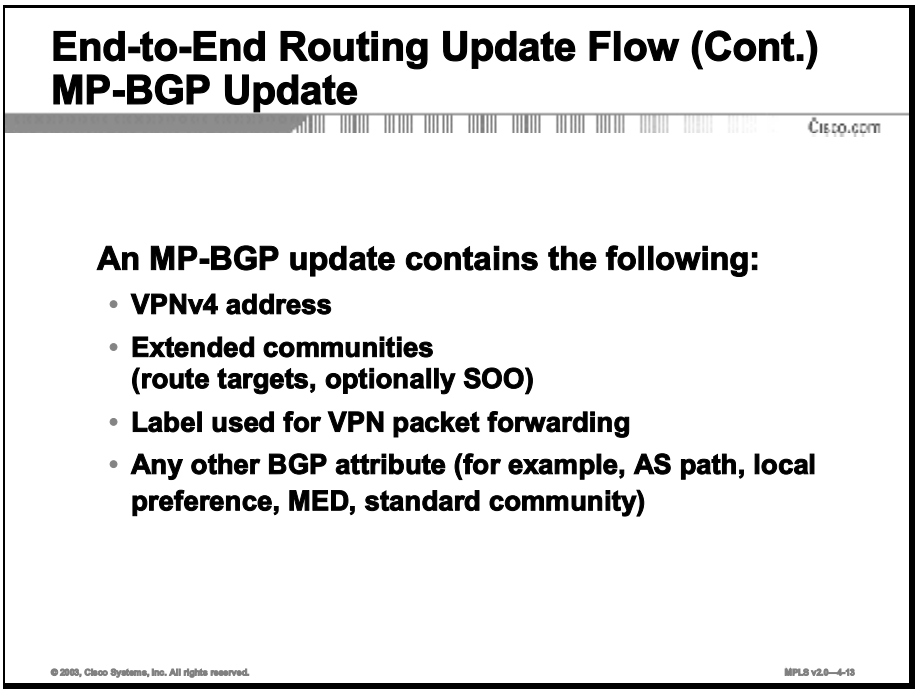

An MP-BGP update exchange between PE routers contains the following:

- VPNv4 address
- Extended BGP communities (RTs required; Site of Origin, or SOO, optional)
- Label used for VPN packet forwarding (the "MPLS VPN Packet Forwarding" lesson follows this lesson and 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.)

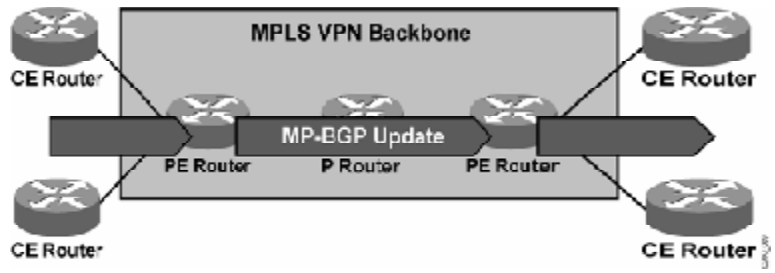

- **Receiving PE router imports incoming VPNv4 routes into the appropriate VRF based on route targets attached to the routes.**
- **Routes installed in VRF are propagated to CE routers.**

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 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.

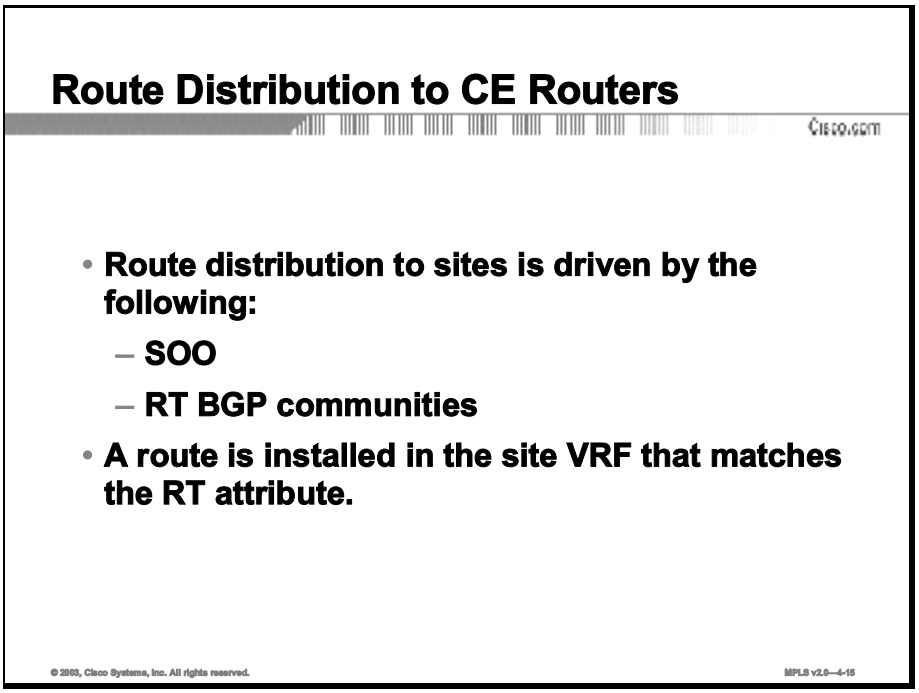

VPNv4 routes are installed into VRFs on the receiving PE router—the incoming VPNv4 route is imported into the VRF only if at least one RT attached to the route matches at least one import RT configured in the VRF.

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.

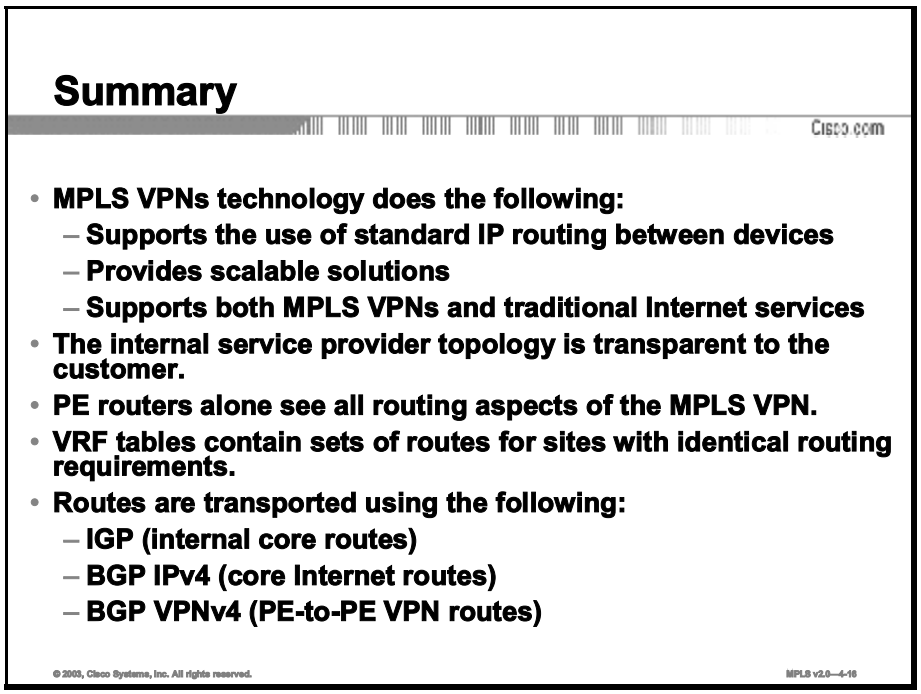

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which routing protocol does the CE router run?

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

Q2) Which routers exchange VPNv4 routes?

A) P

B) CE

C) PE

Q3) Which protocol would a PE router use to support an existing Internet routing scheme?

A) IS-IS

B) EIGRP

C) BGP IPv4

D) BGP VPNv4

Q4) Identify the routing tables implemented in the PE router to support an MPLS VPN and describe their contents.

Q5) What BGP function do MPLS VPNs use to transport RTs?

Q6) How does the PE router know in which VRF table to install received routes for a customer?

Quiz Answer Key

Q1) A

Relates to: MPLS VPN Routing Requirements

Q2) C

Relates to: MPLS VPN Routing

Q3) C

Relates to: Support for Existing Internet Routing

Q4) global IP routing table—contains all core IGP routes and the IPv4 routes; VRFs—contain CE routes and VPNv4 routes

Relates to: Routing Tables on PE Routers

Q5) extended communities

Relates to: End-to-End Routing Update Flow

Q6) the RT contained in the extended BGP community

Relates to: Route Distribution to CE Routers

MPLS VPN Packet Forwarding

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.

Relevance

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

This lesson describes how packets are forwarded in an MPLS VPN environment, identifying how VPN labels get propagated and explaining the effects of summarization in the core.

Upon completing this lesson, you will be able to:

- Describe the end-to-end MPLS VPN forwarding mechanisms, detailing label assignments
- Describe the operation of PHP in an MPLS VPN environment
- Describe the MPLS VPN and backbone label propagation process
- Describe the effects of MPLS VPNs on label propagation
- Describe the effects of MPLS VPNs on packet forwarding

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components, and an understanding of basic VPN implementations

Outline

This lesson includes these topics:

- Overview
- VPN Packet Forwarding Across an MPLS VPN Backbone
- VPN Penultimate Hop Popping
- VPN Label Propagation
- Effects of MPLS VPNs on Label Propagation
- Effects of MPLS VPNs on Packet Forwarding
- Summary
- Quiz

VPN Packet Forwarding Across an MPLS VPN Backbone

This topic describes the end-to-end MPLS VPN forwarding mechanisms, detailing label assignments.

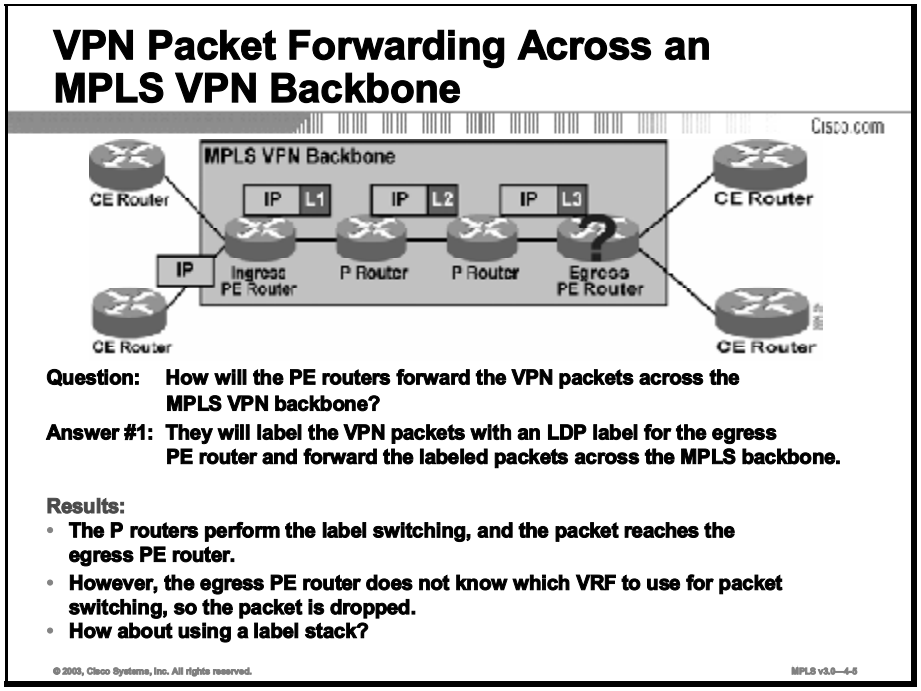

An MPLS-oriented approach to MPLS VPN packet forwarding across the MPLS VPN backbone would be to label the customer packet with the label assigned by Label Distribution Protocol (LDP) for the egress PE router. The core routers consequently would never see the customer IP packet; instead, they would see just a labeled packet targeted toward the egress PE router. They 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.)

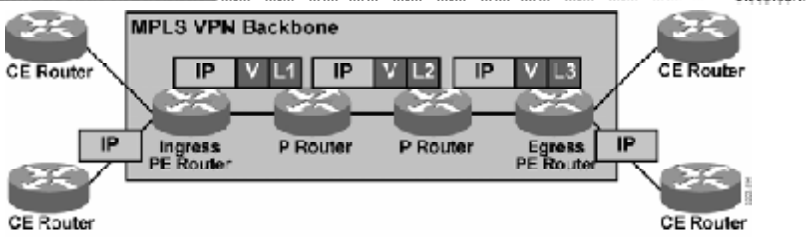

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #2: They 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 all 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.

VPN Penultimate Hop Popping

This topic describes operation of PHP in an MPLS VPN environment.

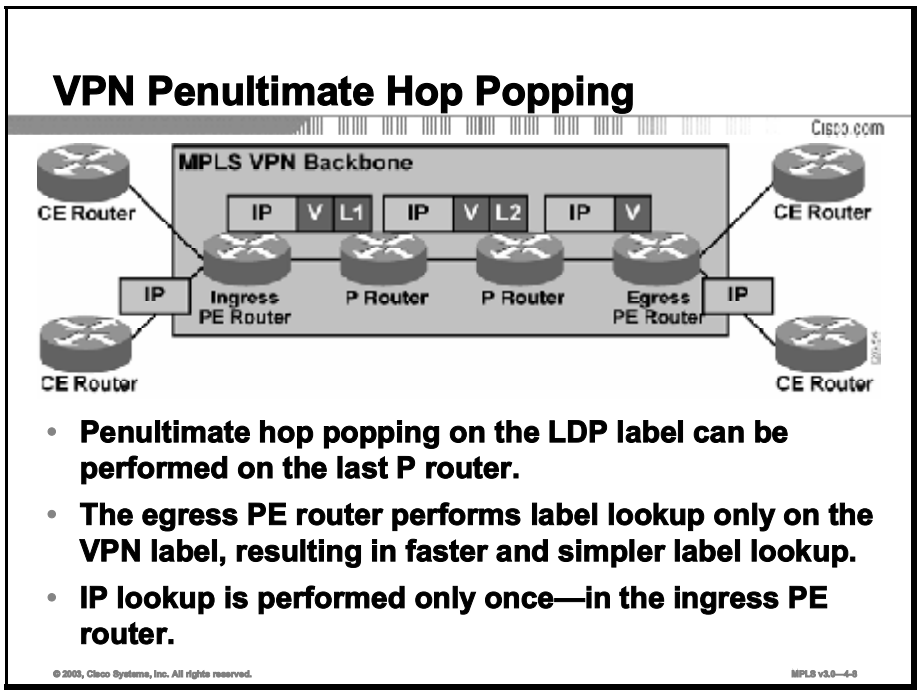

Penultimate hop popping, or PHP (the removal of the top label in the stack on the hop prior to the egress router), can be performed in frame-based MPLS networks. In these networks, the last P router in the label switched path (LSP) tunnel pops the LDP label (as previously requested by the egress PE router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

VPN Label Propagation

This topic describes how labels get propagated between PE routers.

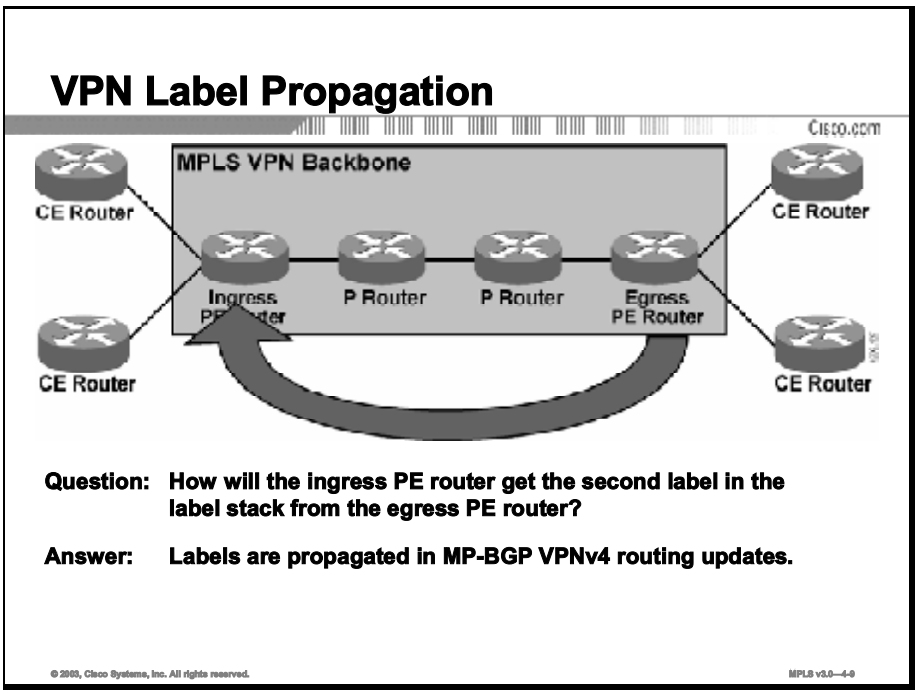

The previous figures showed that an MPLS label stack with the second label is required for proper MPLS VPN operation. This label was allocated by the egress PE router. This label has to be propagated from the egress PD router to the ingress PE routers to enable proper packet forwarding. MP-BGP was chosen 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.

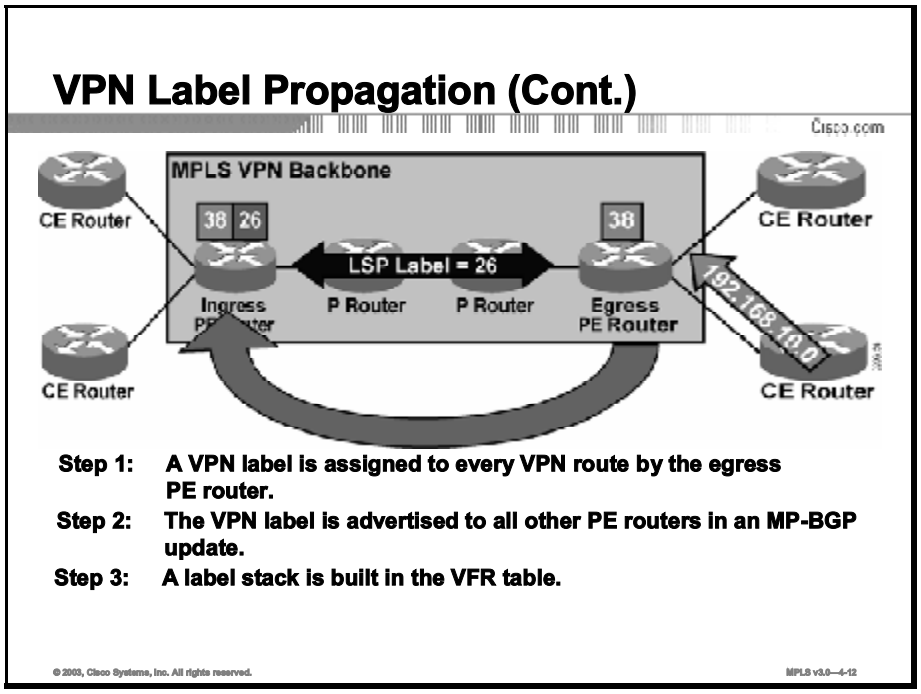

The figure illustrates VPN label propagation between PE routers.

Step 1 The egress PE router assigns a label to every VPN route received from 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 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 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, or LIB) as well as 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.

Effects of MPLS VPNs on Label Propagation

This topic describes the effects of MPLS VPNs on label propagation.

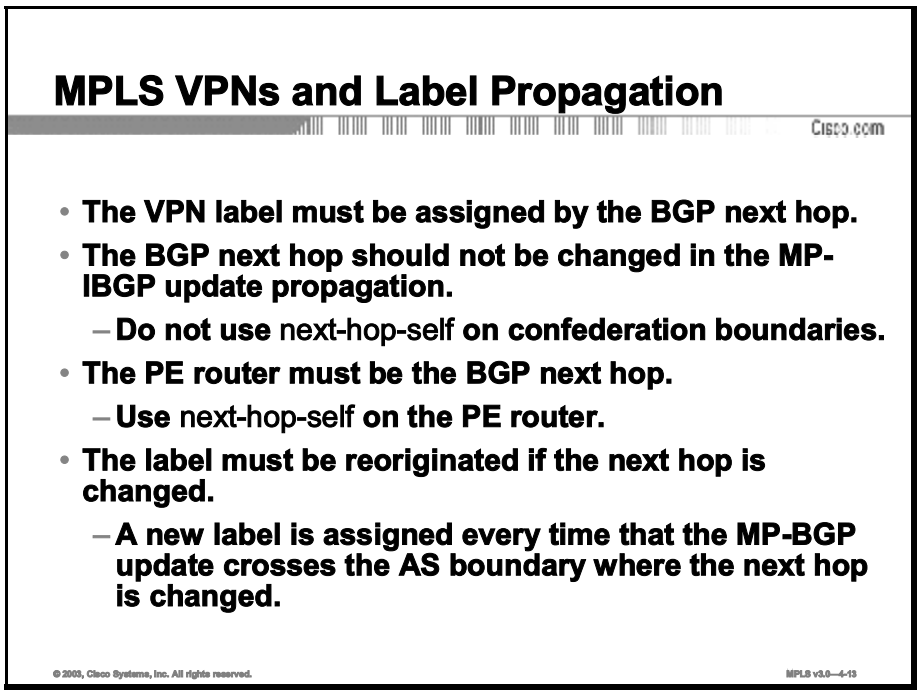

MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same router as the one that assigned the second label in the label stack. Three scenarios can cause the BGP next hop to be different from the IP address of the PE router assigning 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 **next-hop-self** 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 **next-hop-self** on an inter-AS boundary inside a BGP confederation or if you use inbound **route-map** on a PE route to change the next hop (a strongly discouraged practice). To prevent this situation, never change the BGP next hop with **route-map** or **next-hop-self** 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 the BGP next hop is changed. This functionality is supported by Cisco IOS Releases 12.1(4) T, 12.2, and later.

Effects of MPLS VPNs on Packet Forwarding

This topic describes the effects of MPLS VPNs on packet forwarding.

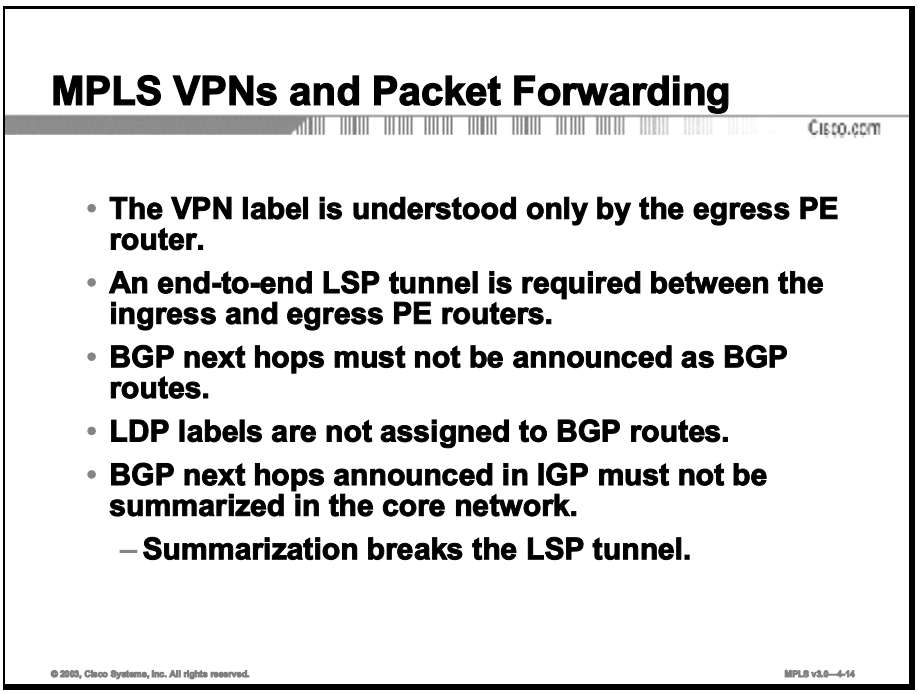

For successful propagation of MPLS VPN packets across an MPLS backbone, there must be an unbroken LSP tunnel between PE routers. This is because the second label in the stack is recognized only by the egress PE router that has originated it and will not be understood by any other router should it ever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

Two scenarios 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 next figure.

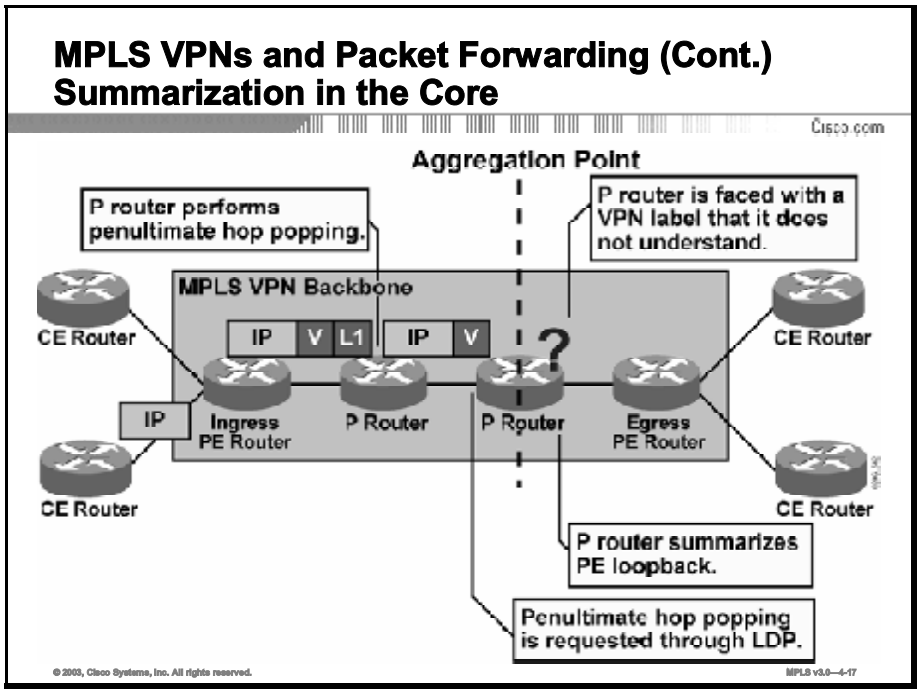

In the figure, the P router summarizes the loopback address of the egress PE router. The LSP tunnel is broken at a summarization point, so the summarizing router needs to perform full IP lookup. In a frame-based MPLS network, the P router would request PHP for the summary route, and the upstream P router (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.

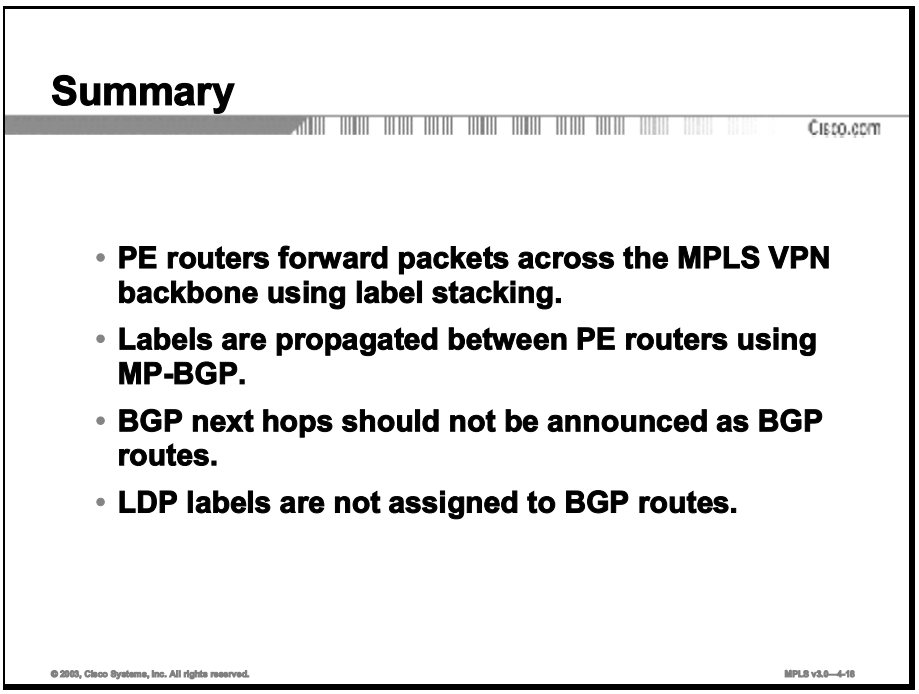

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Why does the label stack contain two labels when supporting MPLS VPNs?

Q2) Why is the VPN label not popped during the PHP process?

Q3) Which protocol is used to transport VPN labels between PE routers?

A) LDP

B) RSVP

C) MP-BGP

D) the core IGP

Q4) In MPLS VPNs, why must the BGP next hop be set to the egress router in all MP-IBGP updates?

Q5) What scenarios would cause the LSP tunnel between PE routers to break?

Quiz Answer Key

Q1) 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.

Relates to: VPN Packet Forwarding Across an MPLS VPN Backbone

Q2) The egress router needs the label to identify which VPN the packet belongs to.

Relates to: VPN Penultimate Hop Popping

Q3) C

Relates to: VPN Label Propagation

Q4) 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.

Relates to: Effects of MPLS VPNs on Label Propagation

Q5) 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.

Relates to: Effects of MPLS VPNs on Packet Forwarding

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers and solutions are found in the Module Assessment Answer Key.

Quiz: MPLS Virtual Private Networks Technology

Complete the quiz to assess what you have learned in this module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Identify the major VPN topologies, their characteristics, and usage scenarios
- Describe the differences between an overlay VPN and peer-to-peer VPN, identifying the implementation, benefits, and drawbacks
- Describe the major VPN topology categories and their implementation
- Describe the major architectural blocks of MPLS VPNs, identifying the functions of route information propagation, route distinguisher, route target, and virtual routing tables
- Identify the routing requirements for MPLS VPNs by describing, from the customer and provider perspective, how routing tables appear on provider edge routers, how customer routes are propagated, and how end-to-end information flows
- Describe how packets are forwarded in an MPLS VPN environment, identifying how VPN labels get propagated, and explaining label imposition and the effect of summarization in the core

Instructions

Complete these steps:

Step 1 Answer all questions in this quiz by selecting the best answer(s) to each question.

Step 2 Verify your results against the answer key at the end of this quiz.

Step 3 Review the related lesson for each question that you answered incorrectly.

Quiz

Answer these questions:

Q1) Which of the following is a characteristic of an overlay VPN?

A) PE routers carry all routes from all customers.

B) It 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.

Q2) Which connectivity category should you use if all sites must have connectivity with each other?

A) simple

B) overlapping

C) peer-to-peer

D) hub-and-spoke

E) central services

Q3) Which connectivity category should you use if all sites must have connectivity to a server provided by the service provider?

A) simple

B) overlapping

C) peer-to-peer

D) hub-and-spoke

E) central services

Q4) What are the connectivity requirements of a managed network VPN?

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.

Q5) In what two ways do MPLS VPNs support overlapping customer address spaces? (Choose two.)

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

Q6) Which of the following is true if you use the P-network IPG to propagate customer routing information across the P-network?

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.

Q7) Why do MPLS VPNs implement route targets?

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

Q8) What is the impact of an MPLS VPN on CE routers?

A) They must support BGP.

B) They must run a link-state protocol.

C) They can run any standard IP routing protocol.

D) Their IGP must be upgraded to a VPN-aware IGP.

Q9) Why would IPv4 routing be enabled on the PE router?

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

Q10) Which two types of route would an MPLS VPN install into the VRF? (Choose two.)

A) those received via an IPv4 update

B) those received via a VPNv4 update

C) those received via the core IGP update

D) those received via the customer IGP update

Q11) What will happen if the SOO attached to the route is equal to the SOO attribute associated with the CE router?

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.

Q12) How can P routers forward VPN packets if they do not have VPN routes?

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.

Q13) Which router assigns the VPN label?

A) P

B) egress CE

C) egress PE

D) ingress CE

E) ingress PE

Q14) What is used to identify the label that will be used to transport the VPN packet to the egress router?

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

Q15) What is the impact of changing a BGP next hop on an MP-BGP update at confederation boundaries?

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.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) D

Relates to: Introduction to Virtual Private Networks; Overlay and Peer-to-Peer VPNs

Q2) A

Relates to: Overlay and Peer-to-Peer VPNs; VPN Categorization

Q3) E

Relates to: Overlay and Peer-to-Peer VPNs; VPN Categorization

Q4) C

Relates to: Overlay and Peer-to-Peer VPNs; VPN Categorization

Q5) A, D

Relates to: MPLS VPN Architecture

Q6) D

Relates to: MPLS VPN Architecture

Q7) B

Relates to: MPLS VPN Architecture

Q8) C

Relates to: MPLS VPN Routing Model

Q9) C

Relates to: MPLS VPN Routing Model

Q10) B, D

Relates to: MPLS VPN Routing Model

Q11) C

Relates to: MPLS VPN Routing Model

Q12) A

Relates to: MPLS VPN Packet Forwarding

Q13) C

Relates to: MPLS VPN Packet Forwarding

Q14) C

Relates to: MPLS VPN Packet Forwarding

Q15) B

Relates to: MPLS VPN Packet Forwarding

Module 5

MPLS VPN Implementation

Overview

This module covers Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) implementation on Cisco IOS platforms. It describes the concepts of virtual routing and forwarding (VRF) tables, the interaction between customer-to-provider routing protocols, and Multiprotocol Border Gateway Protocol (MP-BGP) in the service provider backbone, as well as advanced MPLS VPN-specific routing protocol features. The module continues with a description of MPLS VPN monitoring and debugging commands available on Cisco IOS platforms and concludes with a troubleshooting lesson describing failure scenarios, identifying symptoms, and providing remedial action.

Module Objectives

Upon completing this module, you will be able to use a diagram of a typical simple MPLS VPN solution to identify the Cisco IOS command syntax required to successfully configure, monitor, and troubleshoot VPN operations. This ability includes being able to do the following:

- Describe the usage of VRF tables in an MPLS VPN environment
- Identify the command syntax required to configure VRF tables
- Identify the Cisco IOS command syntax required to successfully configure MP-BGP sessions between PE routers
- Identify the Cisco IOS command syntax required to successfully configure small-scale routing protocols (static, RIP, and EIGRP) between CE and PE routers
- Identify the Cisco IOS command syntax required to successfully configure BGP as the routing protocol between CE and PE routers
- Identify the Cisco IOS command syntax required to successfully configure OSPF as the routing protocol between CE and PE routers
- Identify the Cisco IOS command syntax required to monitor VPN operations
- Identify the Cisco IOS command syntax required to successfully troubleshoot VPN operations

Module Outline

The module contains these lessons:

- MPLS VPN Mechanisms of Cisco IOS Platforms
- Configuring VRF Tables
- Configuring an MP-BGP Session Between PE Routers
- Configuring Small-Scale Routing Protocols Between PE and CE Routers
- Monitoring MPLS VPN Operations
- Configuring OSPF as the Routing Protocol Between PE and CE Routers
- Configuring BGP as the Routing Protocol Between PE and CE Routers
- Troubleshooting MPLS VPNs

MPLS VPN Mechanisms of Cisco IOS Platforms

Overview

This lesson first introduces the VRF table, the major data structure associated with MPLS VPN implementation on Cisco IOS platforms. It describes the other MPLS VPN attributes that are associated with a virtual routing and forwarding instance (VRF), and explains the need for routing protocol contexts and the interaction of routing protocol contexts, VRFs, and MP-BGP.

Relevance

Having a clear understanding of how information is exchanged using VRFs and routing protocol contexts will make it easier to configure VRFs in your network.

Objectives

This lesson describes the usage of VRF tables in an MPLS VPN environment.

Upon completing this lesson, you will be able to:

- Describe the concept of and need for VRF tables in MPLS VPN implementations
- Describe the concept and use of routing protocol contexts in MPLS VPN implementations
- Identify the VRF-aware routing protocols
- Describe the implementation of VRFs in an MPLS VPN implementation
- Describe the outbound interaction between PE-CE routing protocols, backbone MP-BGP, and VRF tables
- Describe the inbound interaction between PE-CE routing protocols, backbone MP-BGP, and VRF tables

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components, and an understanding of basic VPN implementation and routing protocols

Outline

This lesson includes these topics:

- Overview
- Virtual Routing and Forwarding Table
- Need for Routing Protocol Contexts
- VPN-Aware Routing Protocols
- Contents and Use of the VRF Table
- BGP Route Propagation—Outbound
- Non-BGP Route Propagation—Outbound
- Route Propagation—Inbound
- Summary
- Quiz

Virtual Routing and Forwarding Table

This topic describes what a VRF table is and what it does in an MPLS VPN.

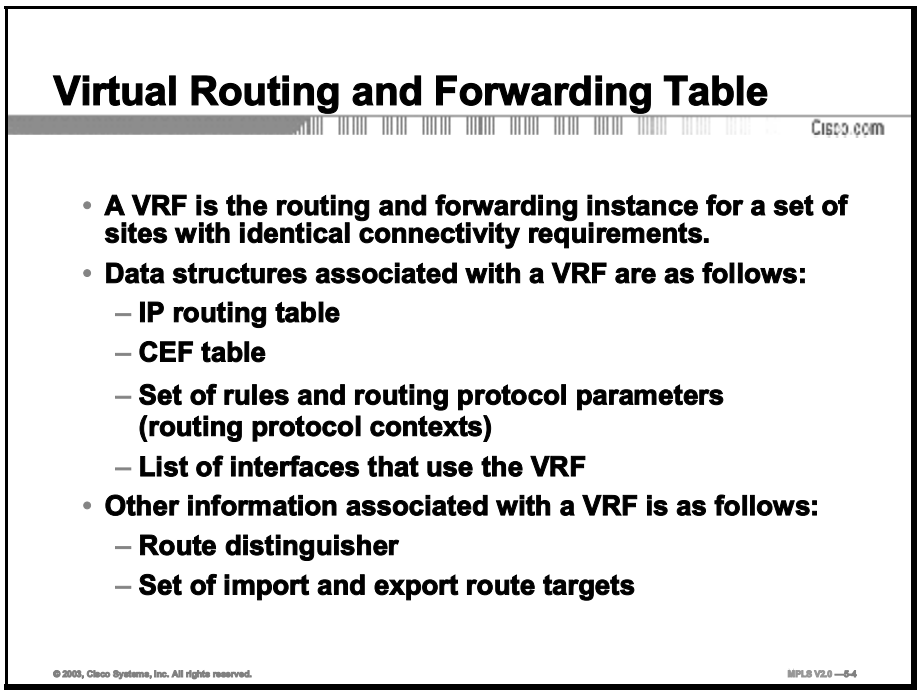

The major data structure associated with MPLS VPN implementation on Cisco IOS platforms is the virtual routing and forwarding instance (VRF) table. This data structure encompasses an IP routing table identical in function to the following:

- The global IP routing table in Cisco IOS software
- A Cisco Express Forwarding (CEF) table identical in function to the global CEF forwarding table (Forwarding Information Base, or FIB)
- Specifications for routing protocols running inside the VRF

A VRF is thus a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same provider edge (PE) router *as long as these sites share exactly the same connectivity requirements*.

Other MPLS VPN attributes associated with a VRF are as follows:

- The route distinguisher (RD), which is prepended (for example, RD + IP Address) to all routes exported from the VRF into the global VPNv4 (also called VPN IPv4) Border Gateway Protocol (BGP) table
- A set of export route targets (RTs), which are attached to any route exported from the VRF
- A set of import route targets, which are used to select VPNv4 routes that are to be imported into the VRF

Need for Routing Protocol Contexts

This topic describes the concept of routing protocol contexts.

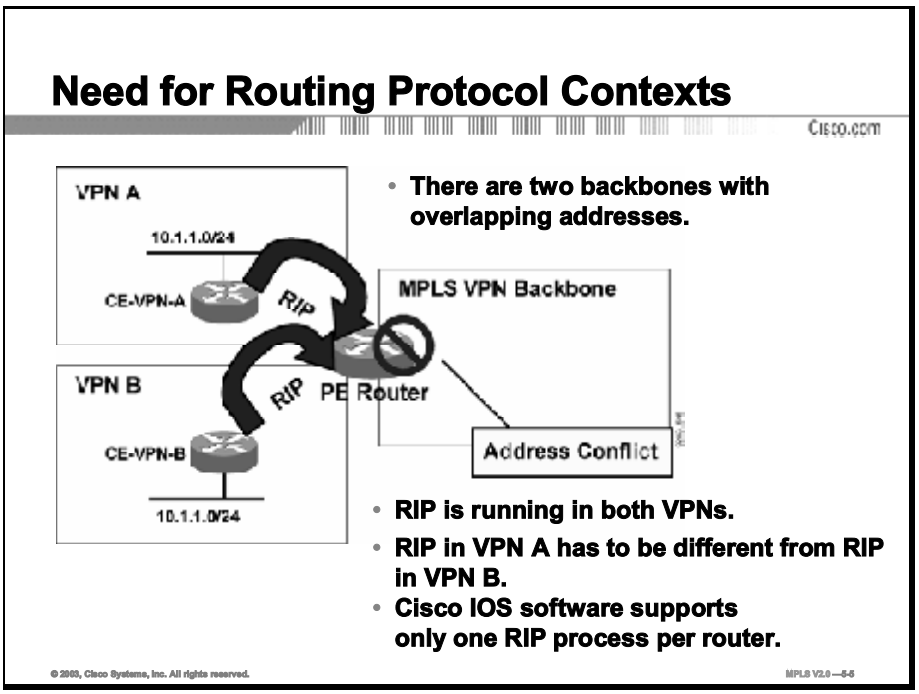

Traditional Cisco IOS software can support a number of different routing protocols. In some cases even several completely isolated copies of the same routing protocol are supported. For example, several Open Shortest Path First (OSPF) processes can be used.

It is important to understand that for several important routing protocols, such as Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), or BGP, Cisco IOS software supports only a single copy of the protocol running in the router. These protocols cannot be used directly between PE and customer edge (CE) routers in VPN environments because each VPN (or, more precisely, each VRF) needs a separate, isolated copy of the routing protocol to prevent undesired route leakage between VPNs. Furthermore, VPNs can use overlapping IP address spaces (for example, each VPN could use subnets of network 10.0.0.0), which would also lead to routing confusions if all VPNs shared the same copy of the routing protocol.

VPN-Aware Routing Protocols

This topic identifies the VPN-aware routing protocols.

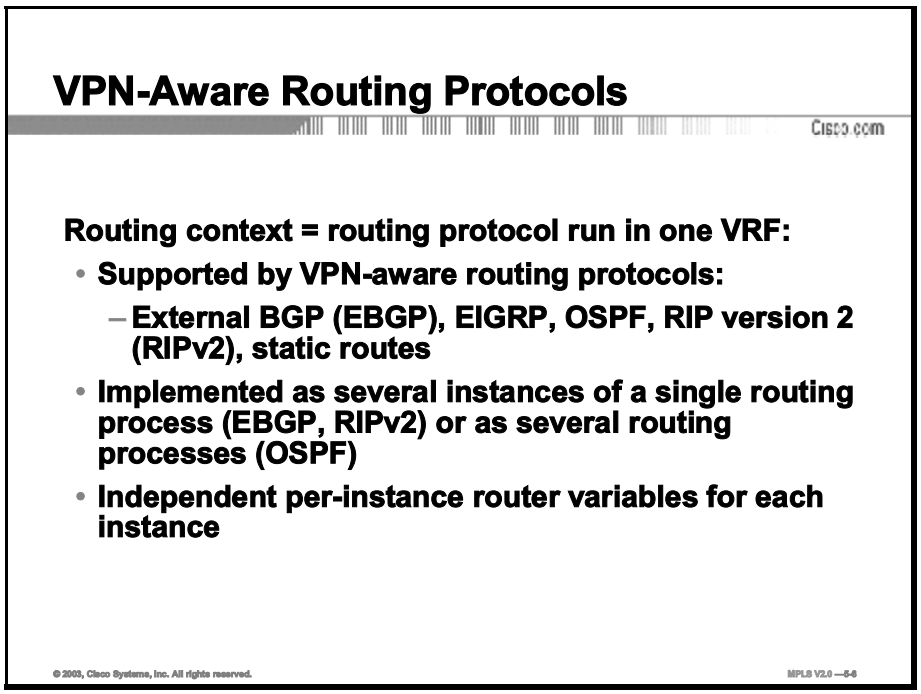

"Routing contexts" were introduced in Cisco IOS software to support the need for separate isolated copies of VPN routing protocols. They can be implemented as separate routing processes (OSPF), similar to traditional Cisco IOS software implementation, or as separate isolated "instances" of the same routing protocol.

If the routing contexts are implemented as instances of the same routing protocol, each instance contains its own independent routing protocol parameters. Examples would include networks over which the routing protocol is run, timers, authentication parameters, passive interfaces, and neighbors. This independence allows the network designer maximum flexibility in implementing routing protocols between PE and CE routers.

Contents and Use of the VRF Table

This topic describes the contents and use of VRF tables.

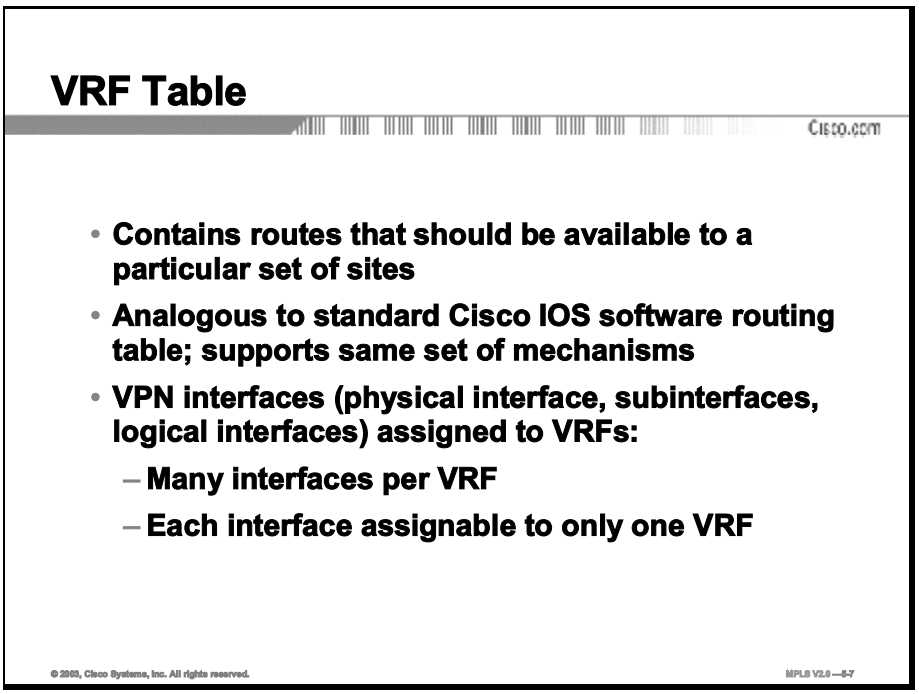

The routes received from VRF routing protocol instances or from dedicated VRF routing processes are inserted into the IP routing table contained within the VRF. This IP routing table supports exactly the same set of mechanisms as the standard Cisco IOS software routing table. These mechanisms include filter mechanisms (distribute lists or prefix lists) and interprotocol route selection mechanisms (administrative distances).

The per-VRF forwarding table (FIB) is built from the per-VRF routing table. This table is used to forward all the packets received through the interfaces associated with the VRF. Any interface can be associated with a VRF, be it a physical interface, subinterface, or a logical interface, as long as it supports CEF switching.

Note The requirement to support CEF switching on inbound VRF interfaces prevents certain media or encapsulation types from being used for VPN connectivity. More notable examples in mainstream Cisco IOS Release 12.1 include dialer interfaces, ISDN interfaces, and Switched Multimegabit Data Service (SMDS) interfaces. Some restrictions are already lifted in Cisco IOS Release 12.1 T. Refer to the release notes of the Cisco IOS platform that you are using for details of interfaces and media types supporting CEF switching.

There is no limit to the number of interfaces associated with one VRF (other than the number of interfaces supported by the router). However, in practice, each interface can be assigned to only one VRF because the router needs to uniquely identify the forwarding table to be used for packets received over an interface.

BGP Route Propagation—Outbound

This topic describes the outbound BGP route propagation process in an MPLS VPN implementation.

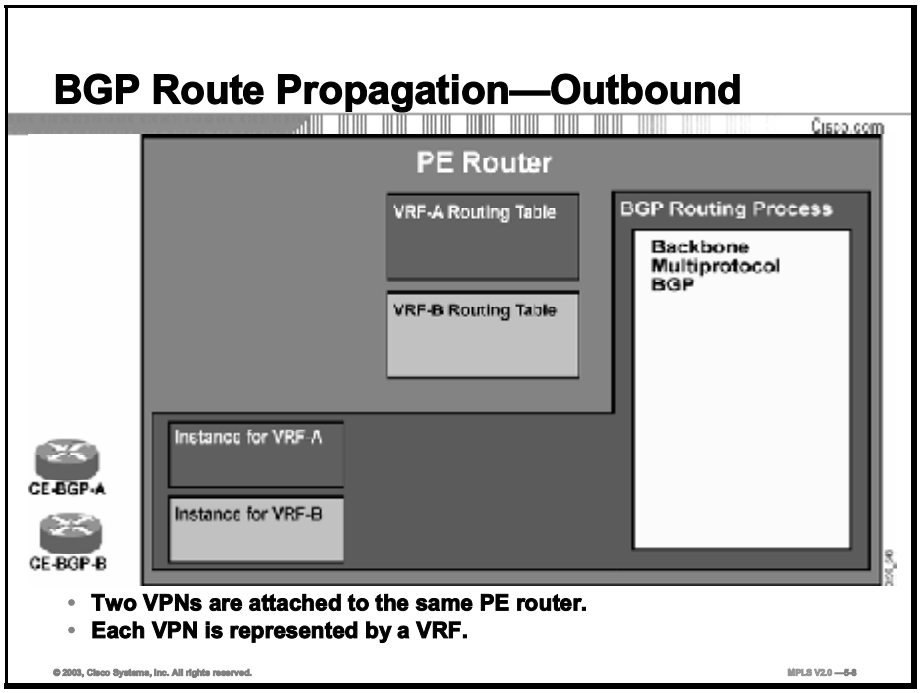

This figure and the following figures illustrate the interactions between VRF instances of routing processes, VRF routing tables, and the global VPNv4 BGP routing process.

The network contains two VPN customers. Ordinarily, the customer sites would be connected to a number of PE routers. This example focuses only on a single PE router, which contains two VRFs—one for each customer. Each customer is connected to the PE router, which is running BGP. CE-BGPA is the CE router for Customer A and is associated with VRF-A (VPN-A). CE-BGPB is the CE router for Customer B and is associated with VRF-B (VPN-B).

BGP Route Propagation—Outbound (Cont.)

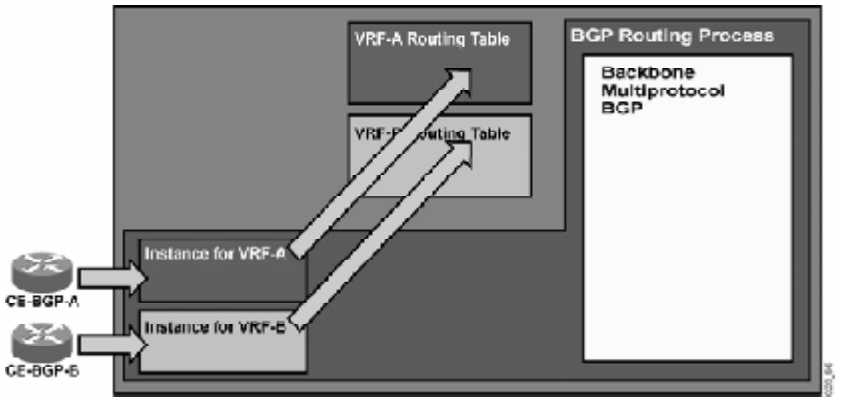

- **BGP-speaking CE routers announce their prefixes to the PE router via BGP.**
- **Instance of BGP process associated with the VRF to which the PE-CE interface belongs collects the routes and inserts them into VRF routing table.**

The BGP-speaking CE routers announce their networks via EBGP sessions to the PE router. The CE BGP neighbors of the PE router are associated with individual VRFs that enable the various instances of the BGP routing process to put the received routing updates into the proper per-VRF routing table.

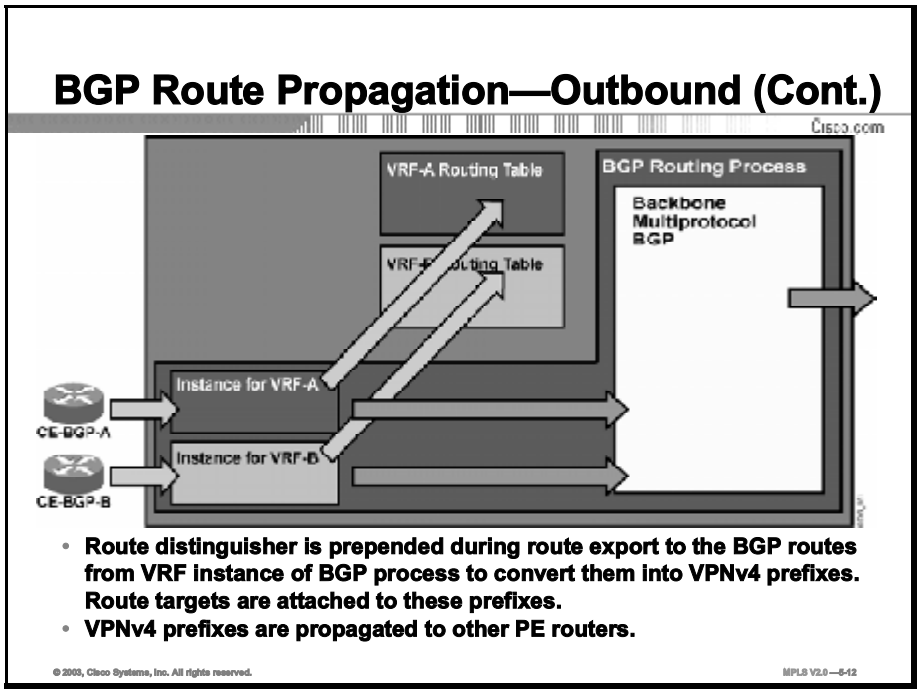

The routes illustrated here are being copied into the Multiprotocol Border Gateway Protocol (MP-BGP) table for further propagation to other PE routers.

The IP prefixes are prepended with the RD, and the set of RTs (extended BGP communities) configured as *export RTs* for the VRF is attached to the resulting VPNv4 route.

Note	The difference between the per-VRF BGP table and the global MP-BGP table holding VPNv4 routes is displayed only to illustrate the steps in the route propagation process. In reality, there is no separate per-VRF BGP table in Cisco IOS software.

Non-BGP Route Propagation—Outbound

This topic describes the outbound non-BGP route propagation process in an MPLS VPN implementation.

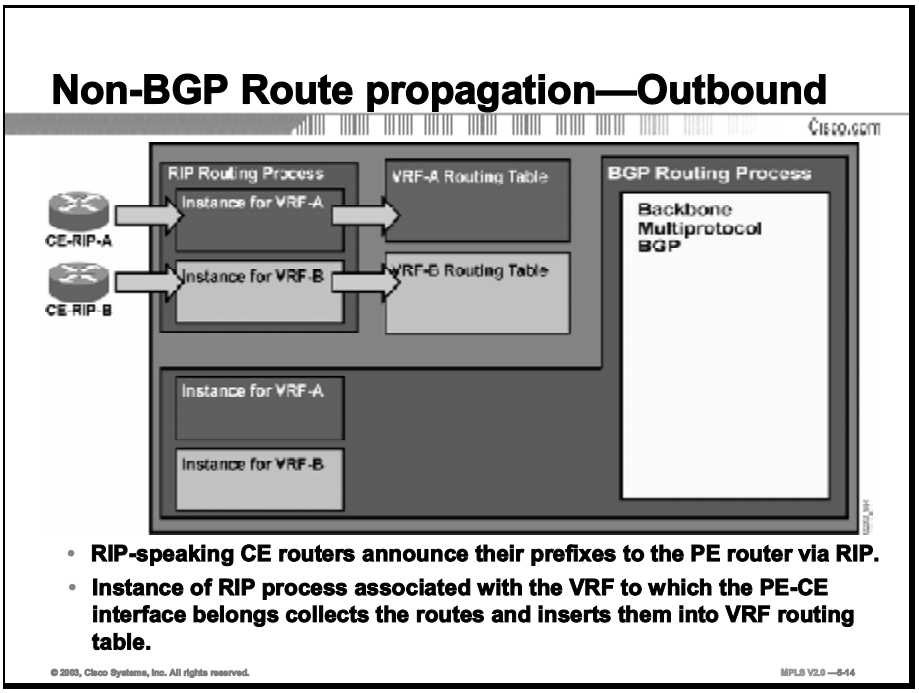

RIP-speaking CE routers identify the correct instance of RIP on the PE router when an inbound PE interface is associated with a VRF. This association allows CE routers to announce their networks to the appropriate per-VRF routing table.

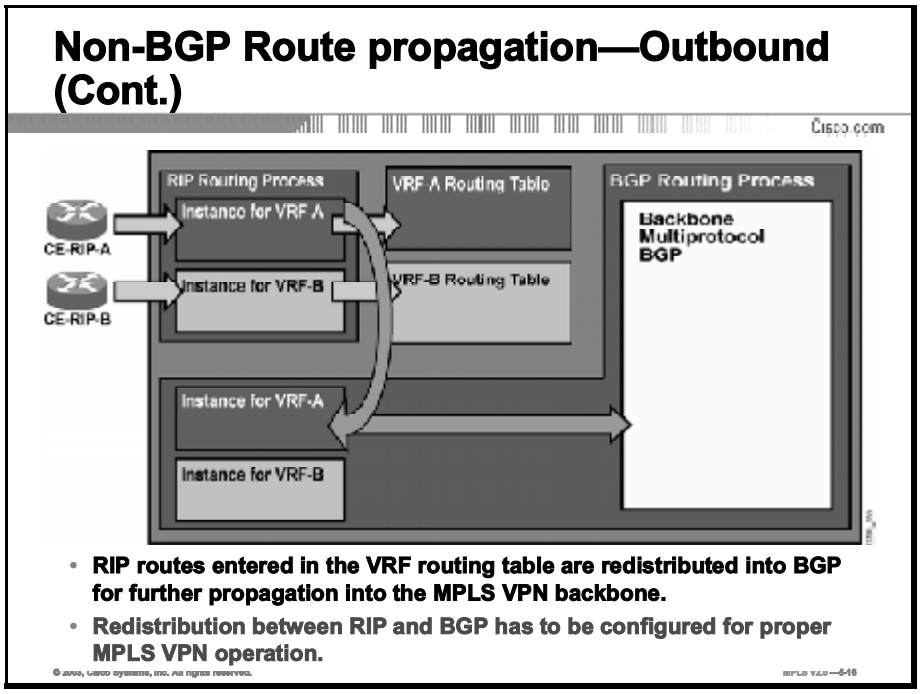

MP-BGP is used in the MPLS VPN backbone to carry VPN routes (prefixed with the RD) as 96-bit VPNv4 routes between the PE routers. The backbone BGP process looks exactly like a standard Internal Border Gateway Protocol (IBGP) setup from the perspective of the VRF. The per-VRF RIP routes therefore *must be redistributed* into the per-VRF instance of the BGP process to allow them to be propagated through the backbone MP-BGP process to other PE routers.

Caution	Failure to redistribute non-BGP routes into the per-VRF instance of BGP is one of the most common MPLS VPN configuration failures.

Should there be an overlap between an inbound RIP update and an inbound EBGP update, the standard route selection mechanism (administrative distance) is used in the per-VRF IP routing table and the EBGP route takes precedence over the RIP route. EBGP precedence results from the fact that the administrative distance of EBGP routes (20) is better than the administrative distance of RIP routes (120).

Route Propagation—Inbound

This topic describes the inbound route propagation process in an MPLS VPN implementation.

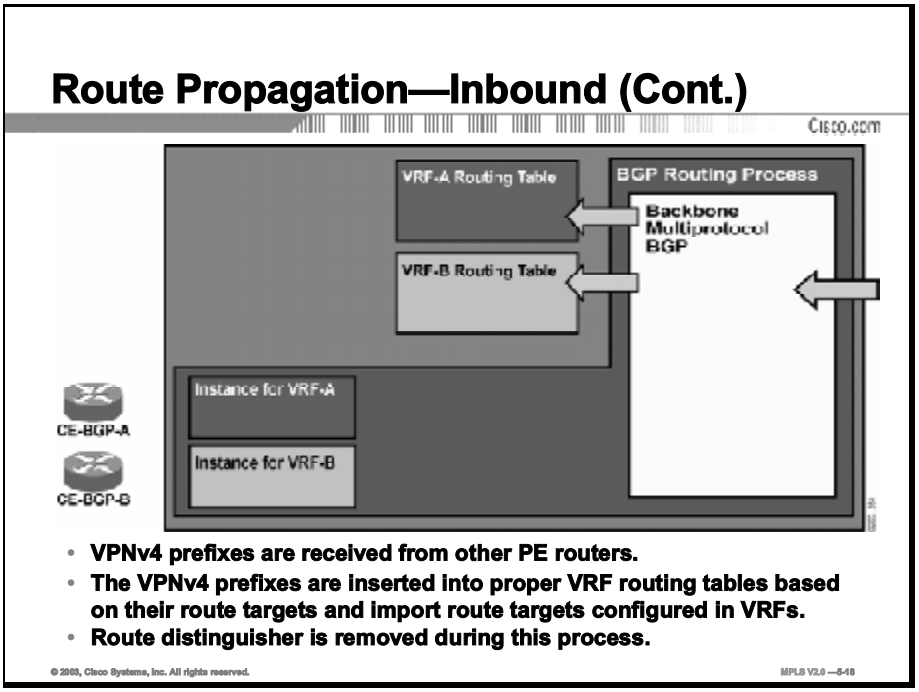

As other PE routers start originating VPNv4 routes, the MP-BGP process in the PE router here receives the routes. The routes are filtered based on RT attributes attached to them, and are inserted into the proper per-VRF IP routing tables based on the *import RTs* configured for individual VRFs. The RD that was prepended by the originating PE router is removed before the route is inserted into the per-VRF IP routing table.

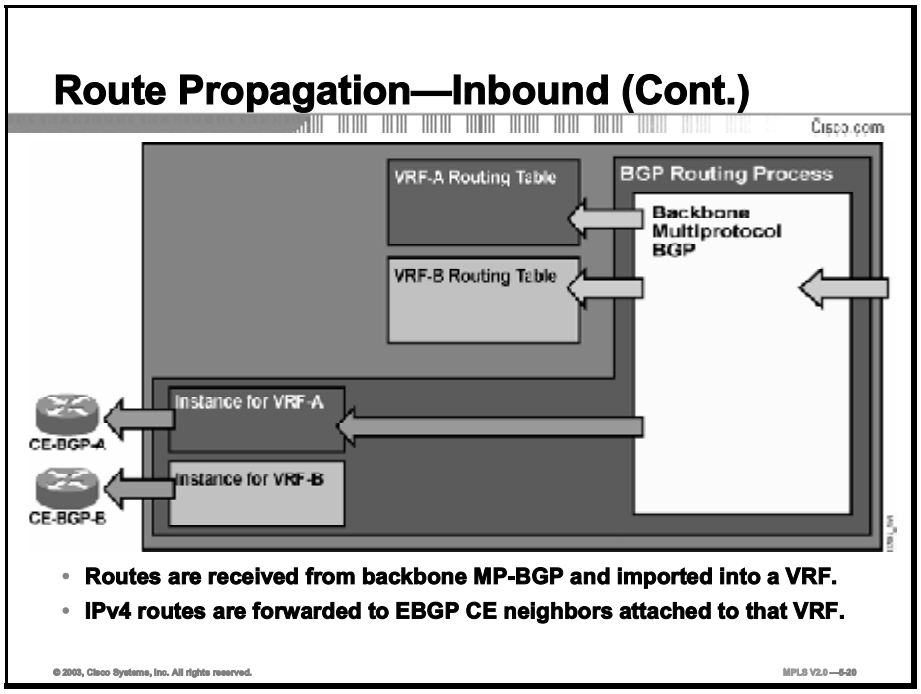

The Multiprotocol Internal Border Gateway Protocol (MP-IBGP) VPNv4 routes received from other PE routers and selected by the import RTs of a VRF are automatically propagated as 32-bit IPv4 (IP version 4) routes to all BGP-speaking CE neighbors of the PE router.

Route Propagation—Inbound (Cont.)

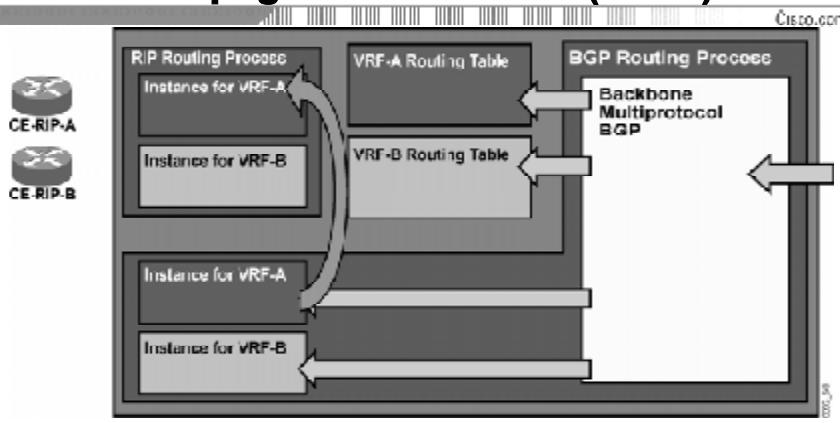

- **MP-IBGP routes imported into a VRF are redistributed into the instance of RIP configured for that VRF.**
- **Redistribution between BGP and RIP has to be configured for end-to-end RIP routing between CE routers.**

The MP-IBGP routes, although they are inserted in the per-VRF IP routing table, are *not* propagated to RIP-speaking CE routers automatically. To propagate these MP-IBGP routes to the RIP-speaking CE routers, you must manually configure the redistribution between per-VRF instance of BGP and per-VRF instance of RIP.

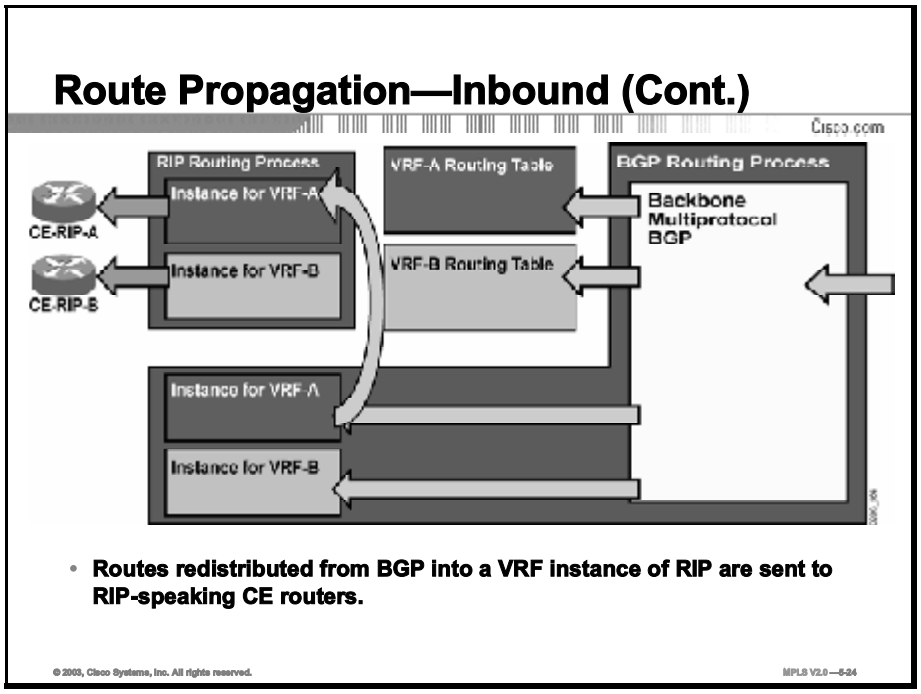

When the IBGP routes from the per-VRF IP routing table are successfully redistributed into the per-VRF instance of the RIP process, the RIP process announces these routes to RIP-speaking CE routers, thus achieving transparent end-to-end connectivity between the CE routers.

Summary

This topic summarizes the key points discussed in this lesson.

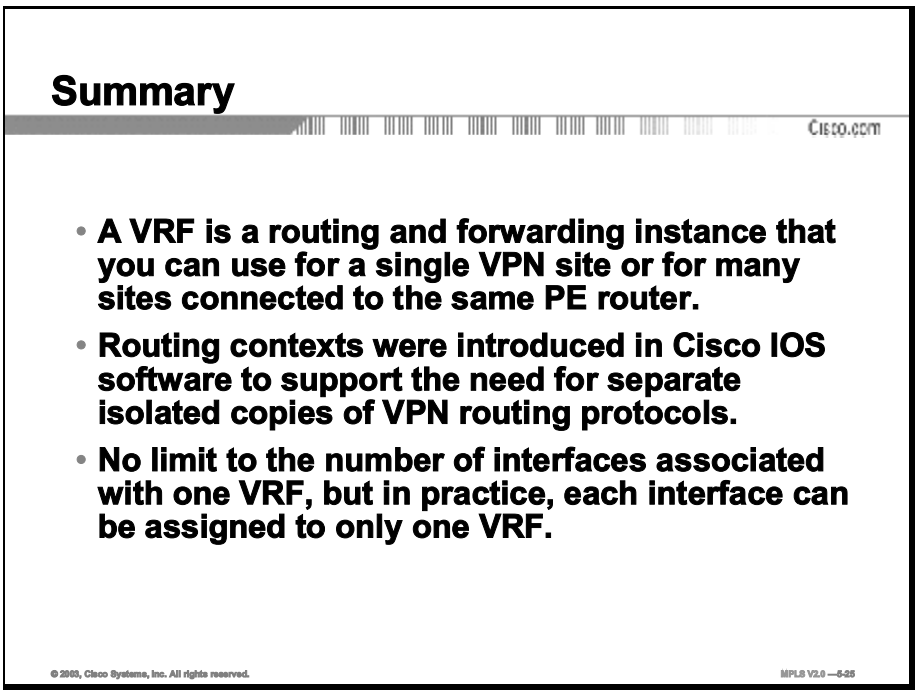

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) In an MPLS VPN implementation, what is a VRF?

A) the routing and forwarding instance for all sites belonging to a single customer

B) the routing and forwarding instance for all sites belonging to a single customer location

C) the routing and forwarding instance for all sites using a common routing protocol

D) the routing and forwarding instance for a set of sites with identical connectivity requirements

Q2) Why are VRFs used to establish separate routing protocol contexts?

_____

_____

Q3) Which two protocols are VPN-aware? (Choose two.)

A) RIPv2

B) IS-IS

C) OSPF

D) EIGRP

Q4) True or False? VRFs are assigned to an interface.

_____

Q5) A PE router is supporting site A for a VPN on one interface using RIP as the routing protocol. Site B belongs to the same VPN and is being supported on a second interface using EBGP as the routing protocol. Why is it necessary to redistribute the RIP-learned route into the per-VRF instance of the BGP process?

A) to allow site A and B to communicate with each other

B) to allow the RIP route to be propagated to the VRF routing tables

C) to allow the RIP routes to be propagated to the local EBGP session

D) to allow the RIP routes to be propagated through the backbone MP-BGP process to other PE routers

Q6) How are VPNv4 routers propagated to a RIP-speaking CE router?

_____

_____

_____

Quiz Answer Key

Q1) D

Relates to: Virtual Routing and Forwarding Table

Q2) because with routing protocols like RIP and BGP, only a single copy of the protocol may be running in the router

Relates to: Need for Routing Protocol Contexts

Q3) A, D

Relates to: VPN-Aware Routing Protocols

Q4) False. Interfaces are assigned to a VRF.

Relates to: Contents and Use of the VRF Table

Q5) D

Relates to: BGP Route Propagation—Outbound

Q6) They are redistributed from the global BGP table to the per-instance BGP table and then to the per-instance RIP, which is propagated to the CE.

Relates to: Route Propagation—Inbound

Configuring VRF Tables

Overview

This lesson explains how to configure VRF tables, listing the configuration tasks, syntax, and definitions of commands used for creation of VRFs. The lesson also provides an example of a VPN configuration.

Relevance

It is important to know how to configure and apply a VRF table onto a routing interface. It is essential to understand the command syntax for the configurations that you want to deploy in your network. This lesson will provide you with the information that will enable you to succeed at such tasks.

Objectives

This lesson describes the command syntax that is required to configure VRF tables.

Upon completing this lesson, you will be able to:

- Identify the tasks that are required to configure a VRF table
- Identify the command syntax that is required to create a VRF table
- Identify the command syntax that is required to specify export and import RTs
- Identify the command syntax that is required to assign an interface to a VRF table
- Describe a typical Cisco IOS configuration that enables VRFs

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- VRF Configuration Tasks
- Creating VRF Tables and Assigning RDs
- Specifying Export and Import RTs
- Assigning an Interface to a VRF Table
- MPLS VPN Network Example
- Summary
- Quiz

VRF Configuration Tasks

This topic identifies the tasks required to configure VRF tables.

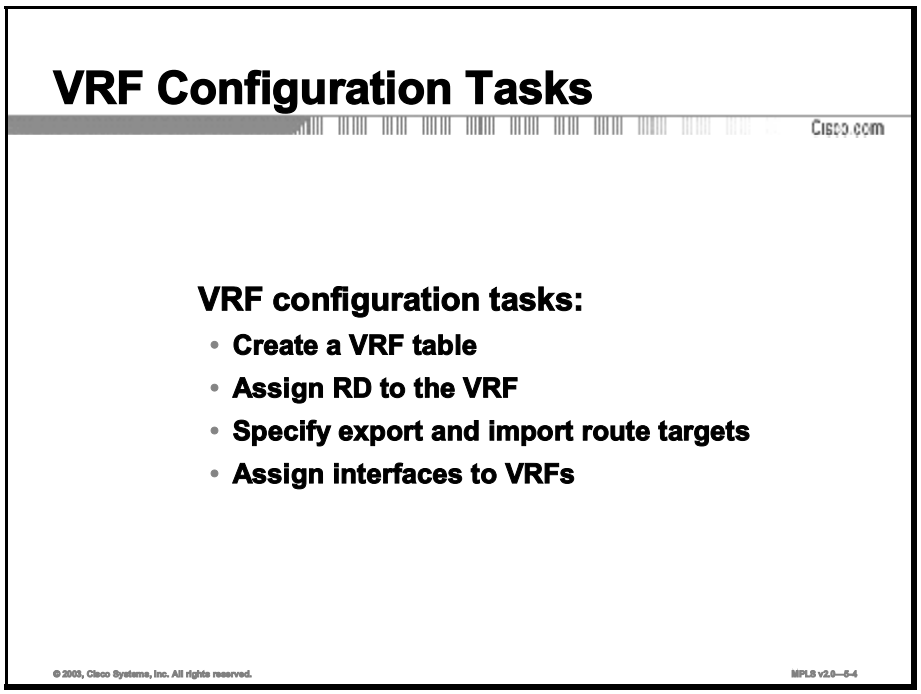

Configuring a VRF table and starting deployment of an MPLS VPN service for a customer consists of four mandatory steps:

- Create a new VRF table.
- Assign a unique RD to the VRF.

Note	You must assign a unique RD to every VRF created in a PE router. The same RD *might* be used in multiple PE routers, based on customer connectivity requirements. The same RD *should* be used on all PE routers for simple VPN service. Refer to the "MPLS VPN Architecture" module for more details on RD assignment for different VPN topologies.

- Specify import and export RTs for the VRF.

Note	Import and export RTs should be equal to the RD for simple VPN service.

- Assign interfaces to VRFs.

Creating VRF Tables and Assigning RDs

This topic identifies the command syntax required to create a VRF table.

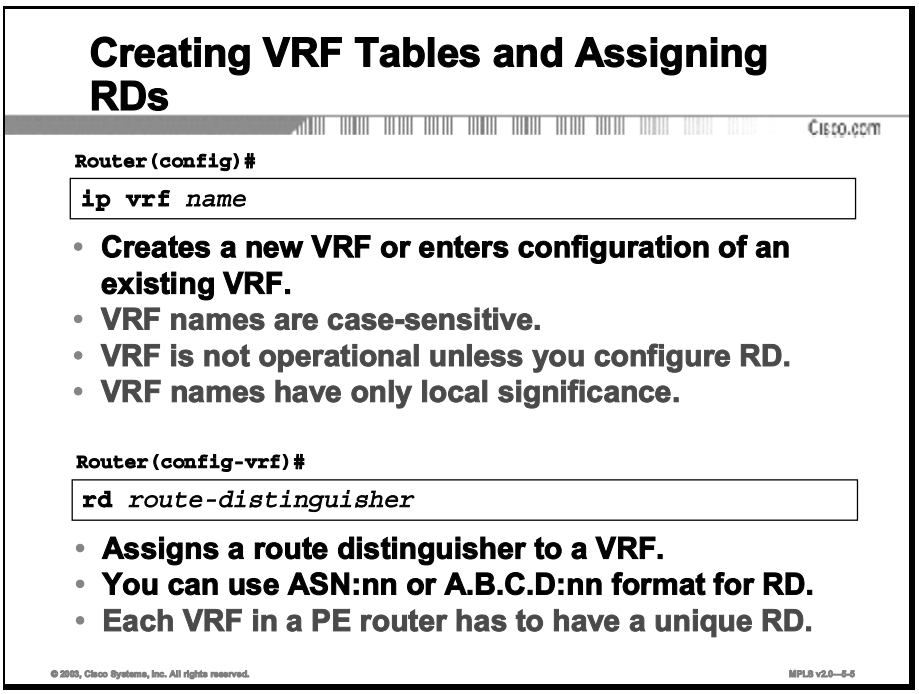

ip vrf

To configure a VRF routing table, use the **ip vrf** command in global configuration mode. To remove a VRF routing table, use the **no** form of this command:

- **ip vrf** *vrf-name*
- **no ip vrf** *vrf-name*

Syntax Description

Parameter	Description
vrf-name	Name assigned to a VRF.

Defaults

No VRFs are defined. No import or export lists are associated with a VRF. No route maps are associated with a VRF.

rd

To create routing and forwarding tables for a VRF, use the **rd** command in VRF configuration submode:

- **rd** *route-distinguisher*

Syntax Description

Parameter	Description
route-distinguisher	Adds an 8-byte value to an IPv4 prefix to create a VPNv4 prefix.

The RD can be specified in one of two formats:

- 16-bit autonomous system (AS) number followed by a 32-bit decimal number (ASN:nn)
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn)

Defaults

There is no default. An RD must be configured for a VRF table to be functional.

Note Once a VRF has been define using the **`ip vrf`** command and a RD has been assigned using the **`rd`** command, the VRF is operational. At this point, any locally active interface will appear in the vrf's routing display.

Specifying Export and Import RTs

This topic identifies the tasks required to configure VRF tables.

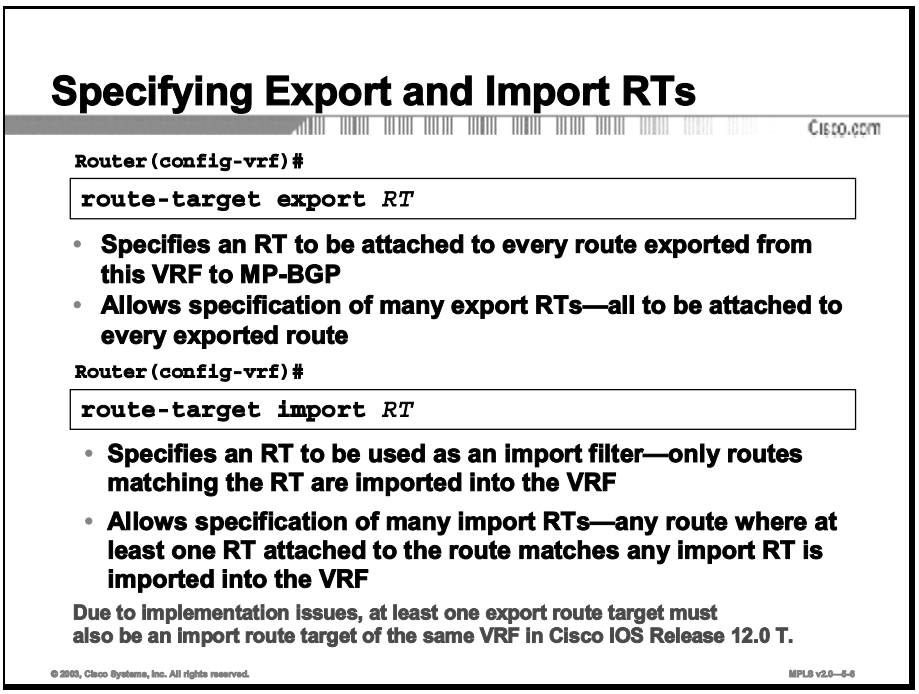

route-target

To create an RT extended community for a VRF, use the **route-target** command in VRF submode. To disable the configuration of an RT community option, use the **no** form of this command:

- **route-target** {**import** | **export** | **both**} *route-target-ext-community*
- **no route-target** {**import** | **export** | **both**} *route-target-ext-community*

Syntax Description

Parameter	Description
import	VPNv4 routes that contain an extended community value that matches the route-target-ext-community field that will be imported into the VRF
export	The value in the route-target-ext-community field that will be inserted into the extended community for routes exported from the VRF to VPNv4
both	Sets the value used by both the import and export process to the valued indicated in the route-target-ext-community field
route-target-ext-community	The RT extended community attribute for the VRF

Similar to RDs, the RTs can be specified in one of two formats:

- 16-bit AS number followed by a 32-bit decimal number (ASN:nn)
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn)

Defaults

There are no defaults. A VRF has no RT extended community attributes associated with it until specified by the **route-target** command.

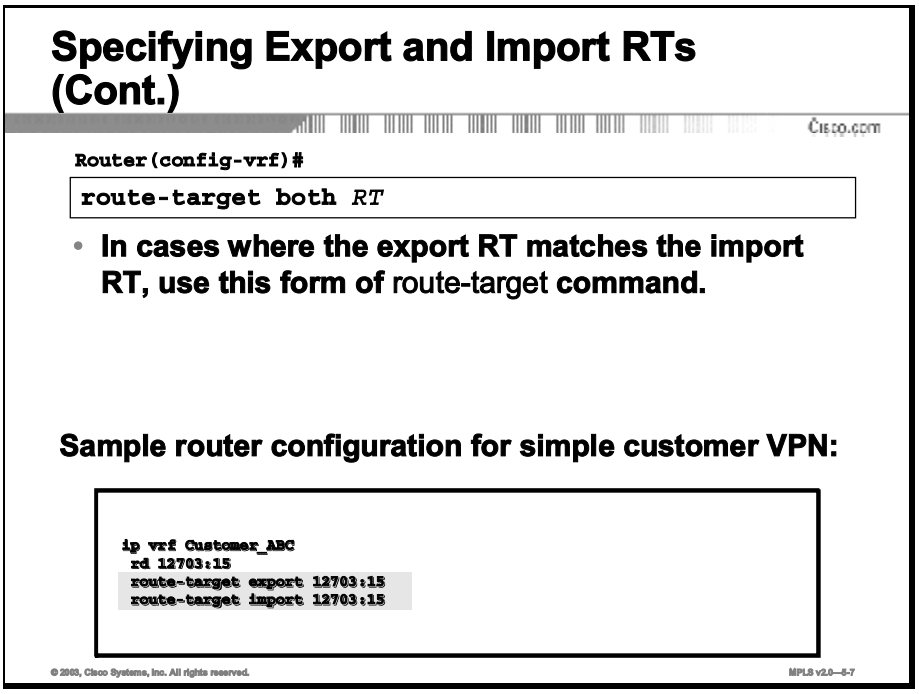

Whenever an RT is both an import and an export RT for a VRF, you can use the **route-target both** command to simplify the configuration. For example, the two **route-target** configuration lines in the sample router configuration in the figure could be reduced to a single command—**route-target both 12703:15**.

Assigning an Interface to a VRF Table

This topic identifies the command syntax required to assign an interface to a VRF table.

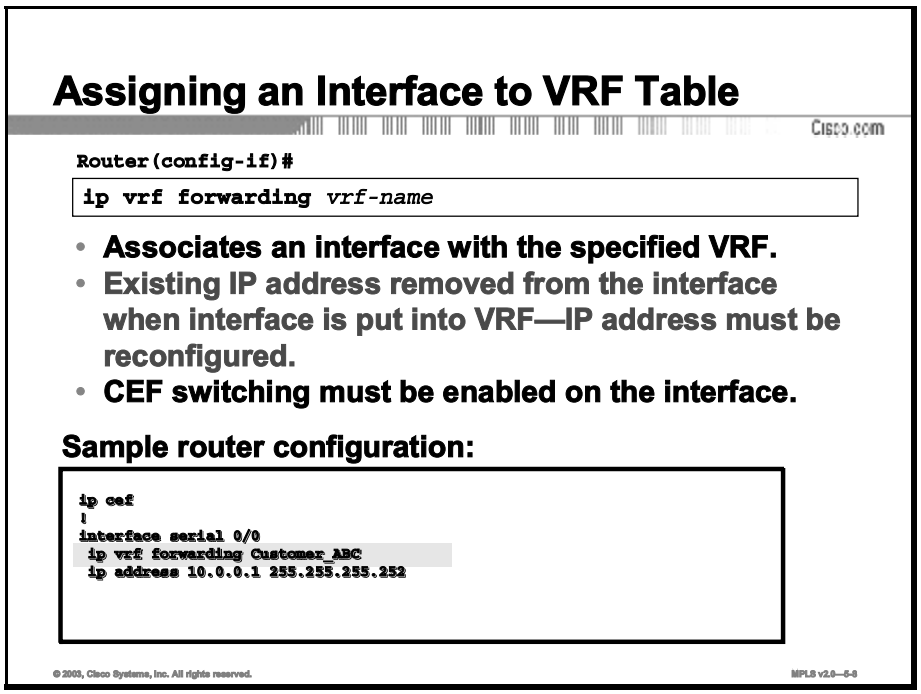

ip vrf forwarding

To associate a VRF with an interface or subinterface, use the **ip vrf forwarding** command in interface configuration mode. To disassociate a VRF, use the **no** form of this command:

- **ip vrf forwarding** *vrf-name*
- **no ip vrf forwarding** *vrf-name*

Syntax Description

Parameter	Description
vrf-name	Name assigned to a VRF.

Defaults

The default for an interface is the global routing table.

MPLS VPN Network Example

This topic provides a sample MPLS VPN network and then describes a typical Cisco IOS configuration that would be used to enable VRFs.

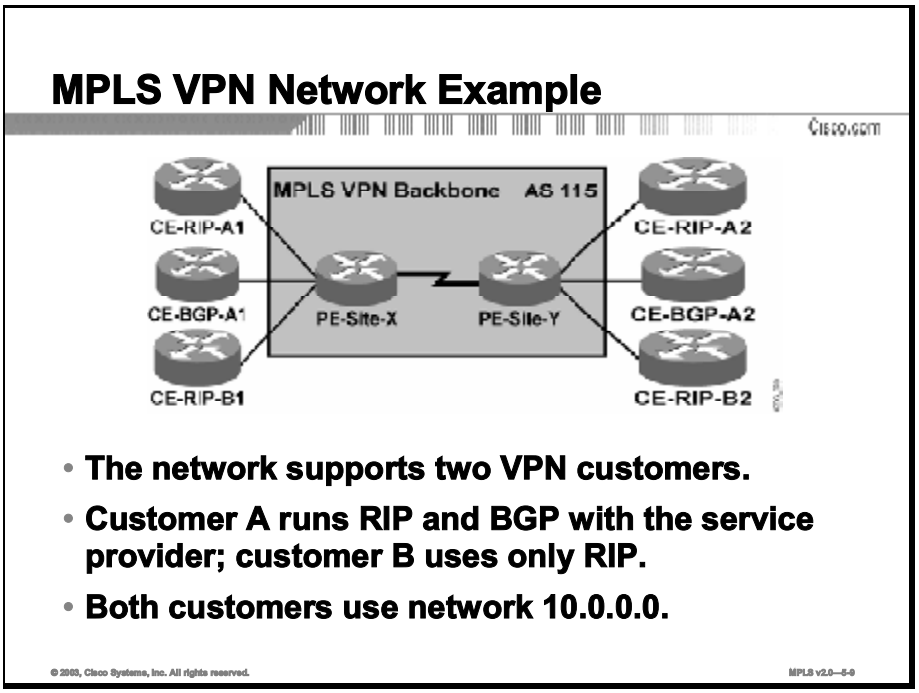

To illustrate the use of MPLS VPN configuration commands, here is a configuration of the PE router in a sample network with two VPN customers. Customer A (with four sites) is using BGP and RIP as the provider edge-customer edge (PE-CE) routing protocol, and customer B (with two sites) is using only RIP. Both customers use private IP address space (subnets of network 10.0.0.0).

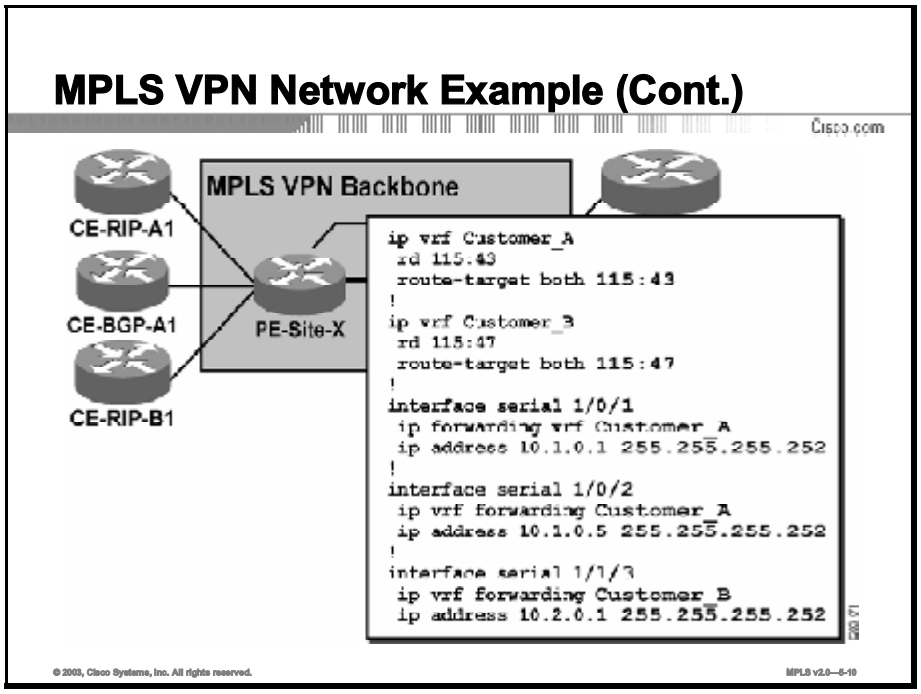

The configuration steps that you can perform on the PE router so far are as follows:

Step 1 Configure VRFs for customer A and customer B.

Step 2 Assign RDs and RTs to the VRFs. Only one RD per customer is used on all PE routers in the MPLS VPN backbone because these customers require only simple VPN connectivity. To simplify the configuration and troubleshooting process, the RTs are made equal to the RDs.

Step 3 Assign PE-CE interfaces to individual VRFs.

Summary

This topic summarizes the key points discussed in this lesson.

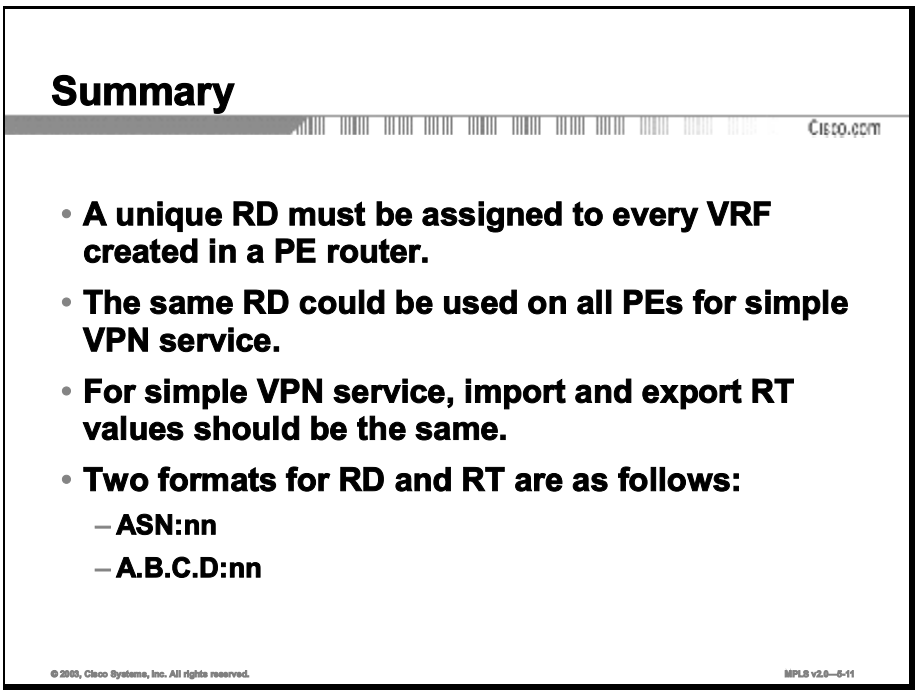

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which command do you use to create a VRF named VPNA?

A) **ip vrf VPNA**

B) **ip rt vrf VPNA**

C) **ip rd vrf VPNS**

D) **ip vrf forwarding VPNA**

Q2) Which two VRF parameters must be specified for a VRF to become operational? (Choose two.)

A) **rd route-distinguisher**

B) **route-target export RT**

C) **route-target import RT**

D) **ip vrf forwarding vrf-name**

Q3) Which command do you use to associate interface e0/0 with a VRF named VPNA?

A) **ip vrf VPNA**

B) **ip vrf VPNA int e0/0**

C) **ip vrf forwarding VPNA**

D) **ip vrf VPNA forwarding e0/0**

Q4) What happens to the interface of an existing IP address when you associate the interface with a VRF?

A) It will remain unchanged.

B) It will be removed from the interface.

C) It will be changed to the loopback 0 address.

D) It will be moved to under the VRF configuration.

Q5) You have created a configuration that defines three import route targets (650001:01, 650002:02, and 650003:03) for a VRF. A route update has three RTs (650003:03, 650004:04, and 650005:05) attached to it. How will this update be processed and why?

A) It will be accepted by the VRF because it matches the import RD of 03.

B) It will be discarded by the VRF because it does not match all of the RTs in the import list.

C) It will be accepted by the VRF because it matches at least one of the RTs in the import list.

D) It will be discarded by the VRF because it does not match all of the RDs in the import list.

Quiz Answer Key

Q1) A

Relates to: Creating VRF Tables and Assigning RDs

Q2) B, C

Relates to: Specifying Export and Import RTs

Q3) C

Relates to: Assigning an Interface to a VRF Table

Q4) B

Relates to: Assigning an Interface to a VRF Table

Q5) C

Relates to: Specifying Export and Import RTs

Configuring an MP-BGP Session Between PE Routers

Overview

This lesson explains the BGP process in an MPLS VPN-enabled router, listing the configuration tasks, steps, syntax, and descriptions. It also discusses BGP community propagation and provides an MP-IBGP configuration example.

Relevance

Most of the configuration in an MPLS VPN depends on how the PE routers are configured. Having a good grasp of exactly what is being configured and why will help greatly to ensure that your MPLS VPN network operates as smoothly as possible.

Objectives

This lesson identifies the command syntax that is required to configure MP-BGP in an MPLS VPN backbone.

Upon completing this lesson, you will be able to:

- Identify the command syntax that is required to configure BGP address families
- Describe the requirements for enabling BGP neighbors in an MPLS VPN environment
- Identify the command syntax that is required to configure MP-BGP in an MPLS VPN environment
- Identify the command syntax that is required to configure MP-IBGP in an MPLS VPN environment
- Identify the command syntax that is required to configure MP-BGP BGP community propagation in an MPLS VPN environment
- Identify the command syntax that is required to disable IPv4 route exchange in an MPLS VPN environment

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Configuring BGP Address Families
- BGP Neighbors
- Configuring MP-BGP
- Configuring MP-IBGP
- MP-BGP BGP Community Propagation
- Disabling IPv4 Route Exchange
- Summary
- Quiz

Configuring BGP Address Families

This topic identifies the command syntax required to configure BGP address families.

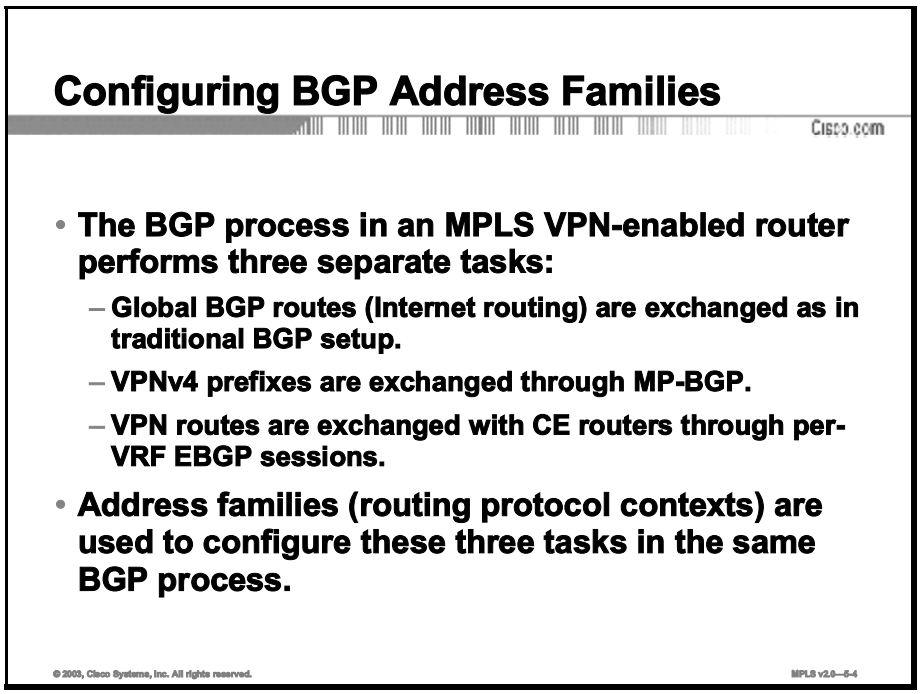

Independently from the MPLS VPN architecture, the PE router can use BGP IPv4 route updates to receive and propagate Internet routes in scenarios where the PE routers are also used to provide Internet connectivity to customers.

The MPLS VPN architecture uses the BGP routing protocol in two different ways:

- VPNv4 routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.
- BGP can be used as the PE-CE routing protocol to exchange VPN routes between the PE routers and the CE routers.

All three route exchange mechanisms take place in one BGP process (because only one BGP process can be configured per router). The routing protocol contexts (called “address families” from the router configuration perspective) are used to configure all three independent route exchange mechanisms.

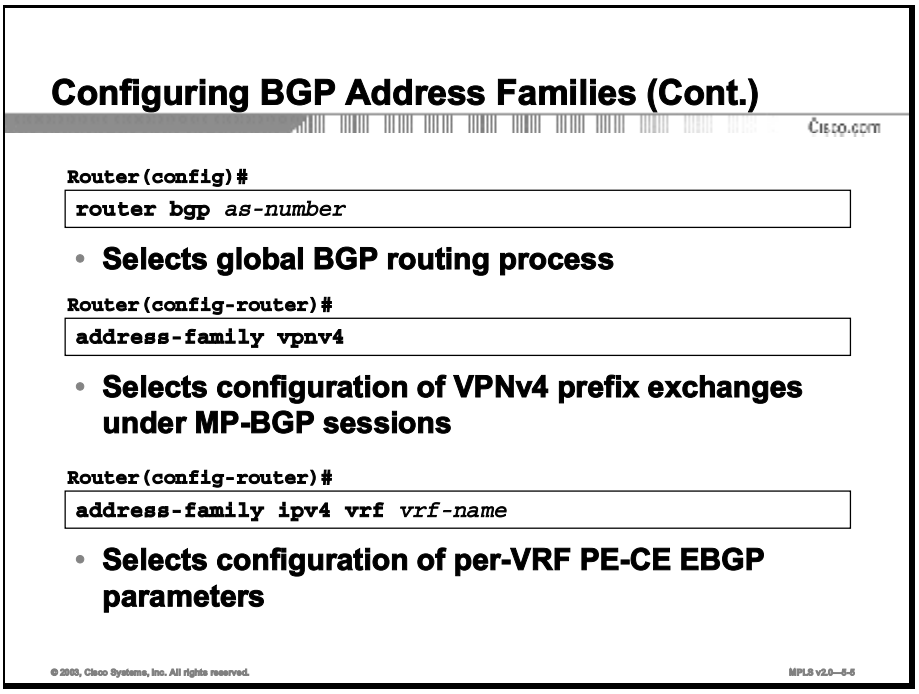

Use the **address-family** command in router configuration mode to select the routing context that you would like to configure:

- Internet routing (global IP routing table) is the default address family that you configure when you start configuring the BGP routing process.
- To configure MP-BGP sessions between the PE routers, use the **vpnv4** address family.
- To configure BGP between the PE routers and the CE routers within individual VRF tables, use the **ipv4 vrf** *vrf-name* address family.

router bgp

To configure the BGP routing process, use the **router bgp** command in global configuration mode. To remove a routing process, use the **no** form of this command:

- **router bgp** *as-number*
- **no router bgp** *as-number*

Syntax Description

Parameter	Description
as-number	Number of an AS that identifies the router to other BGP routers and tags the routing information passed along.

Defaults

No BGP routing process is enabled by default.

address-family

To enter the address family submode for configuring routing protocols, such as BGP, RIP, and static routing, use the **address-family** command in global configuration mode. To disable the address family submode for configuring routing protocols, use the **no** form of this command:

- VPNv4 unicast:
  - **address-family vpnv4 [unicast]**
  - **no address-family vpnv4 [unicast]**
- IPv4 unicast:
  - **address-family ipv4 [unicast]**
  - **no address-family ipv4 [unicast]**
- IPv4 unicast with CE router:
  - **address-family ipv4 [unicast] vrf** *vrf-name*
  - **no address-family ipv4 [unicast] vrf** *vrf-name*

Syntax Description

Parameter	Description
ipv4	Configures sessions that carry standard IPv4 address prefixes.
vpnv4	Configures sessions that carry customer VPNv4 prefixes, each of which has been made globally unique by adding an 8-byte RD.
unicast	(Optional) Specifies unicast prefixes.
vrf *vrf-name*	Specifies the name of a VPN VRF to associate with submode commands.

BGP Neighbors

This topic describes the requirements for enabling BGP neighbors in an MPLS VPN environment.

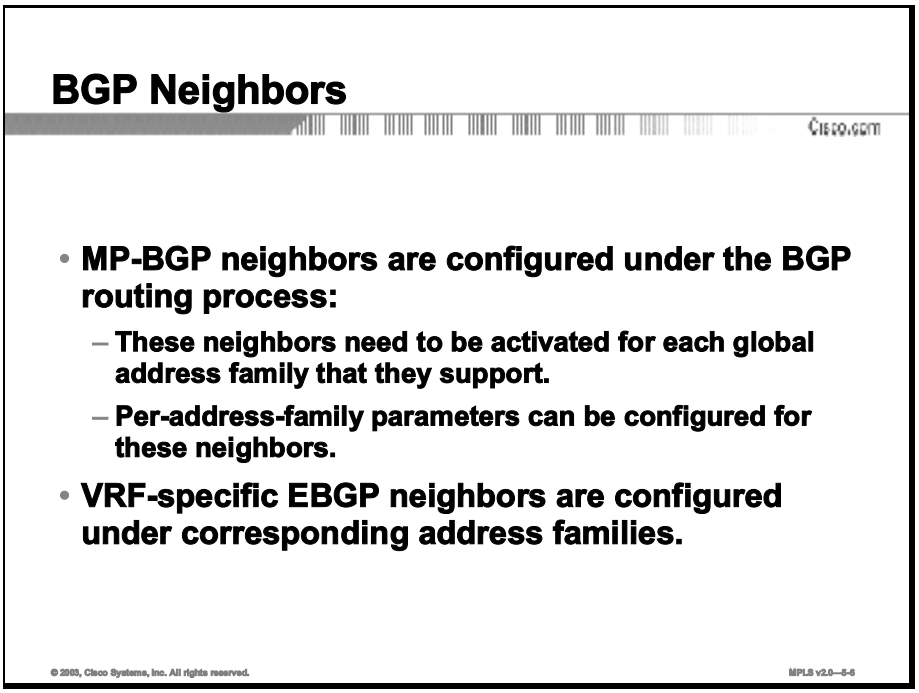

MPLS VPN architecture defines two types of BGP neighbors:

- Global BGP neighbors (other PE routers), with which the PE router can exchange multiple types of routes. These neighbors are defined in the global BGP definition and only have to be *activated* for individual address families.
- Per-VRF BGP neighbors (the CE routers), which are configured and activated within the **ipv4 vrf** *vrf-name* address family.

Configuring MP-BGP

This topic identifies the command syntax required to configure MP-BGP in an MPLS VPN environment.

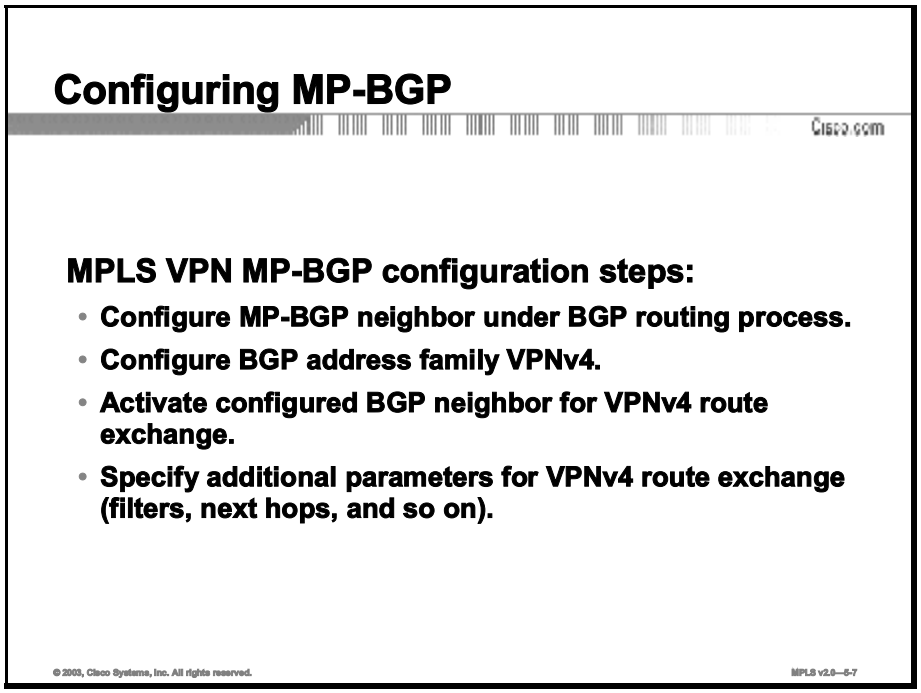

Configure BGP connectivity between two PE routers in four steps:

Step 1 Configure the remote PE router as a global BGP neighbor in BGP router configuration mode.

Step 2 Define the parameters that affect all BGP route exchange (for example, source address for the TCP session) on the global BGP neighbor.

Step 3 Select the VPNv4 address family and activate the BGP neighbor for VPNv4 route exchange.

Step 4 Configure additional VPNv4-specific BGP parameters (filters, next-hop processing, route maps) within the VPNv4 address family.

Note IPv4-specific BGP parameters are still configured under the BGP router configuration mode—there is no special IPv4 address family.

Configuring MP-IBGP

This topic identifies the command syntax required to configure MP-IBGP in an MPLS VPN environment.

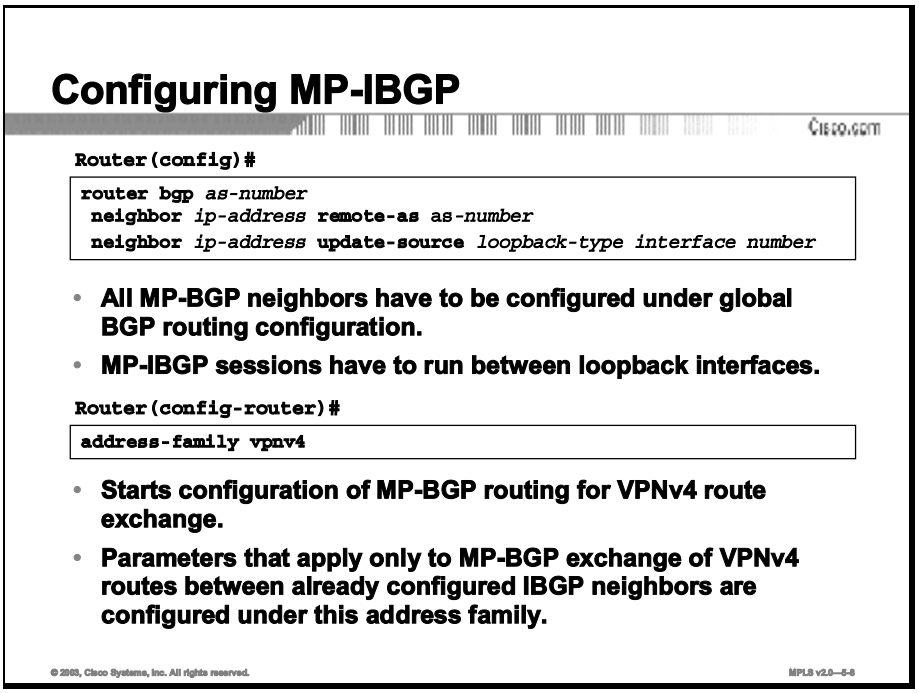

The initial commands needed to configure an MP-IBGP session between PE routers are as follows:

- The **neighbor** *ip-address* **remote-as** *as-number* command configures the neighboring PE router.
- The **neighbor** *ip-address* **update-source** *interface-type interface-number* command configures the source address used for the TCP session carrying BGP updates as well as the IP address used as the BGP next hop for VPNv4 routes.
- The **address-family vpnv4** command allows you to enter VPNv4 configuration mode, where the additional VPNv4-specific parameters have to be configured on the BGP neighbor.

neighbor remote-as

To add an entry to the BGP neighbor table, use the **neighbor remote-as** command in router configuration mode. To remove an entry from the table, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*
- **no neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*

Syntax Description

Parameter	Description
ip-address	Neighbor IP address.
peer-group-name	Name of BGP peer group.
as-number	AS to which the neighbor belongs.

Defaults

There are no BGP neighbor peers.

neighbor update-source

To have the Cisco IOS software allow internal BGP sessions to use any operational interface for TCP connections, use the **neighbor update-source** command in router configuration mode. To restore the interface assignment to the closest interface, which is called the "best local address," use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **update-source** *interface-type*
- **no neighbor** {*ip-address* | *peer-group-name*} **update-source** *interface-type*

Syntax Description

Parameter	Description
ip-address	IP address of the BGP-speaking neighbor.
peer-group-name	Name of BGP peer group.
interface-type	Loopback interface.

Defaults

The default is the best local address.

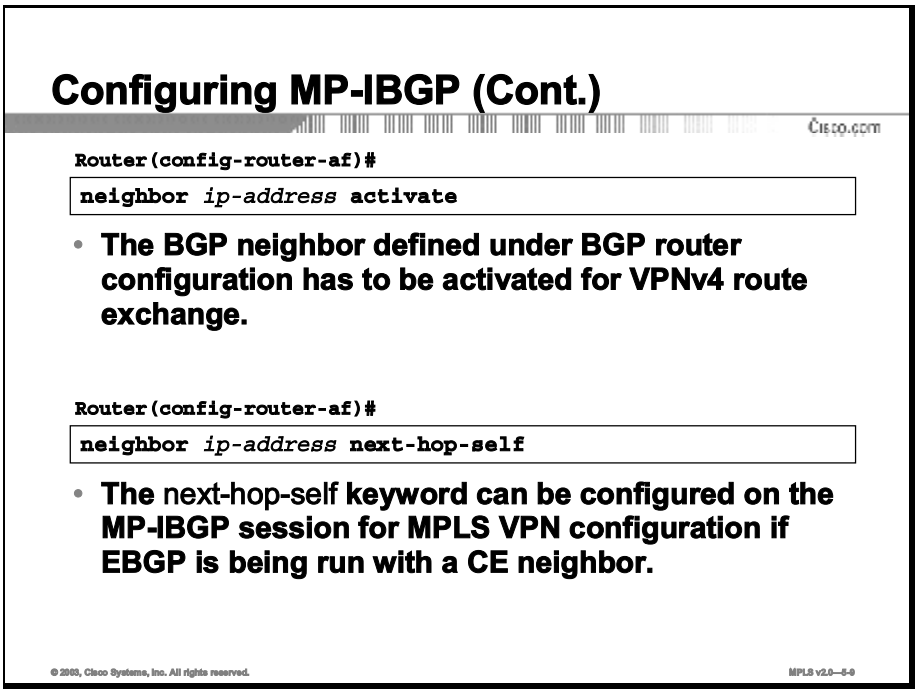

After you define the remote PE router as a global BGP neighbor, you must activate it for VPNv4 route exchange.

neighbor activate

To enable the exchange of information with a BGP neighboring router, use the **neighbor activate** command in router configuration mode. To disable the exchange of an address with a neighboring router, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **activate**
- **no neighbor** {*ip-address* | *peer-group-name*} **activate**

Syntax Description

Parameter	Description
ip-address	IP address of the neighboring router.
peer-group-name	Name of BGP peer group.

Defaults

The exchange of addresses with neighbors is enabled by default for the IPv4 address family. For all other address families, address exchange is disabled by default. You can explicitly activate the default command by using the appropriate address family submode.

neighbor next-hop-self

To disable next-hop processing of BGP updates on the router, use the **neighbor next-hop-self** command in router configuration mode. To disable this feature, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**
- **no neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**

Syntax Description

Parameter	Description
ip-address	IP address of the BGP-speaking neighbor.
peer-group-name	Name of BGP peer group.

Defaults

Default is disabled.

MP-BGP BGP Community Propagation

This topic identifies the command syntax required to configure MP-BGP BGP community propagation in an MPLS VPN environment.

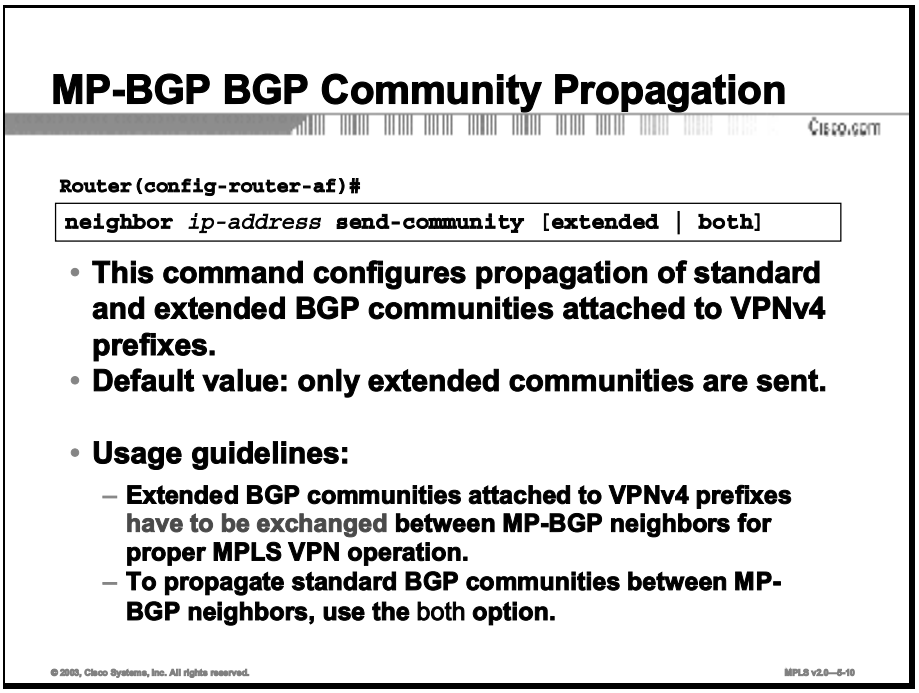

MPLS VPN architecture introduced the "extended community" BGP attribute. BGP still supports the "standard community" attribute, which has not been superseded by the extended communities. The default community propagation behavior for standard BGP communities has not changed. Community propagation still needs to be configured manually. Extended BGP communities are propagated by default because their propagation is mandatory for successful MPLS VPN operation.

The **neighbor send-community** command was extended to support standard and extended communities. Use this command to configure propagation of standard and extended communities if your BGP design relies on use of standard communities. An example of this would be to propagate quality of service (QoS) information across the network.

neighbor send-community

To specify that BGP community attributes that are attached to a BGP route should be sent to a BGP neighbor, use the **neighbor send-community** command in router configuration mode. To remove the entry, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **send-community** [**extended** | **both**]
- **no neighbor** {*ip-address* | *peer-group-name*} **send-community**

Syntax Description

Parameter	Description
ip-address	Neighbor IP address.
peer-group-name	Name of BGP peer group.

Defaults

BGP communities are not propagated to any neighbor.

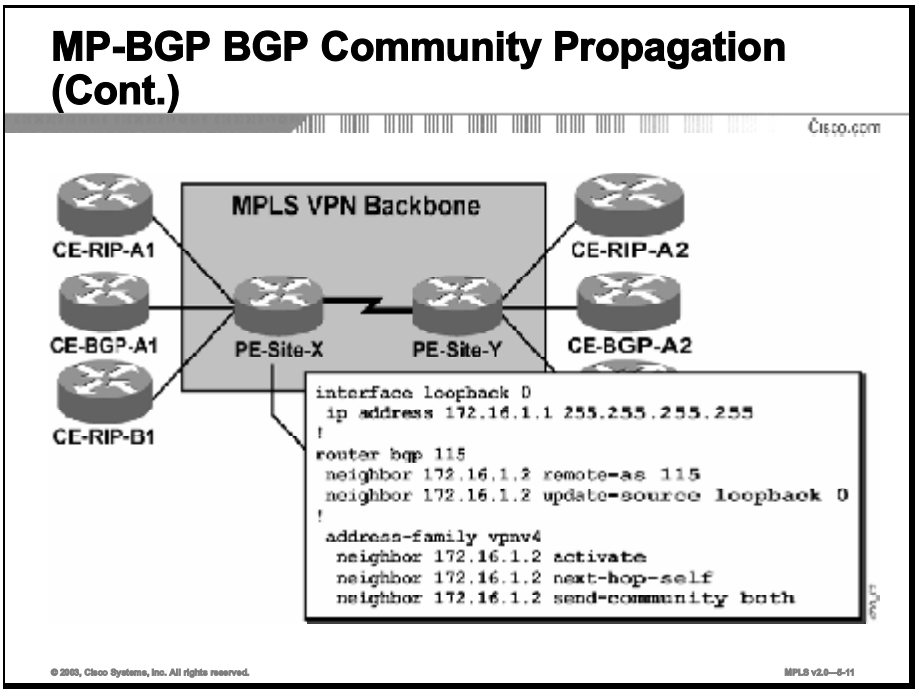

The configuration example that you looked at earlier in the "Configuring VRF Tables" lesson continues here with configuration of MP-IBGP sessions on the PE router. You need to perform these steps:

Configuration of MP-IBGP sessions

Step	Action
1	Define a loopback interface that will serve as the BGP next hop for VPNv4 routes and as the source address for the IBGP session.
2	Configure the remote PE router as the global BGP neighbor.
3	Specify the source address for the TCP session.
4	Select the VPNv4 address family.
5	Activate the remote PE router for VPNv4 route exchange.
6	Disable next-hop processing for VPNv4 route exchange. This action guarantees that the loopback 0 interface will always be the BGP next hop for VPNv4 routes propagated by this router to its MP-IBGP neighbors.
7	Configure propagation of standard and extended communities.

Disabling IPv4 Route Exchange

This topic identifies the command syntax required to disable IPv4 route exchange in an MPLS VPN environment.

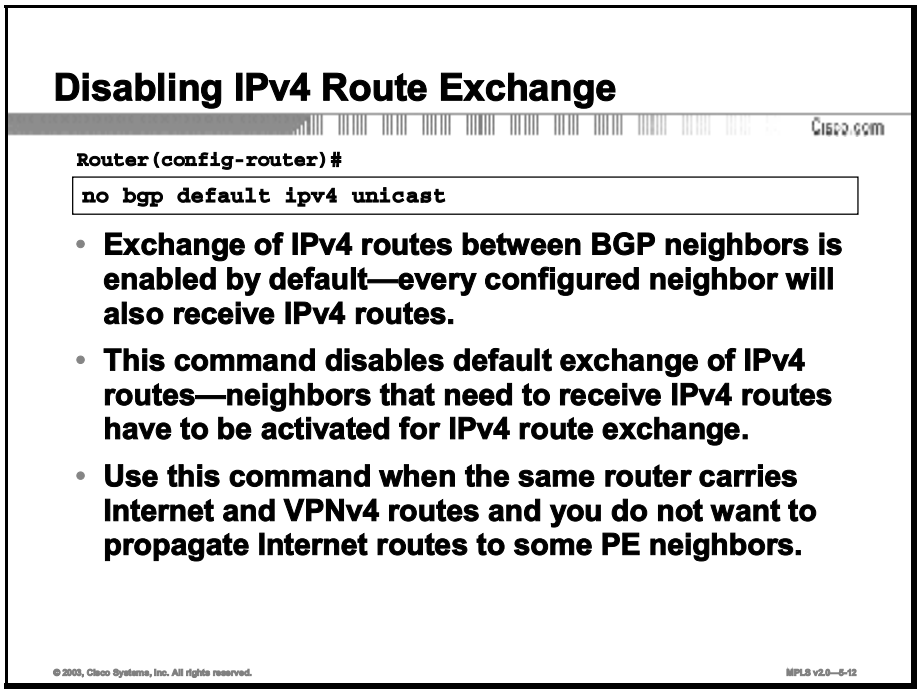

The BGP configuration discussed so far is appropriate for scenarios where the PE routers provide Internet and VPN connectivity. If the PE routers provide only VPN connectivity, they do not need Internet routing and the IPv4 route exchange should be disabled. There are two ways of disabling IPv4 route exchange:

- To disable IPv4 route exchange for only a few neighbors, your best option is to disable the IPv4 route exchange on a neighbor-by-neighbor basis by using the **no neighbor activate** command.
- To disable IPv4 route exchange for most (or all) of the neighbors, you can use the **no bgp default ipv4 unicast** command. After you enter this command, you must manually activate IPv4 route exchange for each configured global BGP neighbor.

Disabling IPv4 Route Exchange (Cont.)

Cisco.com

- Neighbor 172.16.32.14 receives only Internet routes.
- Neighbor 172.16.32.15 receives only VPNv4 routes.
- Neighbor 172.16.32.27 receives Internet and VPNv4 routes.

```
router bgp 12703
 no bgp default ipv4 unicast
 neighbor 172.16.32.14 remote-as 12703
 neighbor 172.16.32.15 remote-as 12703
 neighbor 172.16.32.27 remote-as 12703
! Activate IPv4 route exchange
 neighbor 172.16.32.14 activate
 neighbor 172.16.32.27 activate
! Step#2 - VPNv4 route exchange
 address-family vpnv4
  neighbor 172.16.32.15 activate
  neighbor 172.16.32.27 activate
```


In the figure, only a subset of BGP neighbors needs to receive IPv4 routes. The default propagation of IPv4 routes is thus disabled. IPv4 route exchange, as well as VPNv4 route exchange, is manually activated on a neighbor-by-neighbor basis.

Summary

This topic summarizes the key points discussed in this lesson.

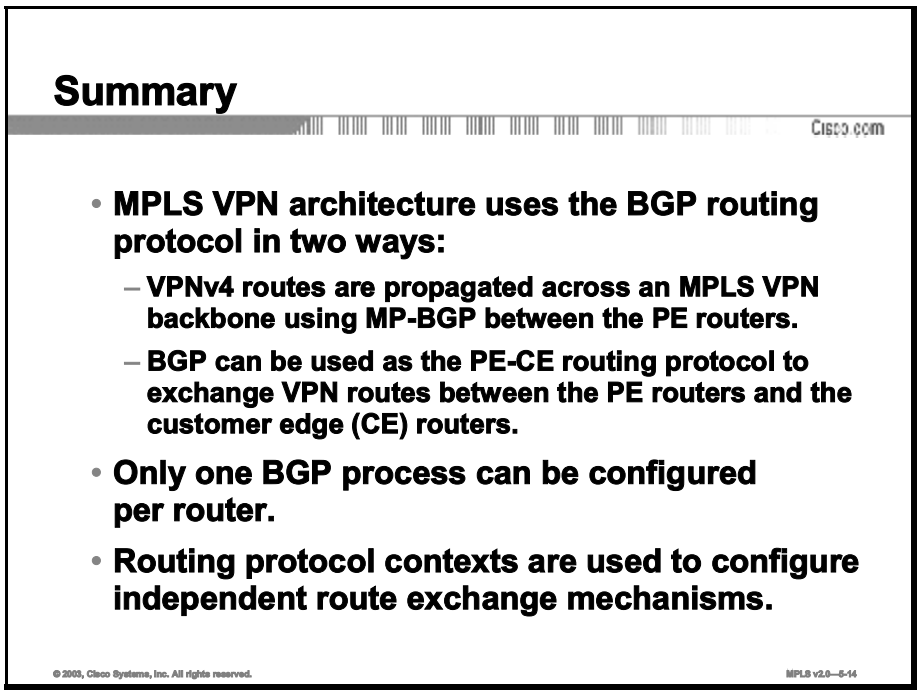

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) In which two ways does the MPLS VPN architecture use the BGP routing protocol?

Q2) What is a BGP address family?

Q3) What are the two types of BGP address families that can be configured on a PE router?

Q4) Which mandatory parameters do you have to configure on an MP-BGP neighbor?

Q5) Why is it necessary to enable extended BGP communities when you are supporting MPLS VPNs?

Q6) Why would you want to disable propagation of IPv4 routing updates between MP-BGP neighbors?

Quiz Answer Key

Q1) EBGP is used to carry routing updates between the PE and the CE.
IBGP (VNPv4) is used to carry VPN route updates between PE routers.

Note IBGP (IPv4) is used to carry *non-VPN* route updates between PE routers.

Relates to: BGP Neighbors; Configuring MP-IBGP; and Configuring MP-IBGP Propagation of All VPNv4 Routes

Q2) A BGP address family is a routing protocol context that is used to configure global BGP routing, VPN routing, and CE-to-PE routing into the same BGP process.

Relates to: Configuring BGP Address Families

Q3) VPNv4 and IPv4

Relates to: Configuring BGP Address Families

Q4) neighbor *ip-address* remote-as *as-number*
neighbor *ip-address* update-source *interface-type interface number*
address-family vpnv4
neighbor *ip-address* activate
neighbor *ip-address* next-hop-self

Relates to: BGP Neighbors

Q5) Extended BGP communities attached to VPNv4 prefixes have to be exchanged between MP-BGP neighbors because they contain the RT information.

Relates to: MP-BGP BGP Community Propagation

Q6) when you are supporting only VPN routes

Relates to: Disabling IPv4 Route Exchange

Configuring Small-Scale Routing Protocols Between PE and CE Routers

Overview

This lesson explains the PE-CE routing protocol configuration steps, and the various routing protocols that you can run between PE and CE routers. These protocols include RIP, EIGRP, and static routes.

Relevance

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need in order to configure MPLS VPN PE-CE routing exchange.

Objectives

The lesson identifies the command syntax that is required to configure PE-CE routing protocols.

Upon completing this lesson, you will be able to:

- Identify the requirements for configuring PE-CE routing protocols
- Identify the command syntax that is used to select VRF routing context within BGP
- Identify the command syntax that is used to configure per-VRF static routes
- Identify the command syntax that is used to configure a RIP PE-CE routing session
- Identify the command syntax that is used to configure an EIGRP PE-CE routing session

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Configuring PE-CE Routing Protocols
- Selecting the VRF Routing Context for BGP and RIP
- Configuring Per-VRF Static Routes
- Configuring RIP PE-CE Routing
- Configuring EIGRP PE-CE Routing
- Summary
- Quiz

Configuring PE-CE Routing Protocols

This topic identifies the requirements for configuring PE-CE routing protocols.

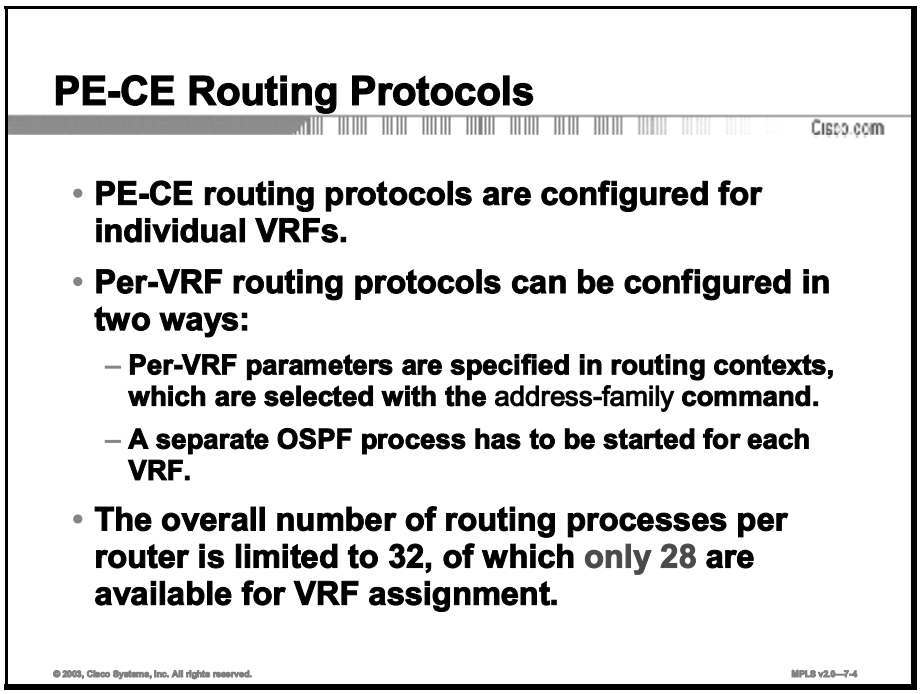

After you configure VRFs and establish MP-IBGP connectivity between PE routers, you have to configure routing protocols between the PE router and the attached CE routers. The PE-CE routing protocols need to be configured for individual VRFs. Sites in the same VPN but in different VRFs cannot share the same PE-CE routing protocol.

Note The per-VRF configuration of the PE-CE routing protocols is another good reason for grouping as many sites into a VRF as possible.

The per-VRF routing protocols can be configured in two ways:

- As individual address families belonging to the same routing process (similar to what you have already seen for BGP).
- As separate routing processes. This option is used for more complex routing protocols that need to maintain a separate topology database for each VRF (for example, OSPF).

Note Current Cisco IOS software implementation limits the overall number of routing protocols in a router to 32. Two routing methods are predefined (static and connected), and two routing protocols are needed for proper MPLS VPN backbone operation (BGP and backbone Interior Gateway Protocol [IGP]). The number of PE-CE routing processes is therefore limited to 28.

Selecting the VRF Routing Context for BGP and RIP

This topic identifies the command syntax used to select the VRF routing context for BGP and RIP.

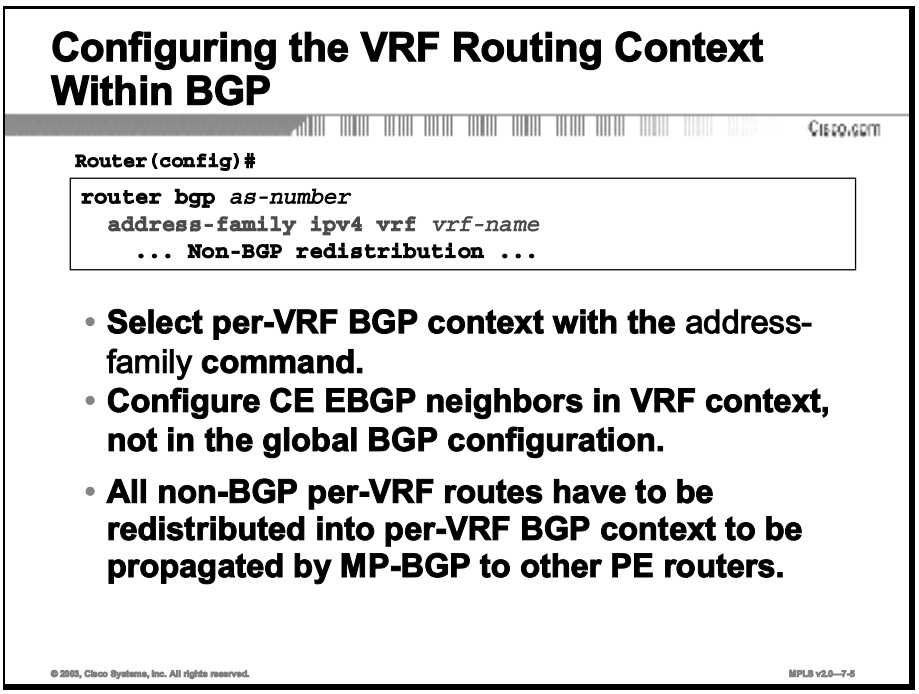

Select the VRF routing context with the **address-family ipv4 vrf** *vrf-name* command in the RIP and BGP routing processes. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

Note	Common parameters defined in router configuration mode are inherited by all address families defined for this routing process and can be overridden for each individual address family.

address-family ipv4

To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the **address-family ipv4** command in router configuration mode. To disable address family configuration mode, use the **no** form of this command:

- **address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]
- **no address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]

Syntax Description

Parameter	Description
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vrf *vrf-name*	(Optional) Specifies the name of the VRF instance to associate with subsequent IPv4 address family configuration mode commands.

Defaults

IPv4 address prefixes are not enabled. Unicast address prefixes are the default when IPv4 address prefixes are configured.

Command Modes

Router configuration

Configuring Per-VRF Static Routes

This topic identifies the command syntax used to configure per-VRF static routes.

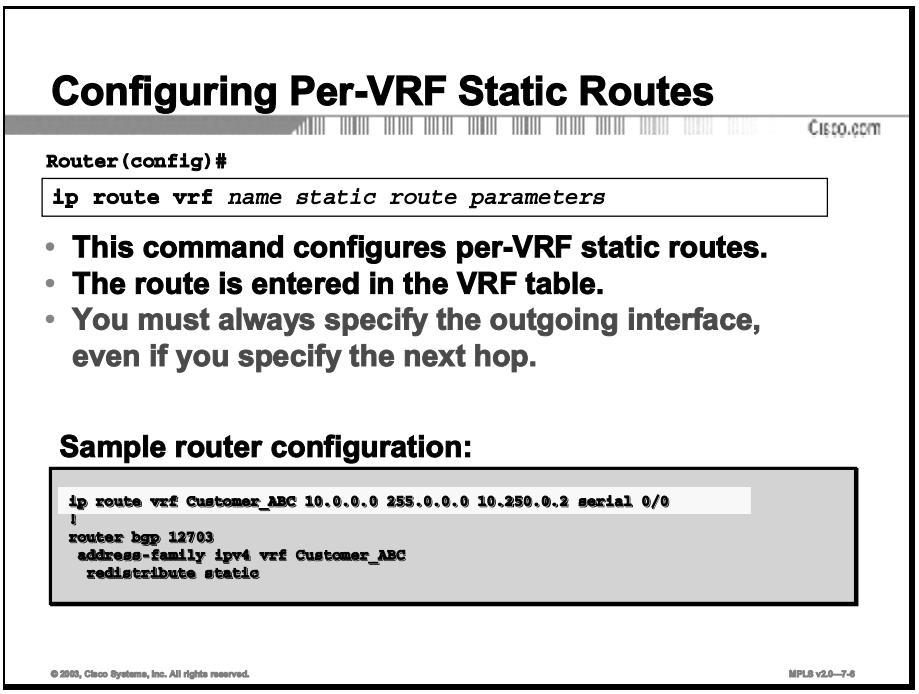

ip route vrf

To establish static routes for a VRF, use the **ip route vrf** command in global configuration mode. To disable static routes, use the **no** form of this command:

- **ip route vrf** *vrf-name prefix mask* [*next-hop-address*] [*interface* {*interface-number*}] [**global**] [*distance*] [**permanent**] [**tag** *tag*]
- **no ip route vrf** *vrf-name prefix mask [next-hop-address*] [*interface* {*interface-number*}] [**global**] [*distance*] [**permanent**] [**tag** *tag*]

Syntax Description

Parameter	Description
vrf-name	Name of the VRF for the static route.
prefix	IP route prefix for the destination, in dotted decimal notation.
mask	Prefix mask for the destination, in dotted decimal notation.
next-hop-address	(Optional) IP address of the next hop (the forwarding router that can be used to reach that network).
interface	Type of network interface to use: ATM, Ethernet, loopback, POS (Packet over SONET), or null.
interface-number	Number identifying the network interface to use.
global	(Optional) Specifies that the given next-hop address be in the non-VRF routing table.
distance	(Optional) An administrative distance for this route.
permanent	(Optional) Specifies that this route will not be removed, even if the interface shuts down.
tag *tag*	(Optional) Label (tag) value that can be used for controlling redistribution of routes through route maps.

Configuring RIP PE-CE Routing

This topic identifies the command syntax used to configure a RIP PE-CE routing session.

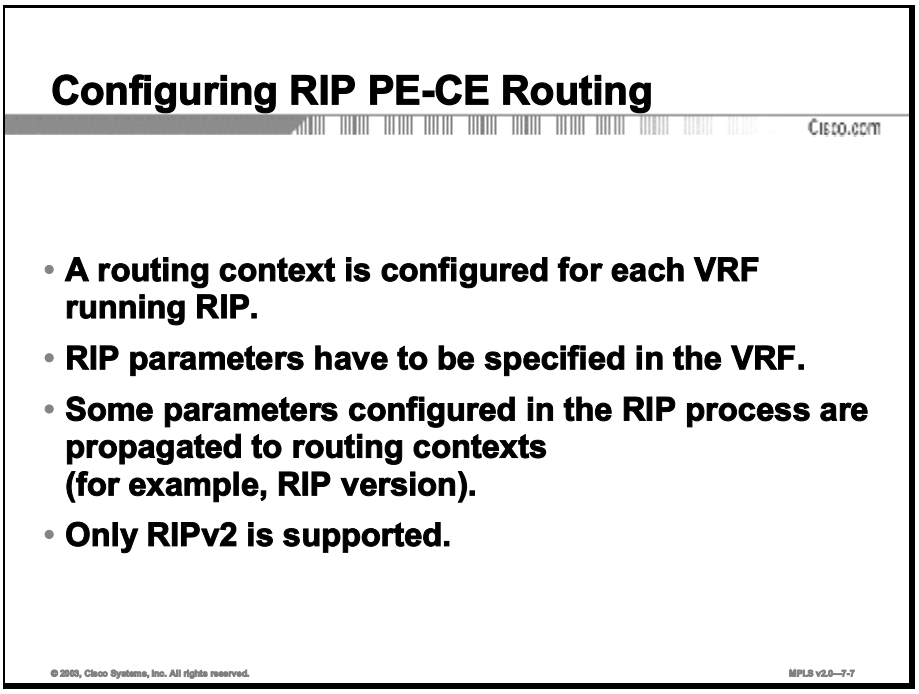

Configuring RIP as the PE-CE routing protocol is even easier than configuring BGP. Start the configuration of individual routing context with the **address-family ipv4 vrf** *vrf-name* command in router configuration mode. You can enter all standard RIP parameters in the per-VRF routing context. Global RIP parameters entered in the scope of RIP router configuration are inherited by each routing context and can be overwritten if needed in each routing context.

Note	Only RIPv2 is supported as the PE-CE routing protocol. It is good configuration practice to configure RIP version as a global RIP parameter using the **version 2** command in router configuration mode.

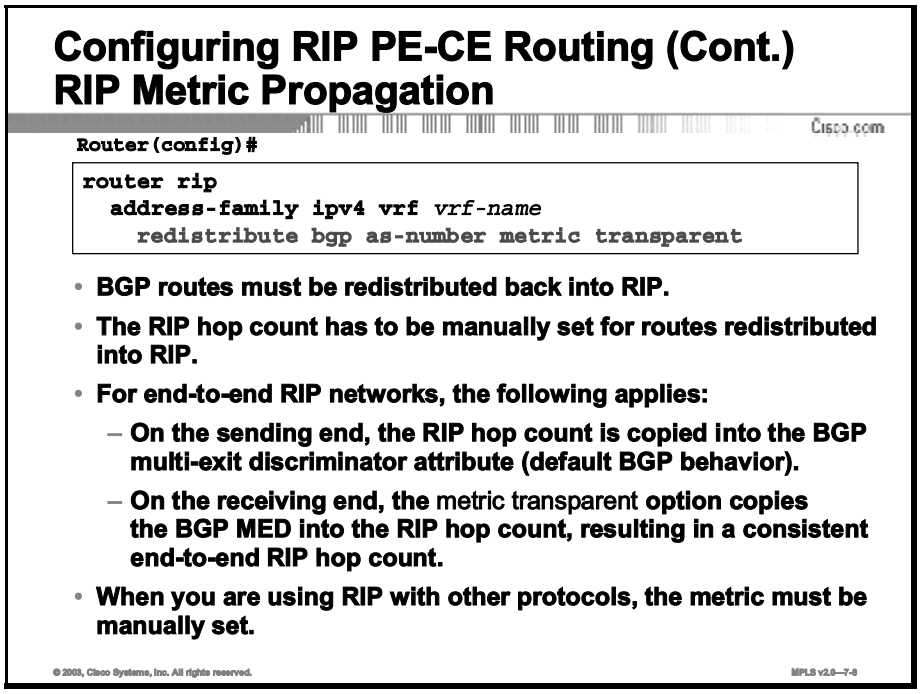

The IGP metric is always copied into the multi-exit discriminator (MED) attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. It is not copied back into the IGP metric. The IGP metric has to be specified in the **redistribute** command or by using the **default-metric** command in router configuration mode.

The MPLS VPN extension to the **redistribute** command (**metric transparent** option) allows the MED attribute to be inserted as the IGP metric of a route redistributed from BGP back into RIP. This extension gives transparent end-to-end (from the customer perspective) RIP routing:

- By default, the RIP hop count is inserted into the BGP attribute MED when the RIP route is redistributed into BGP by the ingress PE router.
- You can configure the value of the MED attribute (the original RIP hop count) to be copied into the RIP hop count when the BGP route is redistributed back into RIP. This action causes the whole MPLS VPN backbone to appear as a single hop to the CE routers.

Note You should *not* change the MED value within BGP if you use the **redistribute metric transparent** command.

Configuring RIP PE-CE Routing (Cont.)

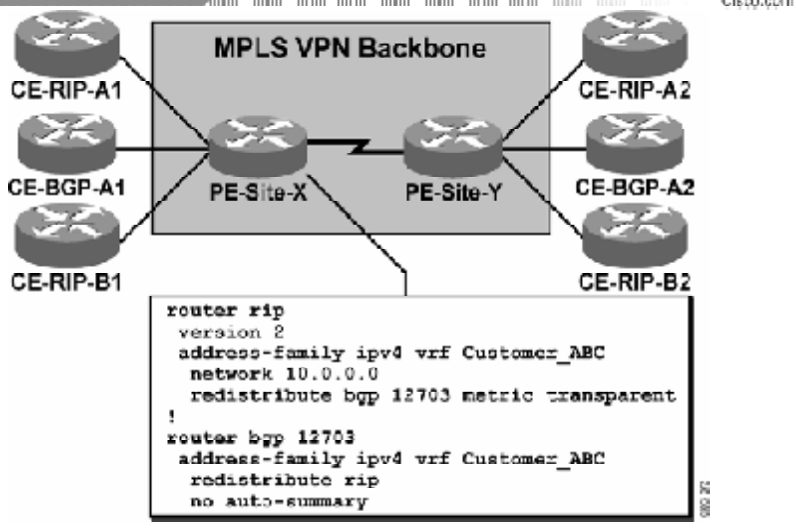

The RIP configuration in this sample network is exceedingly simple:

- The RIP routing process is configured. The RIP version is configured as the global RIP parameter.
- The RIP routing context is configured for every VRF where you want to run RIP as the PE-CE routing protocol. The directly connected networks (configured on interfaces in the VRF) over which you want to run RIP are specified to have standard RIP configuration.
- Redistribution from BGP into RIP with metric propagation is configured.
- The BGP routing context is configured for every VRF. Redistribution of RIP routes into BGP has to be configured for every VRF for which you have configured the RIP routing context.

Configuring EIGRP PE-CE Routing

This topic identifies the command syntax used to configure an EIGRP PE-CE routing session.

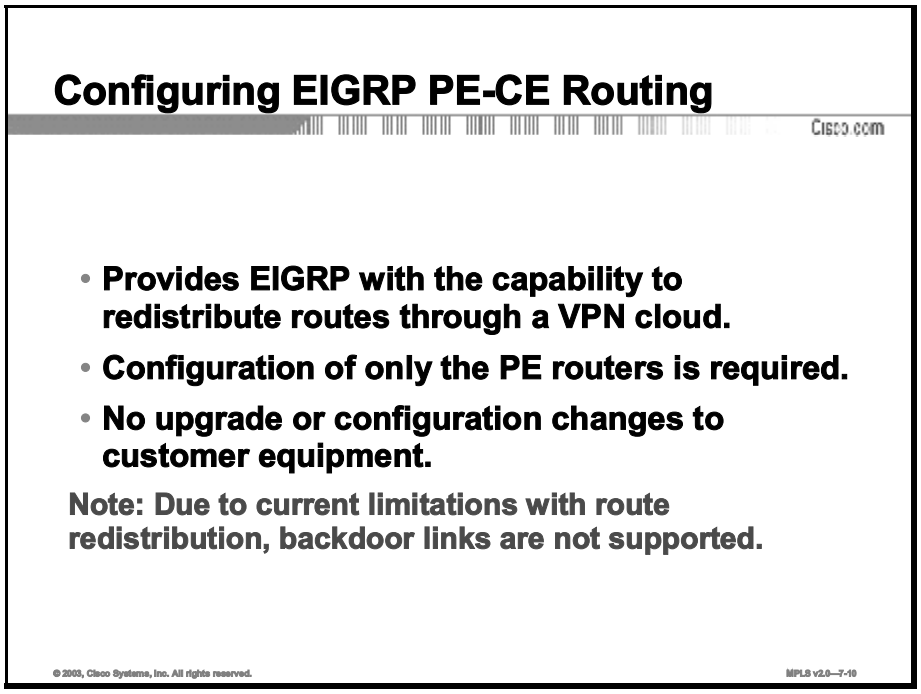

MPLS VPN support for EIGRP between PE and CE provides EIGRP with the capability to redistribute routes through a BGP VPN cloud. This feature is configured only on PE routers, requiring no upgrade or configuration changes to customer equipment. This feature also introduces EIGRP support for MPLS and BGP extended community attributes.

Note	Due to current limitations with route redistribution, backdoor links are not supported. If a backdoor link is implemented, it may become active and override selection of the VPN links.

Configuring EIGRP PE-CE Routing (Cont.)
EIGRP Metric Propagation

Cisco.com

```
Router(config)#
router eigrp as-number
  address-family ipv4 vrf vrf-name
   autonomous-system as-number
   redistribute bgp as-number metric
```

- **Enables the CE's the EIGRP AS number of the CE under the address family.**
- **Configure per-instance AS number.**
- **Configure router redistribution.**
- **External routes received without the configured metric are not to be advertised to the CE router.**
  - **The metric can be configured in the** redistribute **statement using the** redistribute **command or configured with the** default-metric **command.**

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. It is not copied back into the IGP metric. The metric must be configured for routes from external EIGRP autonomous systems and non-EIGRP networks before these routes can be redistributed into an EIGRP CE router. The metric can be configured in the **redistribute** statement using the **redistribute** (IP) command or configured with the **default-metric** (EIGRP) command.

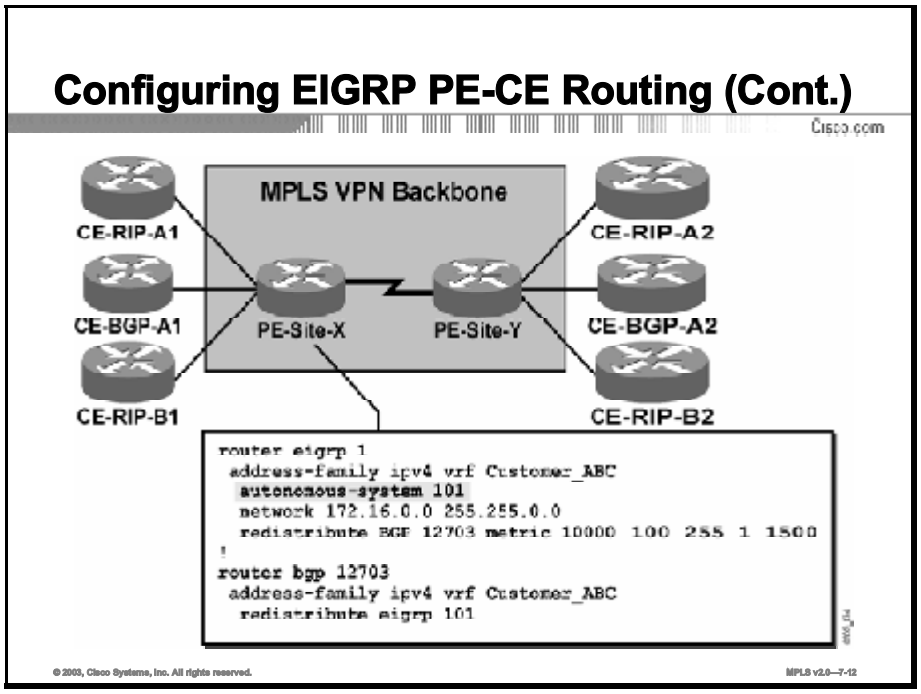

The EIGRP configuration in this sample network is exceedingly simple:

- The EIGRP routing process is configured. The EIGRP version is configured as the global EIGRP parameter.
- The EIGRP routing context is configured for every VRF where you want to run EIGRP as the PE-CE routing protocol. The directly connected networks (configured on interfaces in the VRF) over which you want to run EIGRP are specified to have standard EIGRP configuration.
- Redistribution from BGP into EIGRP with metric propagation is configured.
- The BGP routing context is configured for every VRF. Redistribution of EIGRP routes into BGP has to be configured for every VRF for which you have configured the EIGRP routing context.

Summary

This topic summarizes the key points discussed in this lesson.

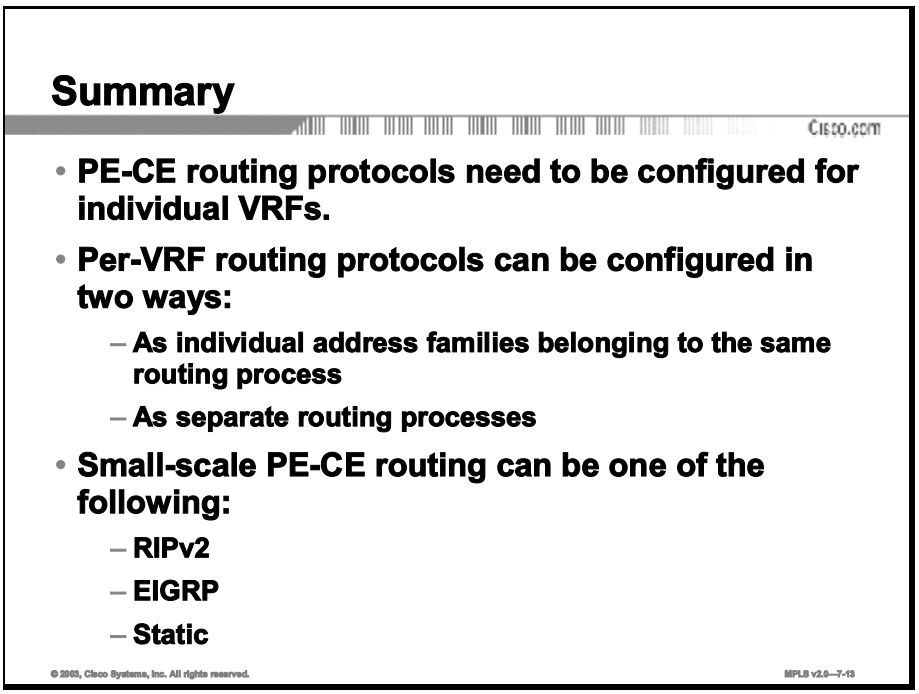

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) How do you configure the routing context in RIP?

Q2) How do you propagate static VRF routes between PE routers?

Q3) How would you configure redistribution to propagate customer RIP routing updates across the MPLS VPN backbone?

Quiz Answer Key

Q1) On the CE enable RIP. On the PE enable RIP and use the **address-family ipv4 vrf** *vrf-name* command under the router RIP section.

Relates to: Selecting the VRF Routing Context for BGP and RIP

Q2) by enabling the static route using the **ip route vrf** *name static route parameters* command

Relates to: Configuring Per-VRF Static Routes

Q3) By using the **redistribute rip** command under the BGP address family on the ingress PE router to redistribute the RIP updates into MP-BGP. MP-BGP uses VPNv4 updates to propagate the updates to the egress PE router. The **redistribute bgp metric transparent** command under the RIP address family is used on the egress PE to redistribute the updates back into RIP.

Relates to: Configuring RIP PE-CE Routing

Monitoring MPLS VPN Operations

Overview

This lesson presents the commands, syntax, and descriptions for monitoring VRF routing, MP-BGP sessions, and VPN status.

Relevance

It is important to understand the network that you have just configured and ensure that it is operating optimally. This lesson will explain how to monitor an MPLS VPN network to ensure that it is functioning smoothly.

Objectives

This lesson identifies the command syntax that is used to monitor MPLS VPN operations.

Upon completing this lesson, you will be able to:

- Identify the command syntax that is used to monitor VRF information
- Identify the command syntax that is used to monitor VRF routing
- Identify the command syntax that is used to monitor MP-BGP sessions
- Identify the command syntax that is used to monitor an MP-BGP VPNv4 table
- Identify the command syntax that is used to monitor per-VRF CEF and LFIB structures
- Identify the command syntax that is used to monitor labels associated with VPNv4 routes
- Identify the command syntax that is used with other MPLS VPN monitoring commands

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Monitoring VRFs
- Monitoring VRF Routing
- Monitoring MP-BGP Sessions
- Monitoring an MP-BGP VPNv4 Table
- Monitoring Per-VRF CEF and LFIB Structures
- Monitoring Labels Associated with VPNv4 Routes
- Other MPLS VPN Monitoring Commands
- Summary
- Quiz

Monitoring VRFs

This topic identifies the command syntax that is used to monitor VRF information.

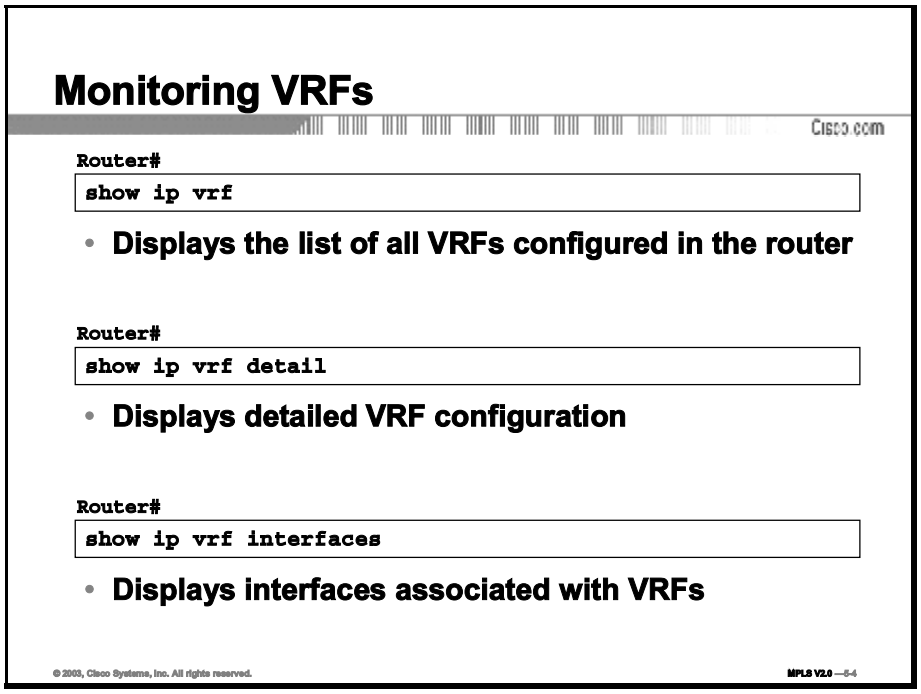

show ip vrf

To display the set of defined VRFs and associated interfaces, use the **show ip vrf** command in EXEC mode:

- **show ip vrf** [{**brief** | **detail** | **interfaces**}] [*vrf-name*] [*output-modifiers*]

Syntax Description

Parameter	Description
brief	(Optional) Displays concise information on the VRF(s) and associated interfaces.
detail	(Optional) Displays detailed information on the VRF(s) and associated interfaces.
interfaces	(Optional) Displays detailed information about all interfaces bound to a particular VRF or to any VRF.
vrf-name	(Optional) Name assigned to a VRF.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Defaults

When no optional parameters are specified, the command shows concise information about all configured VRFs.

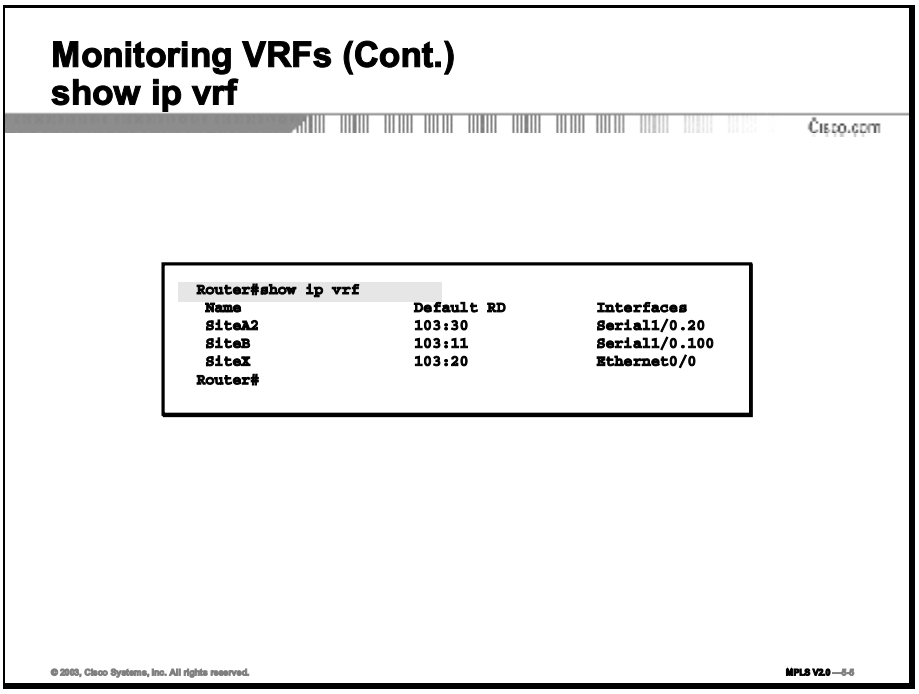

The **show ip vrf** command displays concise information about the VRF(s) and associated interfaces. The following table describes the fields displayed by this command.

Fields	Description
`Name`	Specifies the VRF name.
`Default RD`	Specifies the default RD.
`Interfaces`	Specifies the network interfaces.

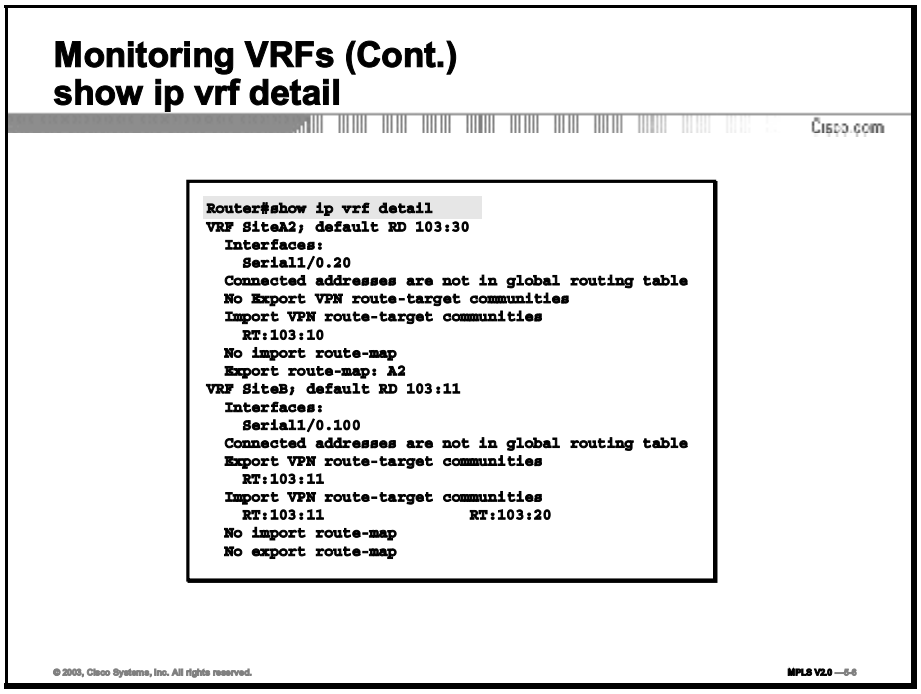

To display detailed information on the VRFs and associated interfaces, use the **show ip vrf detail** command. The following table describes the additional fields shown by this command.

Field	Description
`Interfaces`	Specifies the network interfaces.
`Export`	Specifies VPN RT export communities.
`Import`	Specifies VPN RT import communities.

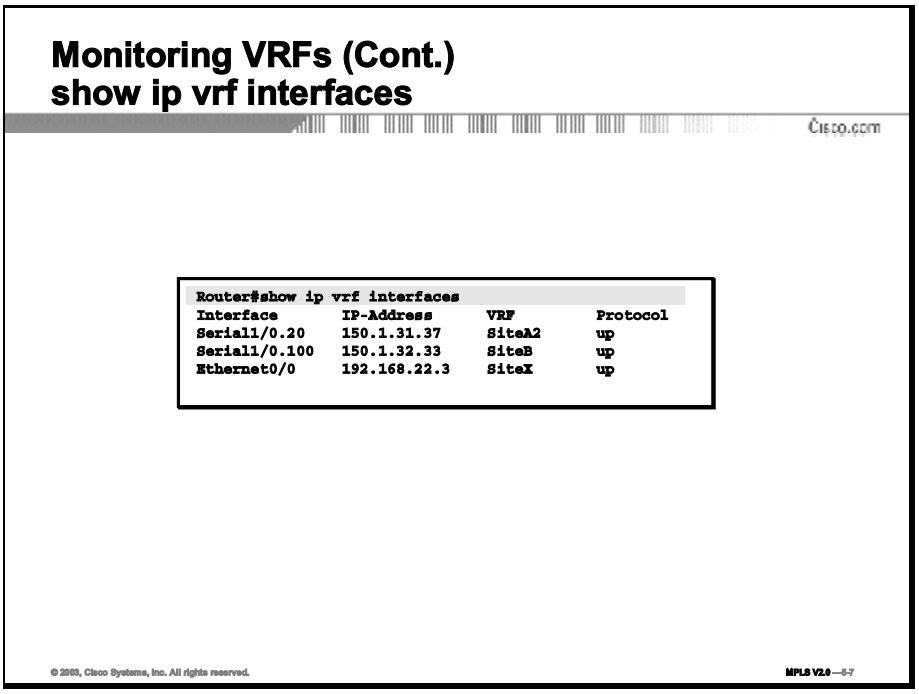

To display the interfaces bound to a particular VRF (or interfaces bound to any VRF), use the **show ip vrf interfaces** command, which displays the fields described in the following table.

Field	Description
`Interface`	Specifies the network interfaces for a VRF.
`IP-Address`	Specifies the IP address of a VRF interface.
`VRF`	Specifies the VRF name.
`Protocol`	Displays the state of the protocol (up/down) for each VRF interface.

Monitoring VRF Routing

This topic identifies the command syntax that is used to monitor VRF routing.

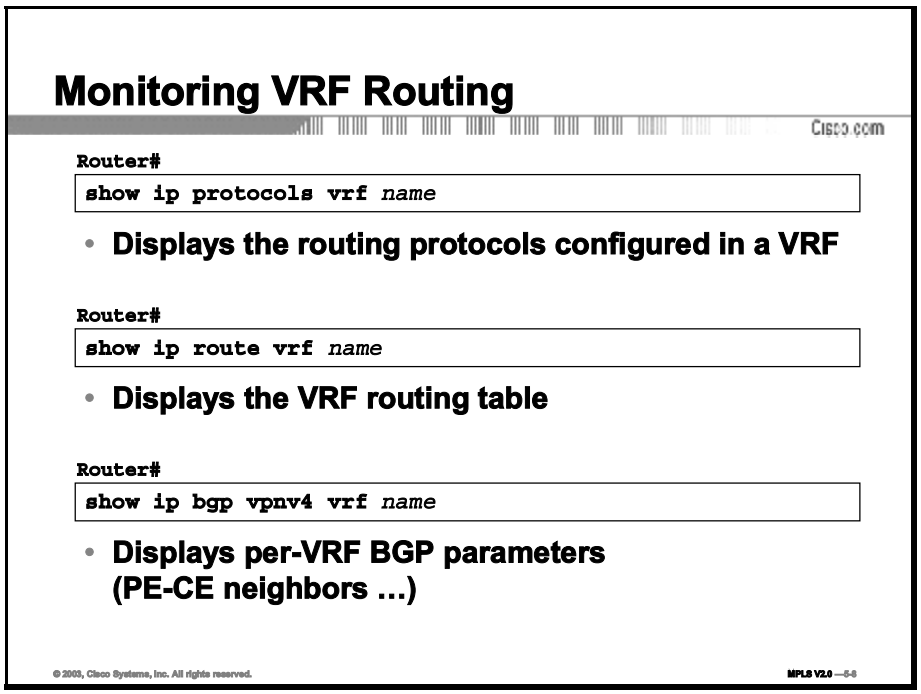

Three commands can be used to monitor VRF routing:

- **show ip protocols vrf** displays the summary information about routing protocols running in a VRF.
- **show ip route vrf** displays the VRF routing table.
- **show ip bgp vpnv4 vrf** displays the VRF BGP table.

show ip protocols vrf

To display the routing protocol information associated with a VRF, use the **show ip protocols vrf** command in EXEC mode:

- **show ip protocols vrf** *vrf-name*

Syntax Description

Parameter	Description
vrf-name	Name assigned to the VRF.

show ip route vrf

To display the IP routing table associated with a VRF, use the **show ip route vrf** command in EXEC mode:

- **show ip route vrf** *vrf-name* [**connected**] [*protocol* [*as-number*] [*tag*] [*output-modifiers*]] [**list** *number* [*output-modifiers*]] [**profile**] [**static** [*output-modifiers*]] [**summary** [*output-modifiers*]] [**supernets-only** [*output-modifiers*]] [**traffic-engineering** [*output-modifiers*]]

Syntax Description

Parameter	Description
vrf-name	Name assigned to the VRF.
connected	(Optional) Displays all connected routes in a VRF.
protocol	(Optional) To specify a routing protocol, use one of the following keywords: **bgp**, **egp**, **eigrp**, **hello**, **igrp**, **isis**, **ospf**, or **rip**.
as-number	(Optional) Autonomous system number.
tag	(Optional) Cisco IOS software routing area label.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
list *number*	(Optional) Specifies the IP access list to display.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
supernets-only	(Optional) Displays supernet entries only.
traffic-engineering	(Optional) Displays only traffic-engineered routes.

show ip bgp vpnv4

To display VPN address information from the BGP table, use the **show ip bgp vpnv4** command in EXEC mode:

- **show ip bgp vpnv4** {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} [*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]] [*network-address* [*mask*] [**longer-prefixes**] [*output-modifiers*]] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**labels**]

Syntax Description

Parameter	Description
all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays Network Layer Reachability Information (NLRI) prefixes that have a matching RD.
vrf *vrf-name*	Displays NLRI prefixes associated with the named VRF.
ip-prefix/length	(Optional) IP prefix address (in dotted decimal notation) and length of mask (0 to 32).
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter, as well as all entries that match the prefix in a "longest-match" sense—that is, prefixes for which the specified prefix is an initial substring.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
network-address	(Optional) IP address of a network in the BGP routing table.
mask	(Optional) Mask of the network address, in dotted decimal notation.
cidr-only	(Optional) Displays only routes that have non-natural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed on account of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes conforming to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP AS paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the AS path "regular expression."
regexp	(Optional) Displays routes matching the AS path "regular expression."
summary	(Optional) Displays BGP neighbor status.
tags	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Monitoring VRF Routing (Cont.)
show ip protocols vrf

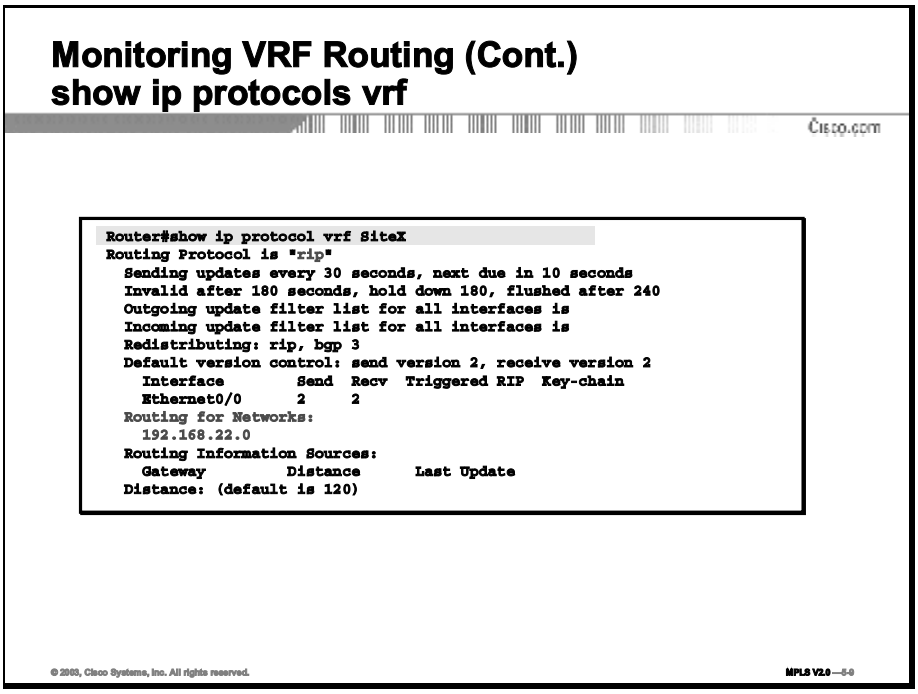

The **show ip protocols vrf** command displays summary information about all routing protocol instances active in the specified VRF. The fields displayed by this command are shown in the following table.

Field	Description
Gateway	Displays the IP address of the router identifier for all routers in the network.
Distance	Displays the metric used to access the destination route.
Last Update	Displays the last time that the routing table was updated from the source.

Monitoring VRF Routing (Cont.)
show ip route vrf

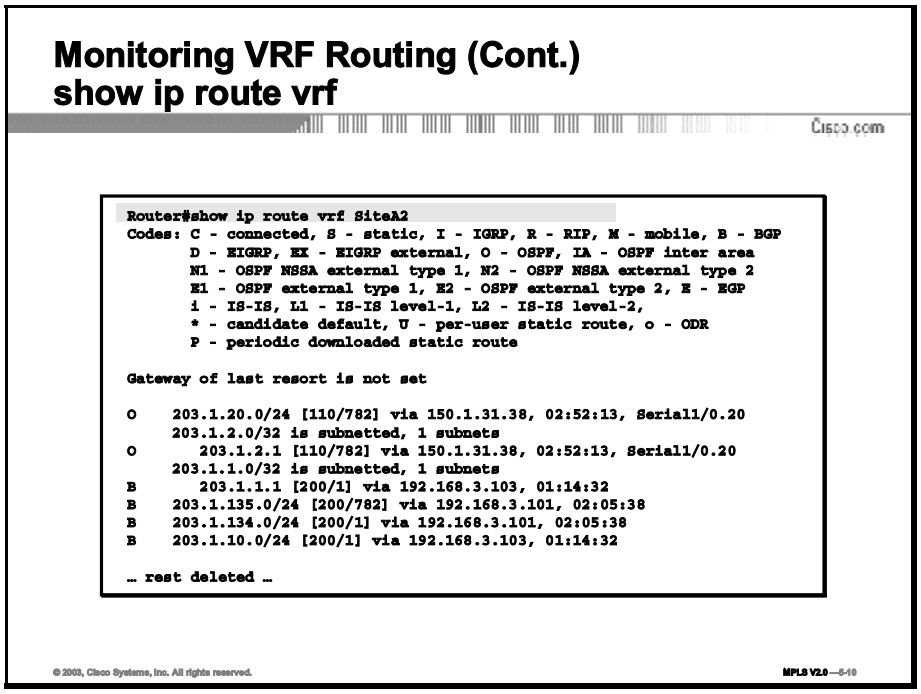

The **show ip route vrf** command displays the contents of the VRF IP routing table in the same format used by the **show ip route** command.

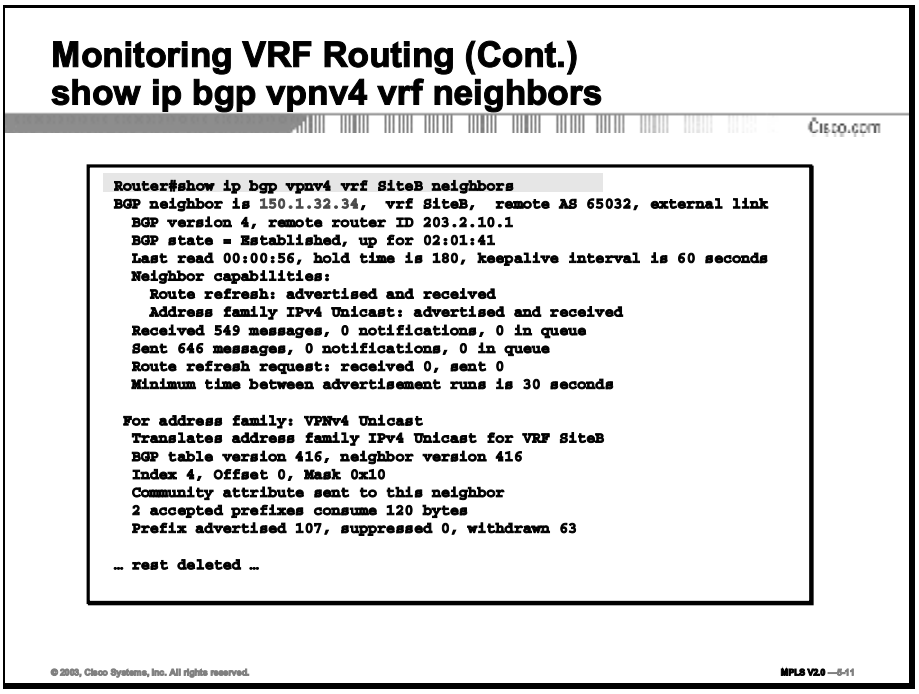

show ip bgp vpnv4 vrf neighbors

To display BGP neighbors configured in a VRF, use the **show ip bgp vpnv4 vrf neighbors** command in privileged EXEC mode:

- **show ip bgp vpnv4** {**all** | **vrf** *vrf-name*} **neighbors**

Syntax Description

Parameter	Description
vpnv4	Specifies VPNv4 information.
all	Displays the complete VPNv4 database.
vrf *vrf-name*	Displays neighbors associated with the named VRF.
neighbors	Displays details about TCP and BGP neighbor connections.

Defaults

This command has no default values.

Usage Guidelines

Use this command to display detailed information about BGP neighbors associated with the MPLS VPN architecture.

Monitoring MP-BGP Sessions

This topic identifies the command syntax that is used to monitor MP-BGP sessions.

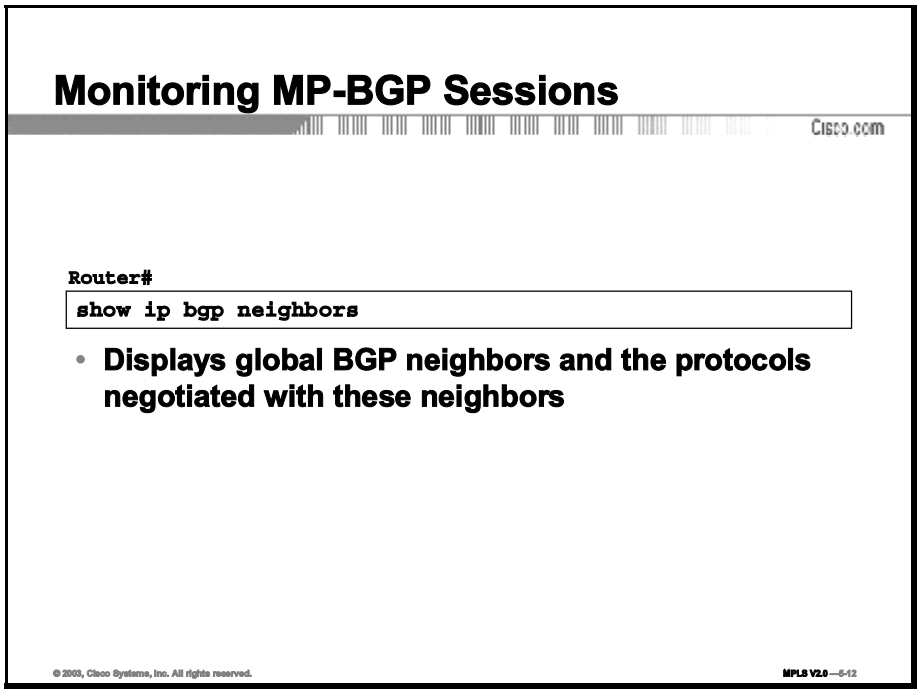

The **show ip bgp neighbors** command is described in detail in the Cisco IOS documentation. It is used to monitor BGP sessions with other PE routers, as well as the address families negotiated with these neighbors.

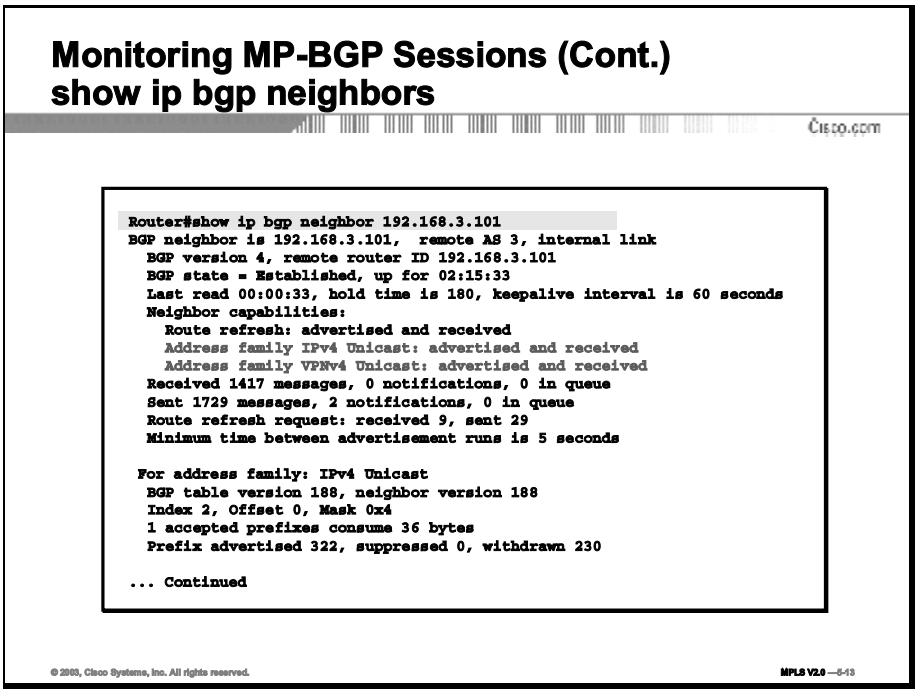

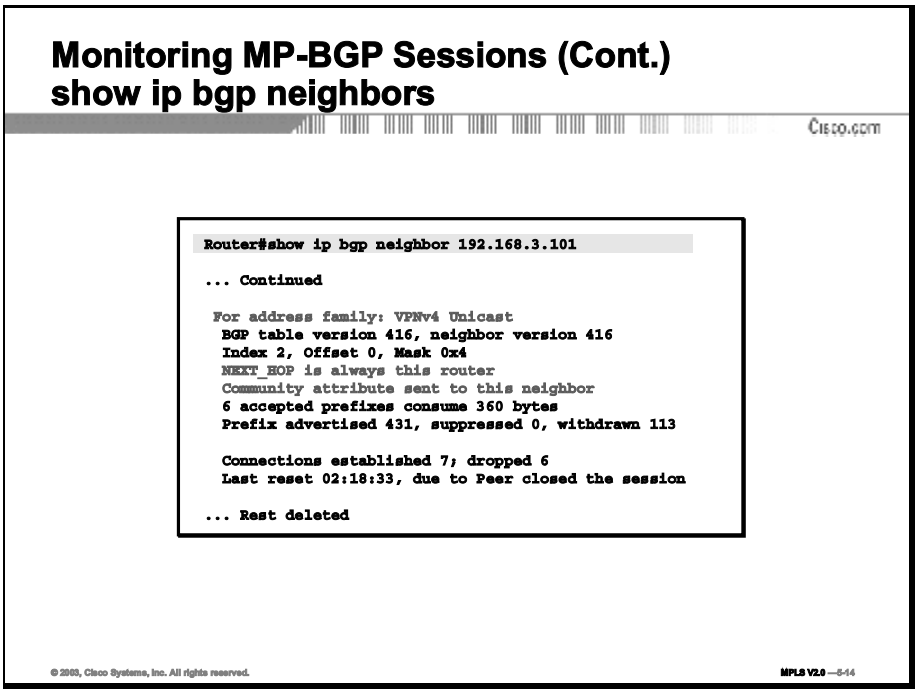

show ip bgp neighbors

To display information about the TCP and BGP connections to neighbors, use the **show ip bgp neighbors** command in EXEC mode:

- **show ip bgp neighbors** [*neighbor-address*] [**received-routes** | **routes** | **advertised-routes** | {**paths** *regexp*} | **dampened-routes**]

Syntax Description

Parameter	Description
neighbor-address	(Optional) Address of the neighbor whose routes you have learned from. If you omit this argument, all neighbors will be displayed.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. This parameter is a subset of the output from the **received-routes** keyword.
advertised-routes	(Optional) Displays all the routes that the router has advertised to the neighbor.
paths *regexp*	(Optional) Regular expression that is used to match the paths received.
dampened-routes	(Optional) Displays the dampened routes to the neighbor at the IP address specified.

Example

The following is sample output from the **show ip bgp neighbors** command:

```
Router# sh ip bgp nei 192.168.100.129
BGP neighbor is 192.168.100.129,  remote AS 65001, internal
link
 BGP version 4, remote router ID 192.168.100.129
  BGP state = Established, up for 5d01h
  Last read 00:00:56, hold time is 180, keepalive interval is
60 seconds
Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received


(output continued)
```

The following table describes the fields shown in the sample output.

Field	Description
BGP neighbor	IP address of the BGP neighbor and its AS number. If the neighbor is in the same AS as the router, the link between them is internal; otherwise, it is considered external.
remote AS	Autonomous system of the neighbor.
external link internal link	Indicates that this peer is either an EBGP peer or an IBGP peer.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	IP address of the neighbor.
BGP state	Internal state of this BGP connection.
up for	Amount of time, in seconds, that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period, in seconds, between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. Note: A state of **advertised and received** indicates a active neighbor relationship.

Router# show ip bgp neighbors 192.168.100.129 (Example continued)

```
For address family: IPv4 Unicast
  BGP table version 31, neighbor version 31
  Index 1, Offset 0, Mask 0x2
                              Sent       Rcvd
  Prefix activity:            ----       ----
    Prefixes Current:            0         30 (Consumes
1440 bytes)
    Prefixes Total:              0         30
    Implicit Withdraw:           0          0
    Explicit Withdraw:           0          0
    Used as bestpath:          n/a         30
    Used as multipath:         n/a          0
                              Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:              30        n/a
    Total:                                30          0
Number of NLRIs in the update sent: max 0, min 0

For address family: VPNv4 Unicast
  BGP table version 30, neighbor version 30
  Index 4, Offset 0, Mask 0x10
  NEXT_HOP is always this router
  Community attribute sent to this neighbor
                              Sent       Rcvd
```



```
  Prefix activity:                ----       ----
    Prefixes Current:               9          1 (Consumes 256
bytes)
    Prefixes Total:                18          1
    Implicit Withdraw:              9          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          4
    Used as multipath:            n/a          0

                                    Outbound    Inbound
  Local Policy Denied Prefixes:     --------    -------
    VPN Imported prefix:                   3        n/a
    Bestpath from this peer:               1        n/a
    Total:                                 4          0
  Number of NLRIs in the update sent: max 4, min 0
(output omitted)
```

The following table describes the fields shown in the sample output.

Field	Description
`Address family IPv4 Unicast:`	IP Version 4 unicast-specific properties of this neighbor.
`Address family VPNv4`	VPNv4-specific properties of this neighbor.

Note For detailed information, please consult the Cisco IOS reference manual.

Monitoring an MP-BGP VPNv4 Table

This topic identifies the command syntax used to monitor an MP-BGP VPNv4 table.

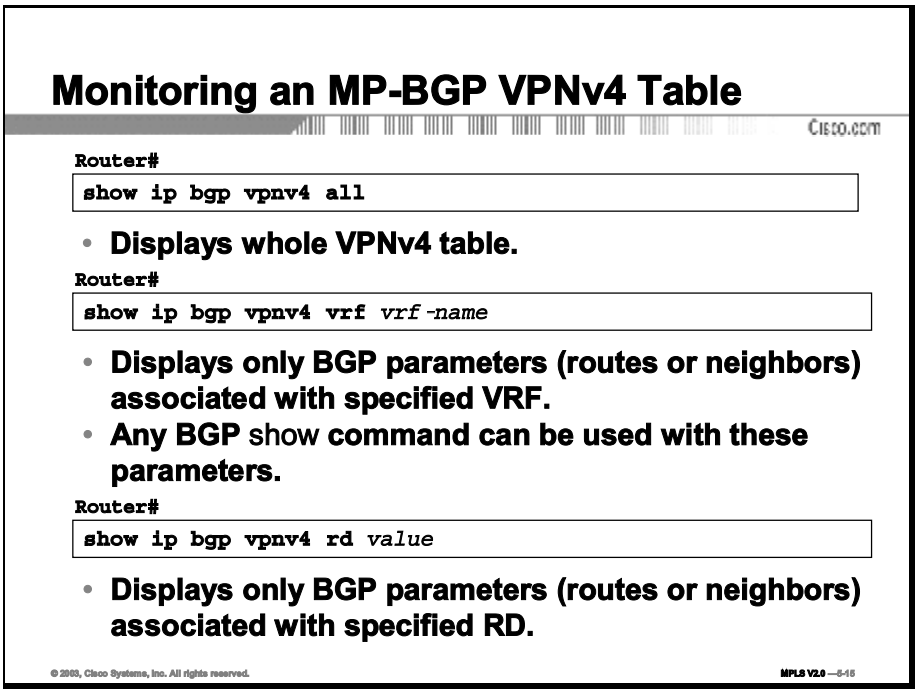

The **show ip bgp vpnv4** command displays IPv4 BGP information as well as VPNv4 BGP information. To display VPNv4 BGP information, use one of these keywords:

- **all** to display the whole contents of the VPNv4 BGP table
- **vrf** *vrf-name* to display VPNv4 information associated with the specified VRF
- **rd** *route-distinguisher* to display VPNv4 information associated with the specified RD

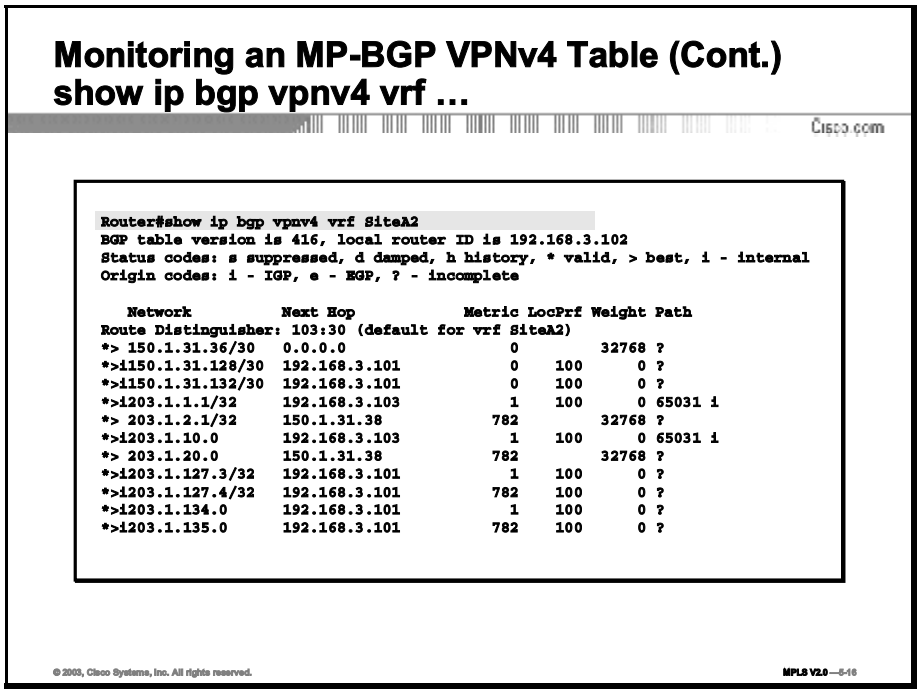

show ip bgp vpnv4 vrf

To display VPNv4 information from the BGP database associated with a VRF, use the **show ip bgp vpnv4 vrf** command in privileged EXEC mode:

- **show ip bgp vpnv4 vrf** *vrf-name* [*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]] [*network-address* [*mask*] [**longer-prefixes**] [*output-modifiers*]] [**cidr-only**] [**community**][**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**tags**]

Syntax Description

Parameter	Description
vrf *vrf-name*	Displays NLRI prefixes associated with the named VRF.

Defaults

This command has no default values.

Usage Guidelines

Use this command to display VPNv4 information that is associated with a VRF from the BGP database. A similar command—**show ip bgp vpnv4 all**—displays all available VPNv4 information. The **show ip bgp vpnv4 summary** command displays BGP neighbor status.

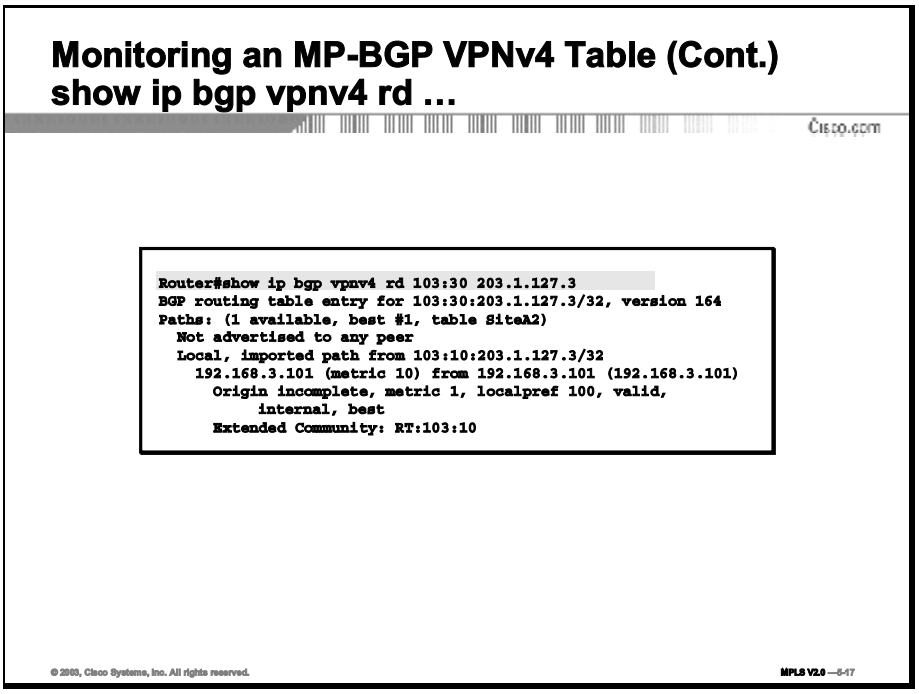

show ip bgp vpnv4 rd *route-distinguisher*

To display all VPNv4 routes that contain a specified RD, use the **show ip bgp vpnv4 rd** command in privileged EXEC mode:

- **show ip bgp vpnv4 rd** *route-distinguisher* [*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]] [*network-address* [*mask*] [**longer-prefixes**] [*output-modifiers*]] [**cidr-only**] [**community**][**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**paths** [*line*]] [**quote-regexp**] [**regexp**] [**summary**] [**tags**]

Syntax Description

Parameter	Description
rd *route-distinguisher*	Displays NLRI prefixes that have a matching RD.

Defaults

There is no default. A route distinguisher (RD) must be configured for a VRF to be functional.

Usage Guidelines

A RD creates routing and forwarding tables and specifies the default route distinguisher for a VPN. The RD is added to the beginning of the customer's IPv4 prefixes to change them into globally unique VPN-IPv4 prefixes.

Either RD is an ASN-relative RD, in which case it is composed of an autonomous system number and an arbitrary number, or it is an IP-address-relative RD, in which case it is composed of an IP address and an arbitrary number.

You can enter an RD in either of these formats:

16-bit AS number: your 32-bit number
For example, 101:3.

32-bit IP address: your 16-bit number
For example, 192.168.122.15:1.

Examples

The following example shows how to configure a default RD for two VRFs. It illustrates the use of both AS-relative and IP-address-relative RDs:

```
Router(config)# ip vrf vrf_blue
Router(config-vrf)# rd 100:3
Router (config-vrf)# exit
Router(config)# ip vrf vrf_red
Router(config-vrf)# rd 173.13.0.12:200
```

Monitoring Per-VRF CEF and LFIB Structures

This topic identifies the command syntax that is used to monitor per-VRF CEF and label forwarding information base (LFIB) structures.

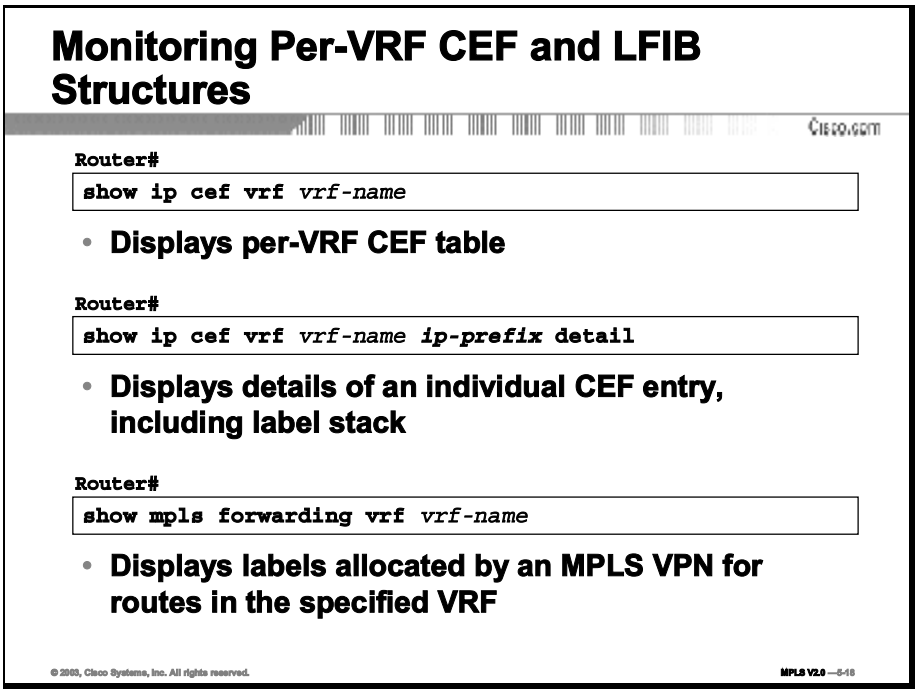

Three commands can be used to display per-VRF FIB and LFIB structures:

- The **show ip cef vrf** command displays the VRF FIB.
- The **show ip cef vrf detail** command displays detailed information about a single entry in the VRF FIB.
- The **show mpls forwarding vrf** command displays all labels allocated to VPN routes in the specified VRF.

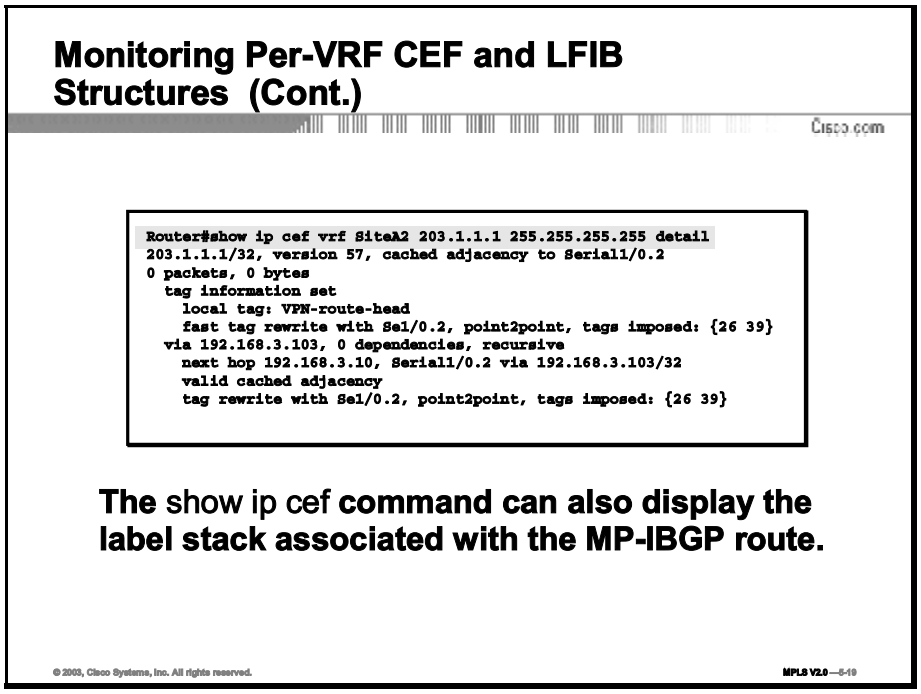

show ip cef vrf

To display the CEF forwarding table associated with a VRF, use the **show ip cef vrf** command in privileged EXEC mode:

- **show ip cef vrf** *vrf-name* [*ip-prefix* [*mask* [**longer-prefixes**]] [**detail**] [*output-modifiers*]] [*interface interface-number*] [**adjacency** [*interface interface-number*] [**detail**] [**discard**] [**drop**] [**glean**] [**null**] [**punt**] [*output-modifiers*]] [**detail** [*output-modifiers*]] [**non-recursive** [**detail**] [*output-modifiers*]] [**summary** [*output-modifiers*]] [**traffic** [**prefix-length**] [*output-modifiers*]] [**unresolved** [**detail**] [*output-modifiers*]]

Syntax Description

Parameter	Description
vrf-name	Name assigned to the VRF.
ip-prefix	(Optional) IP prefix of entries to show, in dotted decimal notation (A.B.C.D).
mask	(Optional) Mask of the IP prefix, in dotted decimal notation.
longer-prefixes	(Optional) Displays table entries for all the more specific routes.
detail	(Optional) Displays detailed information for each CEF table entry.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
interface	(Optional) Type of network interface to use: ATM, Ethernet, loopback, POS, or null.
interface-number	Number identifying the network interface to use.
adjacency	(Optional) Displays all prefixes resolving through adjacency.
discard	Discards adjacency.
drop	Drops adjacency.
glean	Gleans adjacency.
null	Null adjacency.
punt	Punts adjacency.
non-recursive	(Optional) Displays only nonrecursive routes.
summary	(Optional) Displays a CEF table summary.
traffic	(Optional) Displays traffic statistics.
prefix-length	(Optional) Displays traffic statistics by prefix size.
unresolved	(Optional) Displays only unresolved routes.

Defaults

This command has no default values.

Usage Guidelines

Used with the *vrf-name* argument, the **show ip cef vrf** command shows a shortened display of the CEF table.

Used with the **detail** keyword, the **show ip cef vrf** command shows detailed information for all CEF table entries.

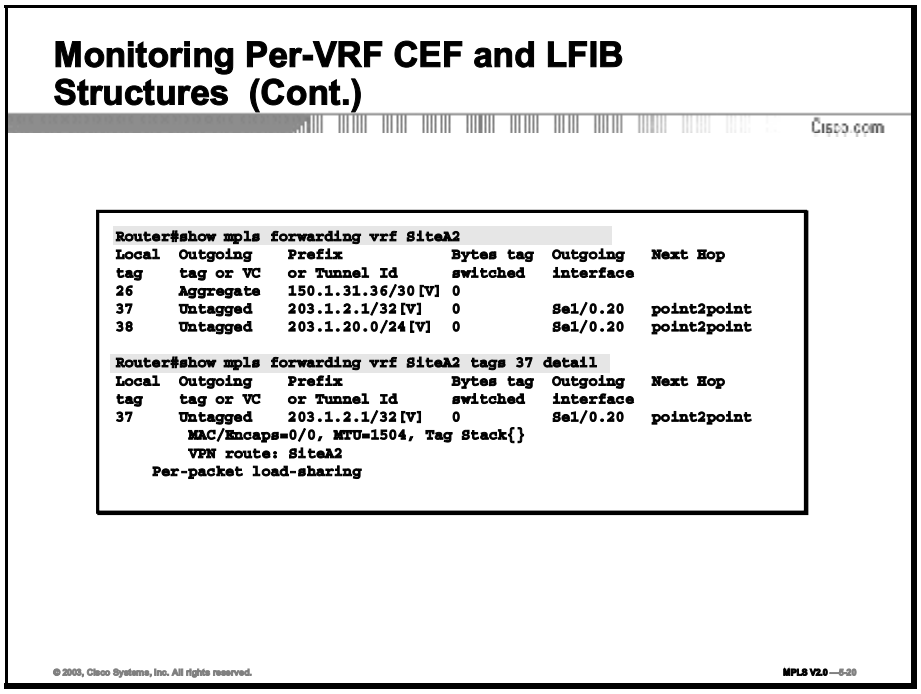

show mpls forwarding vrf

To display label-forwarding information for advertised VRF routes, use the **show mpls forwarding vrf** command in EXEC mode. To disable the display of label-forwarding information, use the **no** form of this command:

- **show mpls forwarding vrf** *vrf-name* [*ip-prefix/length* [*mask*]] [**detail**] [*output-modifiers*]
- **no show mpls forwarding vrf** *vrf-name* [*ip-prefix/length* [*mask*]] [*detail*] [*output-modifiers*]

Syntax Description

Parameter	Description
vrf-name	Displays NLRI prefixes associated with the named VRF.
ip-prefix/length	(Optional) IP prefix address (in dotted decimal notation) and length of mask (0 to 32).
mask	(Optional) Destination network mask in dotted decimal notation.
detail	(Optional) Displays detailed information on the VRF routes.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Defaults

This command has no default behavior or values.

Usage Guidelines

Use this command to display label-forwarding entries associated with a particular VRF or IP prefix.

Monitoring Labels Associated with VPNv4 Routes

This topic identifies the command syntax that is used to monitor labels associated with VPNv4 routes.

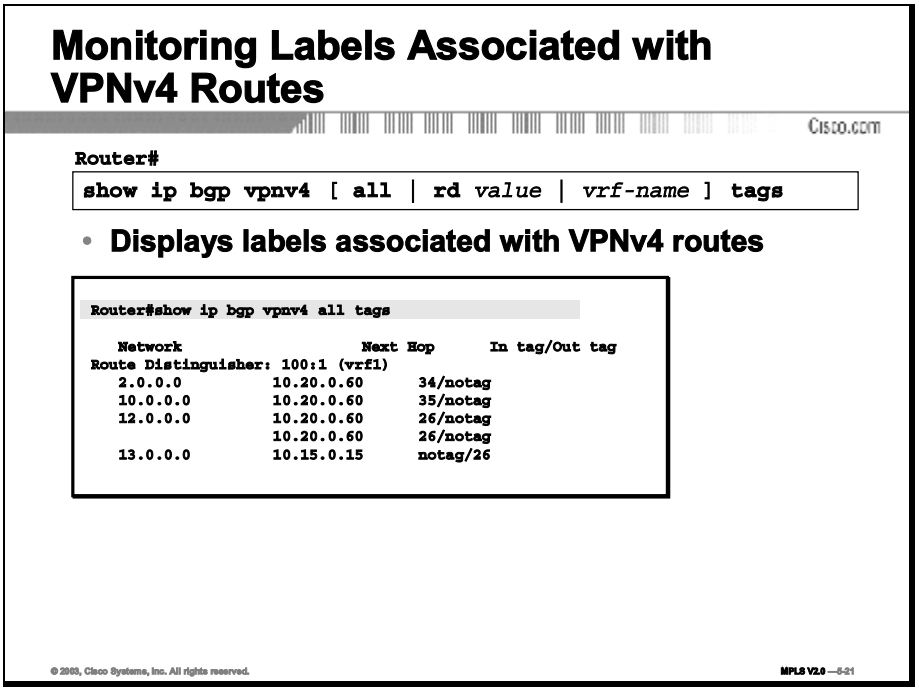

You can use the **show ip bgp vpnv4 tags** command to display tags assigned to local or remote VRF routes by the local or remote PE router. This command displays tags associated with all VPNv4 routes when you use the **all** keyword. It can also display tags associated with a specified RD or VRF.

The following table describes the fields displayed by this command.

Field	Description
`Network`	Displays the network address from the BGP table.
`Next Hop`	Specifies the BGP next-hop address.
`In tag`	Displays the label (if any) assigned by this router.
`Out tag`	Displays the label assigned by the BGP next-hop router.

Other MPLS VPN Monitoring Commands

This topic identifies the command syntax that is used with other MPLS VPN monitoring commands.

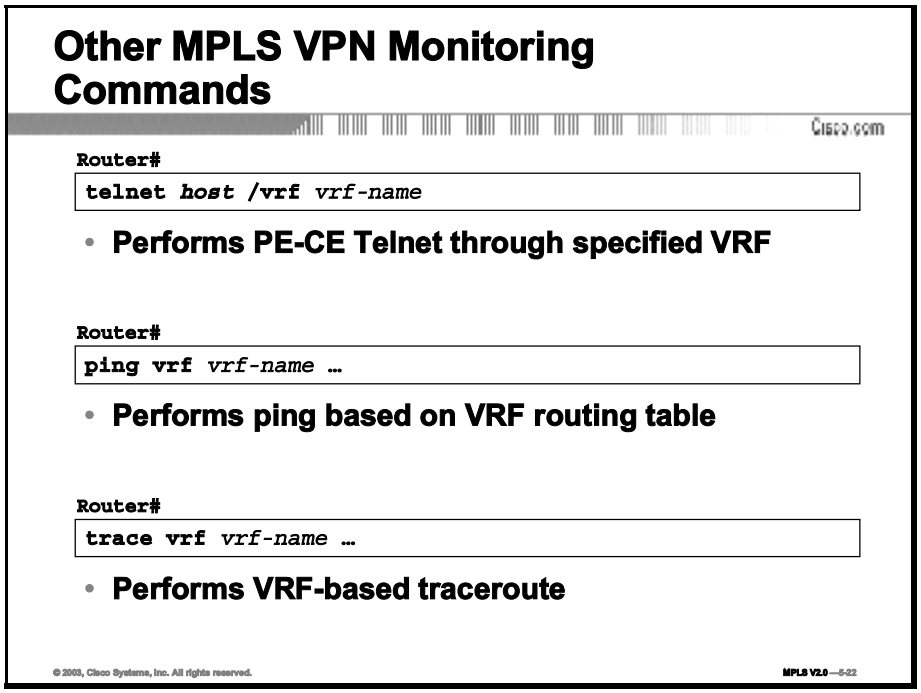

Three additional Cisco IOS software monitoring commands are VRF-aware:

- The **telnet** command can be used to connect to a CE router from a PE router using the **/vrf** option.
- The **ping vrf** command can be used to ping a destination host reachable through a VRF.
- The **trace vrf** command can be used to trace a path toward a destination reachable through a VRF.

Summary

This topic summarizes the key points discussed in this lesson.

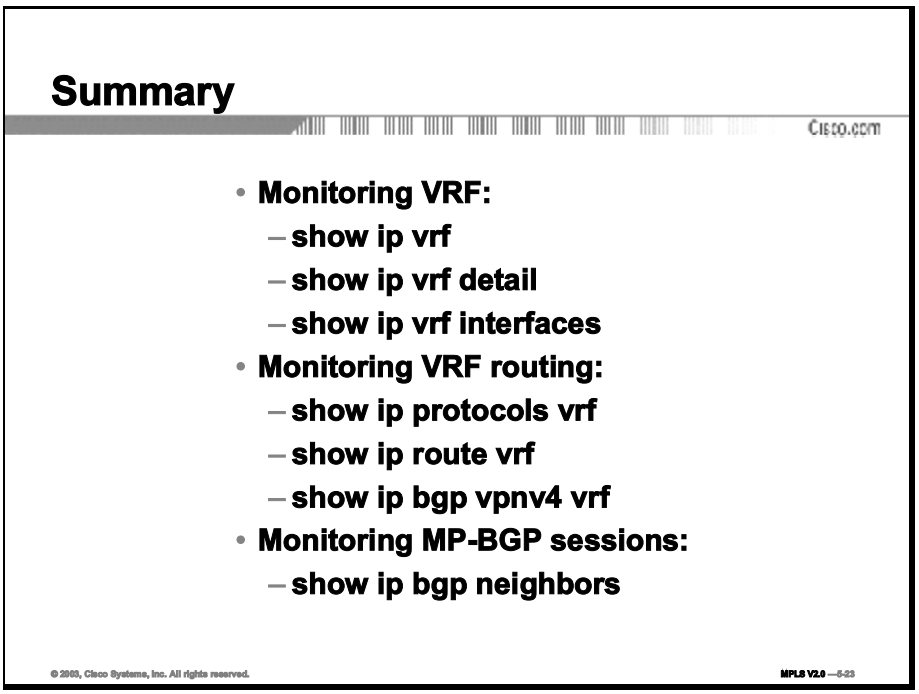

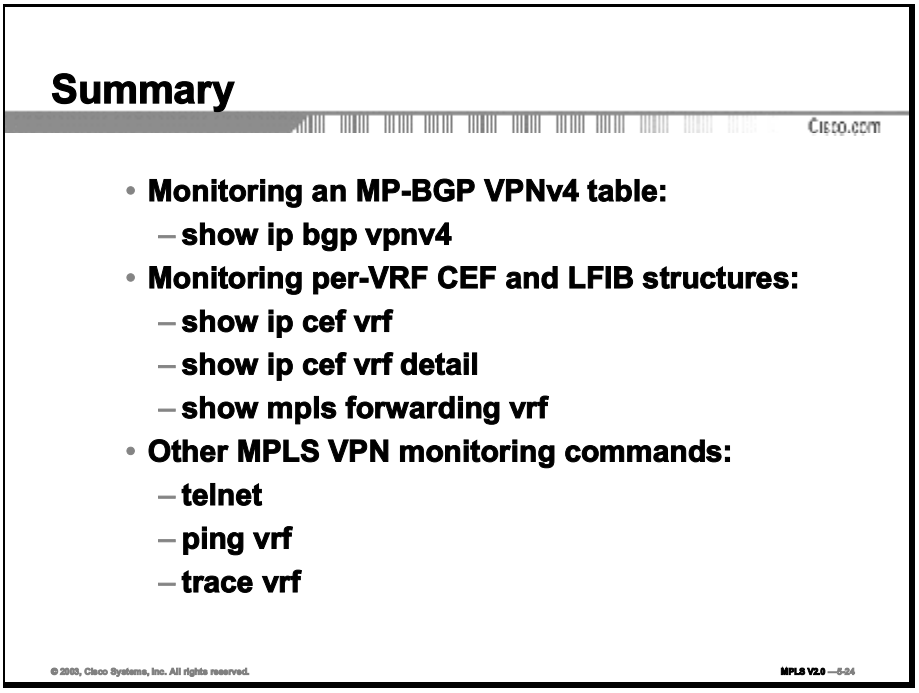

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three commands do you use to display all configured VRFs on the router?

Q2) How do you verify the contents of a VRF routing table?

Q3) Why is the BGP protocol always running in every VRF and how would you display the BGP parameter related to a VRF?

Q4) How do you verify that a session has been established between two VPNv4 neighbors?

Q5) How do you verify the contents of a BGP VPNv4 routing table?

Q6) Which three commands can be used to display per-VRF FIB and LFIB information?

Q7) Which command can be used to display tags assigned to local or remote VRF routes by the local or remote PE router?

_____

Q8) Which command do you use to perform the following traceroutes?

Ingress CE to egress PE— _____

Ingress CE to egress CE— _____

Ingress PE to egress PE— _____

Ingress PE to egress CE— _____

Ingress P to egress PE— _____

Ingress P to egress CE— _____

Quiz Answer Key

Q1) **show ip vrf**
show ip vrf detail
show ip vrf interfaces

Relates to: Monitoring VRFs

Q2) Use the **show ip route vrf** command.

Relates to: Monitoring VRF Routing

Q3) It is needed to carry the VPNv4 routes. Use the **show ip bgp vpnv4 vrf** command.

Relates to: Monitoring MP-BGP Sessions

Q4) Use the **show ip bgp neighbors** command and verify that the VPNv4 status is "advertised and received."

Relates to: Monitoring MP-BGP Sessions

Q5) Use the **show ip bgp vpnv4 vrf** *vrf-name* command.

Relates to: Monitoring an MP-BGP VPNv4 Table

Q6) **show ip cef vrf**
show ip cef vrf detail
show mpls forwarding vrf

Relates to: Monitoring Per-VRF CEF and LFIB Structures

Q7) **show ip bgp vpnv4 tags**

Relates to: Monitoring Labels Associated with VPNv4 Routes

Q8) Ingress CE to egress PE—trace
Ingress CE to egress CE—trace
Ingress PE to egress PE—trace
Ingress PE to egress CE—trace vrf *vrf-name*
Ingress P to egress PE—trace
Ingress P to egress CE—You cannot do a traceroute from a P router to any CE router. The P router does not have the CE routing information in its routing table.

Relates to: Other MPLS VPN Monitoring Commands

Configuring OSPF as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the issues encountered and the PE-CE routing protocol configuration steps required when you are running OSPF between PE and CE routers.

Relevance

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need in order to configure MPLS VPN PE-CE routing exchange.

Objectives

The lesson identifies the command syntax that is required to configure PE-CE routing protocols.

Upon completing this lesson, you will be able to:

- Describe the enhanced OSPF hierarchical model
- Describe the propagation of OSPF customer routes across the MPLS VPN backbone and the issues that might be encountered
- Describe how an MPLS VPN is implemented as an OSPF superbackbone
- Identify the command syntax that is used to configure a PE-CE OSPF routing session
- Describe the route loop issue that may result from the redistribution of routes between OSPF and BGP in an MPLS VPN environment and how the OSPF down bit is used to address this issue
- Describe how optimizing of packet forwarding is accomplished across the MPLS VPN backbone

- Describe the route loop issue that may result from advertising routes across multiple OSPF domains in an MPLS VPN environment and how the OSPF tag field is used to address this issue
- Describe the function of a sham link and how it is implemented
- Identify the command syntax that is used to configure a sham link

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- OSPF Hierarchical Model
- OSPF in an MPLS VPN Routing Model
- OSPF Superbackbone
- Configuring OSPF PE-CE Routing
- OSPF Down Bit
- Optimizing of Packet Forwarding Across the MPLS VPN Backbone
- OSPF Tag Field
- Sham Link
- Configuring a Sham Link
- Summary
- Quiz

OSPF Hierarchical Model

This topic describes the enhanced OSPF hierarchical model.

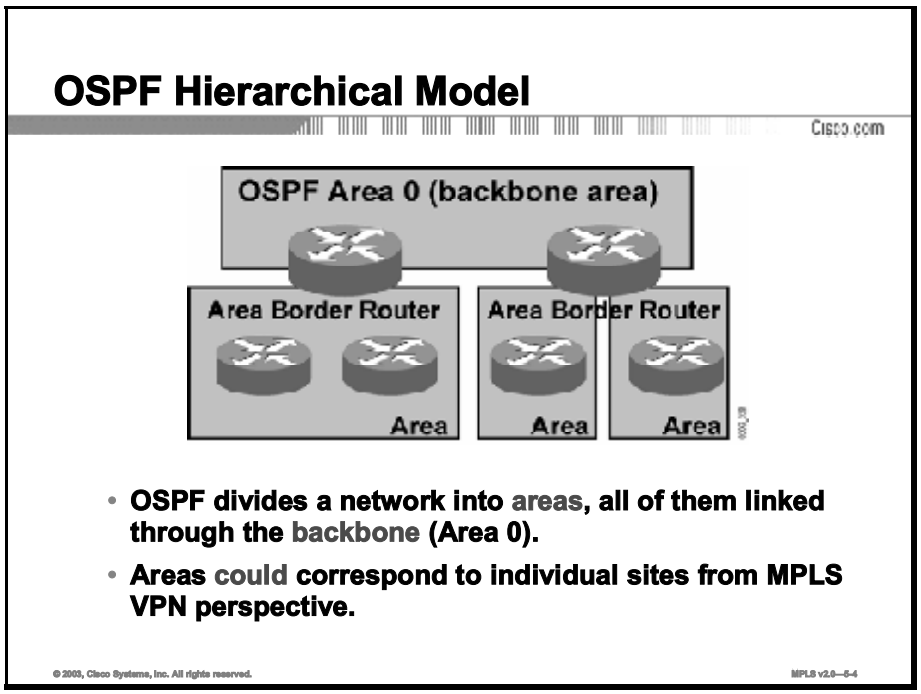

The Open Shortest Path First (OSPF) routing protocol was designed to support hierarchical networks with a central backbone. The network running OSPF is divided into "areas." All areas have to be directly connected to the "backbone area" (Area 0). The whole OSPF network (backbone area and any other areas connected to it) is called the OSPF domain.

The OSPF areas in the customer network could correspond to individual sites, but there are also other options that are often encountered:

- A single area could span multiple sites (for example, the customer decides to use an area per region, but the region contains multiple sites)
- The backbone area could be extended into individual sites

Note Please refer to the *Building Scalable Cisco Networks* (BSCN) course or OSPF curriculum for background information on OSPF.

OSPF in an MPLS VPN Routing Model

This topic describes the propagation of OSPF customer routes across the MPLS VPN backbone and the issues that might be encountered.

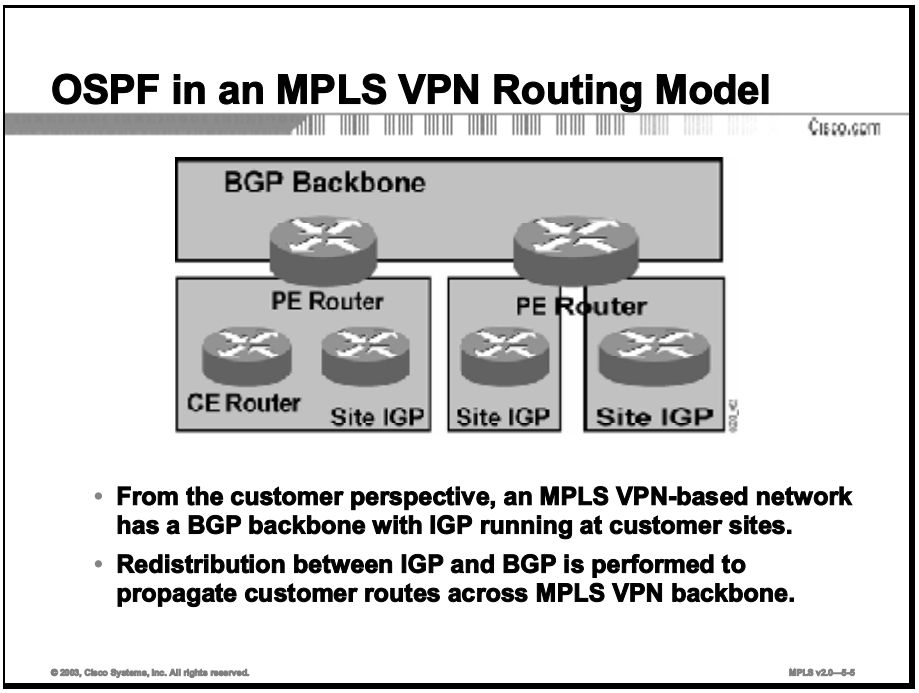

The MPLS VPN routing model introduces a BGP backbone into the customer network. Isolated copies of IGP run at every site, and MP-BGP is used to propagate routes between sites. Redistribution between customer IGP—running between PE routers and CE routers—and the backbone MP-BGP, is performed at every PE router.

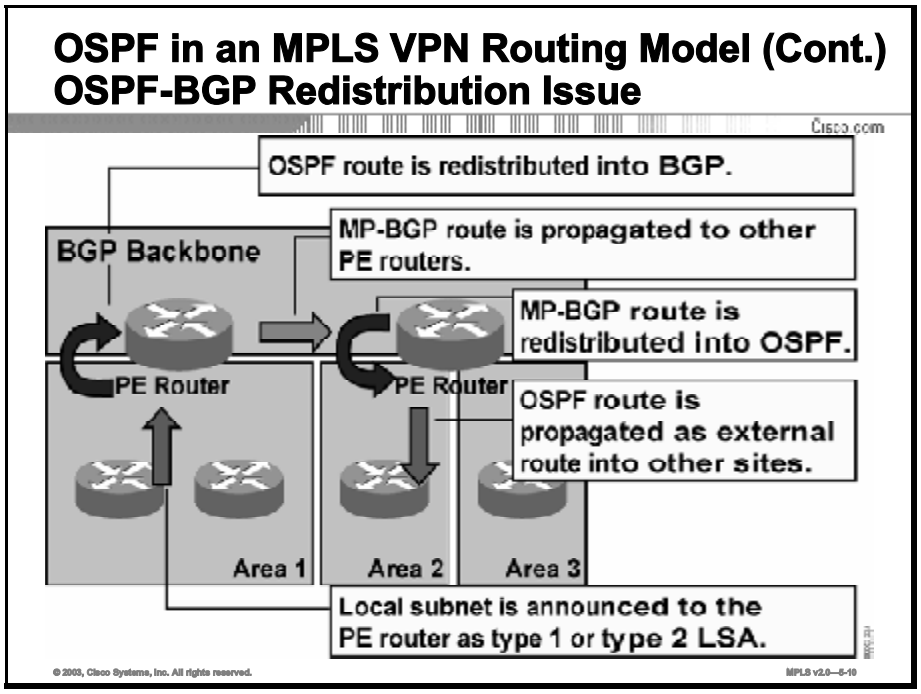

The IGP-BGP redistribution introduced by the MPLS VPN routing model does not fit well into customer networks running OSPF. When an OSPF customer is migrated to an MPLS VPN service, any route that is redistributed into OSPF from another routing protocol will now be redistributed as an *external* OSPF route. The OSPF routes received by one PE router are propagated across the MPLS backbone and redistributed back into OSPF at another site as external OSPF routes.

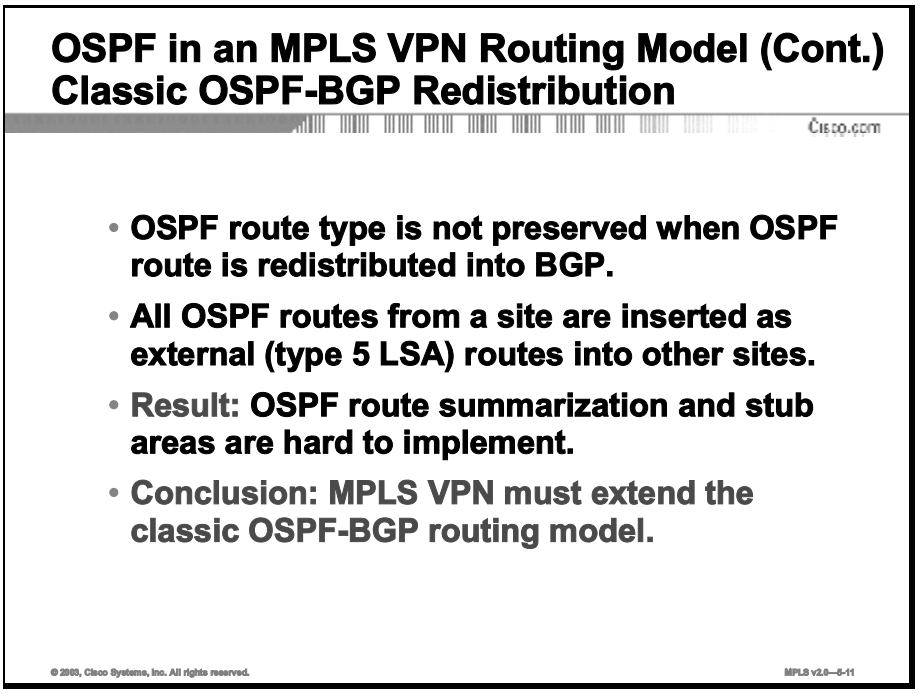

With the traditional OSPF-BGP redistribution, the OSPF route type (internal or external route) is not preserved when the OSPF route is redistributed into BGP. When that same route is redistributed back into OSPF, it is always redistributed as an external OSPF route.

There are a number of caveats associated with external OSPF routes:

- External routes cannot be summarized.
- External routes are flooded across all OSPF areas.
- External routes could use a different metric type that is not comparable to OSPF cost.
- External routes are not inserted in stub areas or not-so-stubby areas (NSSAs).
- Internal routes are always preferred over external routes, regardless of their cost.

Because of all these caveats, migrating an OSPF customer toward an MPLS VPN service might severely impact the routing of that customer. The MPLS VPN architecture must therefore extend the classic OSPF-BGP routing model to support transparent customer migration.

OSPF Superbackbone

This topic describes how an MPLS VPN is implemented as an OSPF superbackbone.

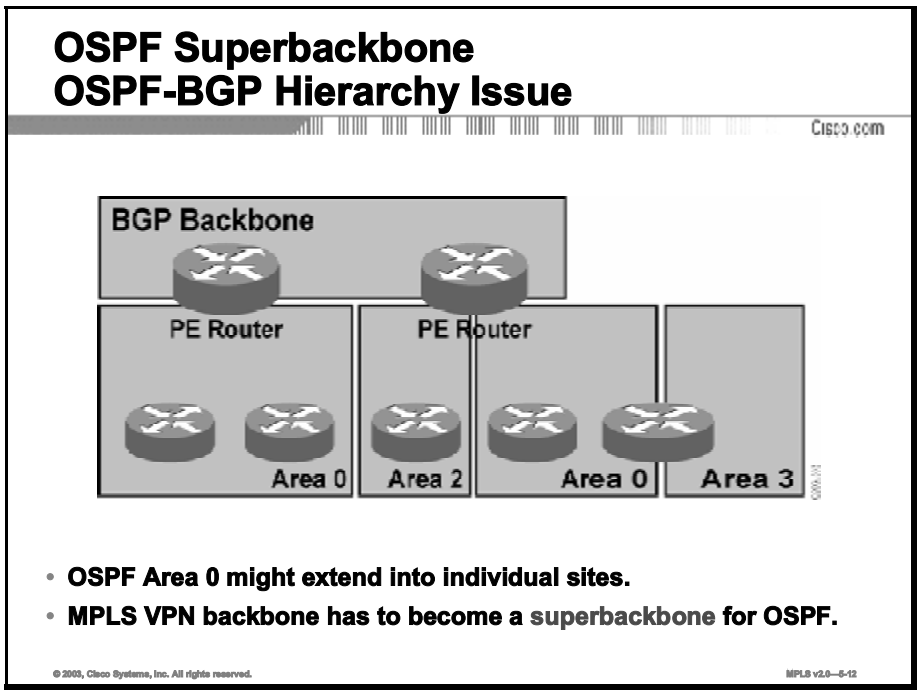

The MPLS VPN architecture extends the OSPF architecture by introducing another backbone above OSPF Area 0, the superbackbone. The OSPF superbackbone is implemented with MP-BGP between the PE routers but is otherwise completely transparent to the OSPF routers. The architecture even allows disjoint OSPF backbone areas (Area 0) at MPLS VPN customer sites.

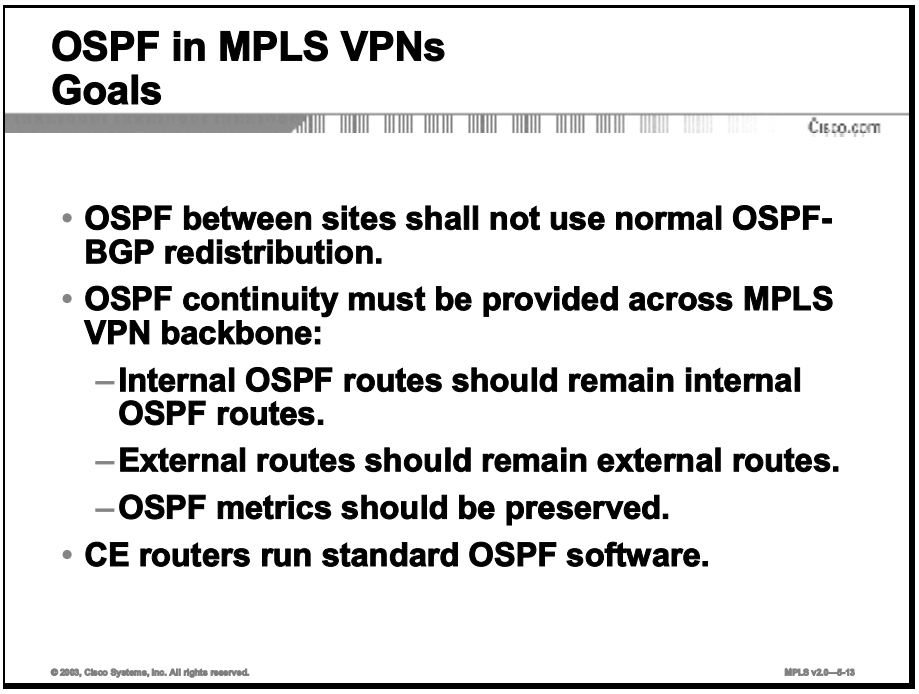

The goals that have to be met by the OSPF superbackbone are as follows:

- The superbackbone shall not use standard OSPF-BGP redistribution.
- OSPF continuity must be provided between OSPF sites:
  - Internal OSPF routes must remain internal OSPF routes.
  - External OSPF routes must remain external OSPF routes.
  - Non-OSPF routes redistributed into OSPF must appear as external OSPF routes in OSPF.
  - OSPF metrics and metric types (external 1 or external 2) have to be preserved.
- The OSPF superbackbone shall be transparent to the CE routers that run standard OSPF software.

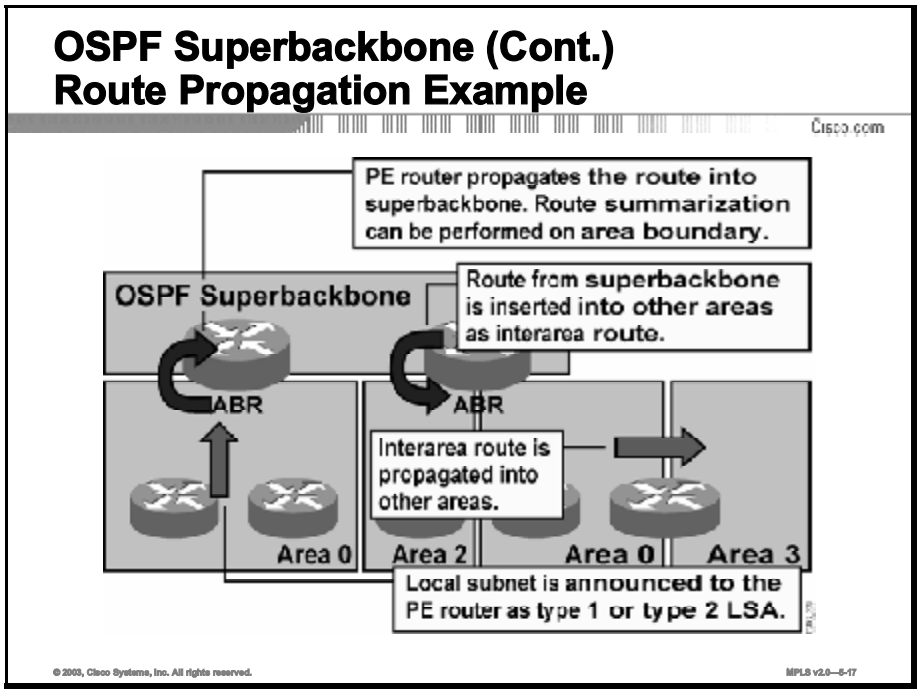

The MPLS VPN superbackbone appears as another layer of hierarchy in the OSPF architecture. The PE routers that connect regular OSPF areas to the superbackbone therefore appear as OSPF Area Border Routers (ABRs) in the OSPF areas to which they are attached. In Cisco IOS implementation, they also appear as Autonomous System Boundary Routers (ASBRs) in nonstub areas.

From the perspective of a standard OSPF-speaking CE router, the PE routers insert interarea routes from other areas into the area in which the CE router is present. The CE routers are not aware of the superbackbone or of other OSPF areas present beyond the MPLS VPN superbackbone.

With the OSPF superbackbone architecture, the continuity of OSPF routing is preserved:

- The OSPF intra-area route (described in the OSPF router link-state advertisement [LSA] or network LSA) is inserted into the OSPF superbackbone by redistributing the OSPF route into MP-BGP. Route summarization can be performed on the redistribution boundary by the PE router.
- The MP-BGP route is propagated to other PE routers and inserted as an OSPF route into other OSPF areas. Because the superbackbone appears as another area behind the PE router (acting as ABR), the MP-BGP route derived from the intra-area route is always inserted as an interarea route. The interarea route can then be propagated into other OSPF areas by ABRs within the customer site.

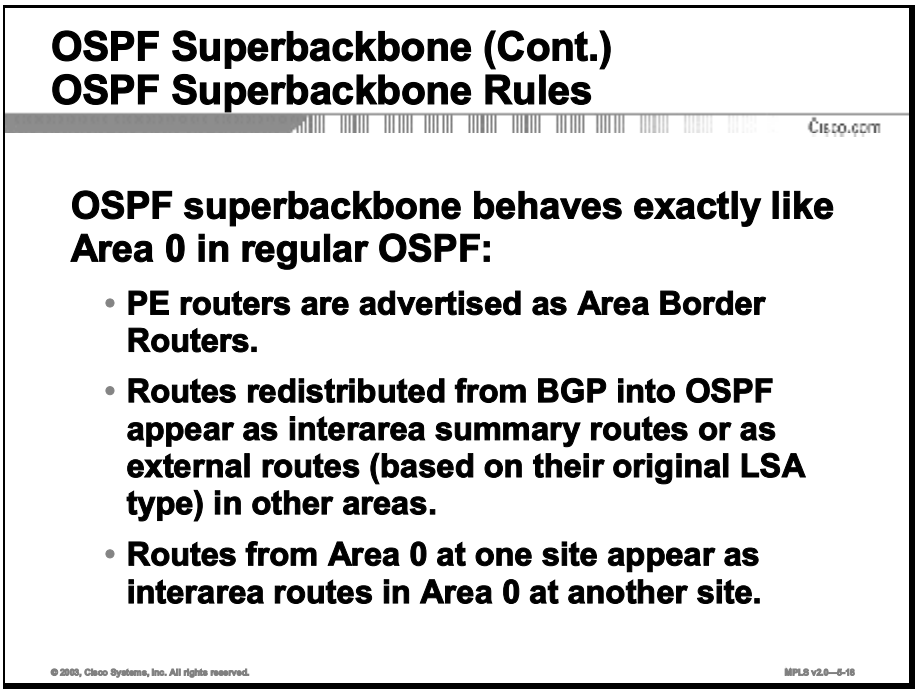

The OSPF superbackbone rules could be summarized as follows:

- PE routers advertise themselves as ABRs. The superbackbone appears as another area to the CE routers.
- Routes redistributed into MP-BGP from OSPF will appear as interarea routes in other OSPF sites if the original route was an intra-area or interarea route and as external routes if the original route was an external route.

As a consequence of the second rule, routes from the backbone area at one site appear as interarea routes (not as backbone routes) in backbone areas at other sites.

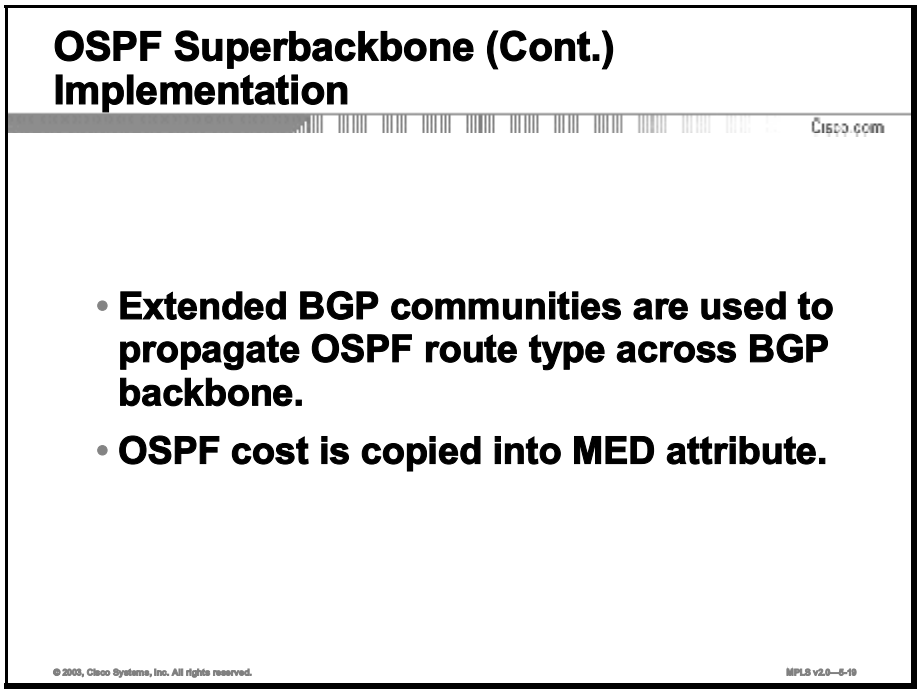

The OSPF superbackbone is implemented with the help of several BGP attributes:

A new BGP extended community was defined to carry OSPF route type and OSPF area across the BGP backbone. The format of this community is defined in the following table.

Field	Number of Bytes	Comments
Community type	2	The community type is 0x8000.
OSPF area	4	This field carries the OSPF area from which the route was redistributed into MP-BGP.
LSA type	1	This field carries the OSPF LSA type from which the route was redistributed into MP-BGP.
Option	1	This field is used for external metric type. The low-order bit is set for external 2 routes.

Note The option field in the OSPF route type extended community is not equivalent to the option field in the OSPF LSA.

As in the standard OSPF-BGP redistribution, the OSPF cost is carried in the MED attribute.

OSPF Superbackbone (Cont.) Implementation

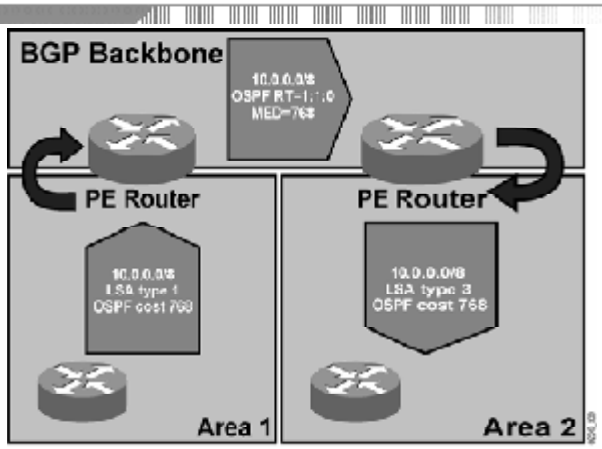

- **OSPF route type is copied into extended BGP community on redistribution into BGP.**
- **Egress PE router performs interarea transformation.**

This figure illustrates the propagation of internal OSPF routes across the MPLS VPN superbackbone. The sending PE router redistributes the OSPF route into MP-BGP, copies OSPF cost into the MED attribute, and sets the BGP extended community to indicate the LSA type from which the route was derived.

The receiving PE router redistributes the MP-BGP route back into OSPF and uses the original LSA type and the MED attribute to generate an interarea summary LSA. An interarea summary LSA is always generated, because the receiving PE router acts as an ABR between the superbackbone and the OSPF area(s).

OSPF Superbackbone (Cont.) External Routes

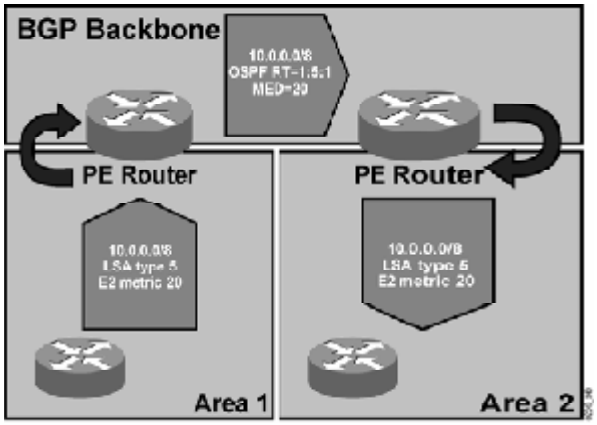

- **External OSPF routes are propagated in the same way as internal OSPF routes across superbackbone.**
- **External metric and route type are preserved.**

OSPF Superbackbone (Cont.) External Routes

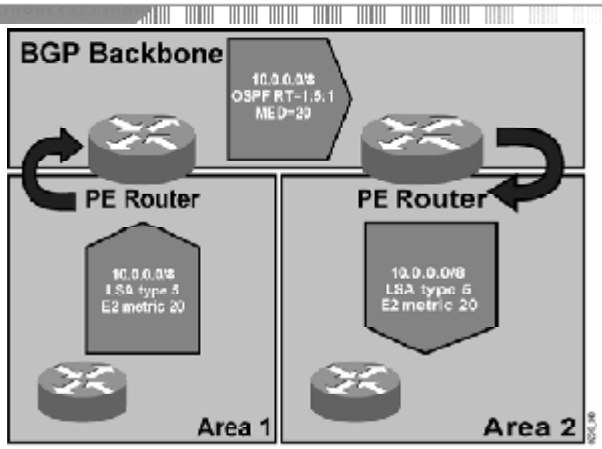

- **External OSPF routes are propagated in the same way as internal OSPF routes across superbackbone.**
- **External metric and route type are preserved.**

The external OSPF routes are redistributed into the MP-BGP in exactly the same way as the internal OSPF routes. The process changes slightly on the receiving PE router:

- For external routes (type 5 LSA), the LSA is reoriginated, with the receiving PE router being the ASBR. The external metric type is copied from the BGP extended community, and the external cost is copied from the MED.
- For NSSA external routes (type 7 LSA), the route is announced to the other OSPF sites as a type 5 LSA external route, because the route has already crossed the area boundary.

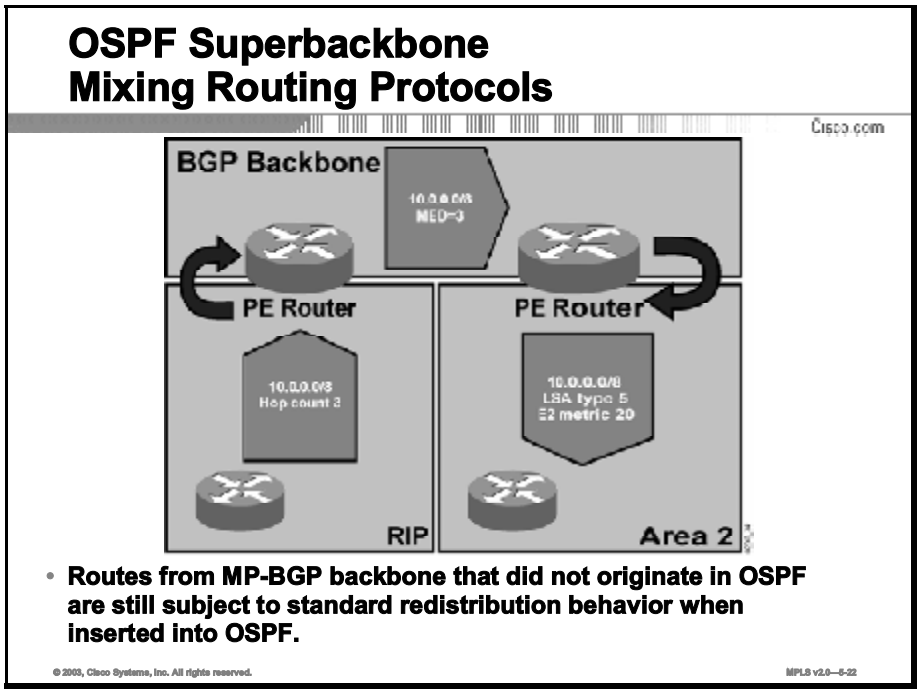

The MPLS VPN superbackbone still retains the traditional OSPF-BGP route redistribution behavior for routes that did not originate in OSPF at other sites (and therefore do not carry the OSPF extended BGP community). These routes are inserted into the OSPF topology database as type 5 external routes (or type 7 external routes for NSSA areas), with the default OSPF metric (not the value of MED).

Configuring OSPF PE-CE Routing

This topic identifies the command syntax that is used to configure a PE-CE OSPF routing session.

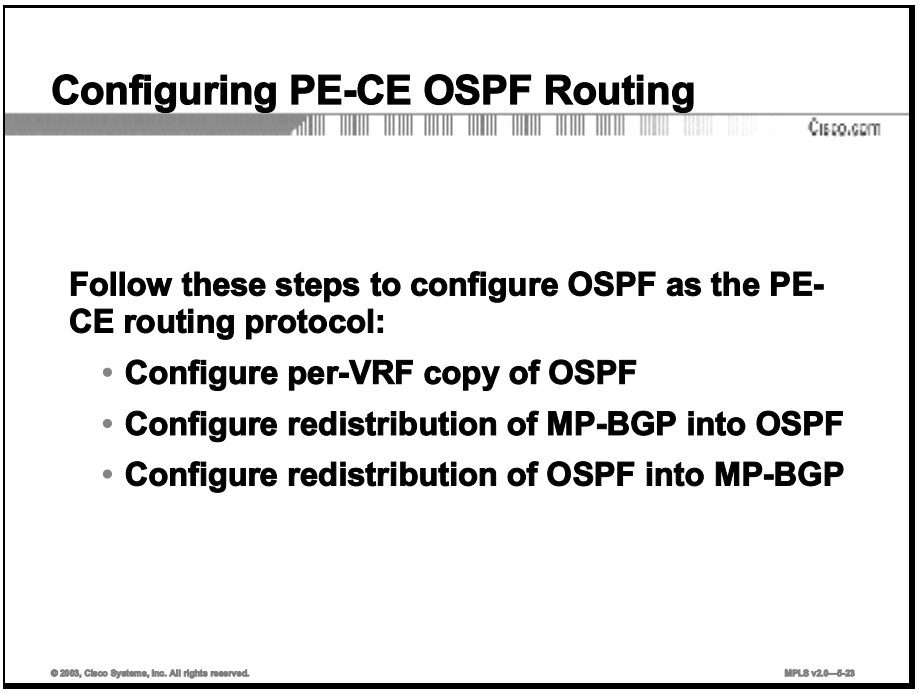

To configure OSPF as a PE-CE routing protocol, you need to start a separate OSPF process for each VRF in which you want to run OSPF. The per-VRF OSPF process is configured in the same way as a standard OSPF process. You can use all OSPF features available in Cisco IOS software.

You need to redistribute OSPF routes into BGP and redistribute BGP routes into OSPF if necessary. Alternatively, you can originate a default route into a per-VRF OSPF process by using the **default-information originate always** command in router configuration mode.

MP-BGP propagates more than just OSPF cost across the MPLS VPN backbone. The propagation of additional OSPF attributes into MP-BGP is automatic and requires no extra configuration.

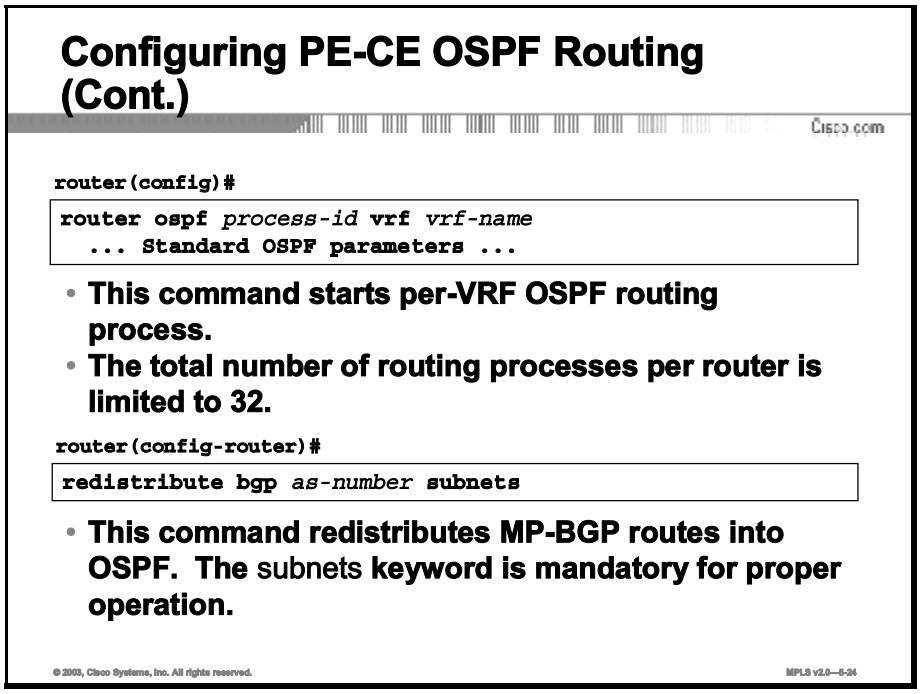

OSPF is the only PE-CE routing protocol that is not fully VPN-aware. A separate OSPF process is run for every VRF.

router ospf

To configure an OSPF routing process within a VRF, use the **router ospf** command in global configuration mode. To terminate an OSPF routing process, use the **no** form of this command:

- **router ospf** *process-id* **vrf** *vrf-name*
- **no router ospf** *process-id* **vrf** *vrf-name*

Syntax Description

Parameter	Description
process-id	Internally used identification parameter for an OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process.
vrf-name	The name of the VRF where the OSPF process will reside.

Defaults

No OSPF routing process is defined.

Configuring PE-CE OSPF Routing (Cont.)

Cisco.com

```
router(config)#
router bgp as-number
  address-family ipv4 vrf vrf-name
    redistribute ospf process-id [match [internal]
[external-1] [external-2]]
```

- **OSPF-BGP route redistribution is configured with the** redistribute **command under the proper** address-family.
- **Without the OSPF** match **keyword specified, only internal OSPF routes are redistributed into OSPF.**

Use the standard BGP redistribution commands.

OSPF Down Bit

This topic describes the route loop issue that may result from the redistribution of routes between OSPF and BGP in an MPLS VPN environment and how the OSPF down bit is used to address this issue.

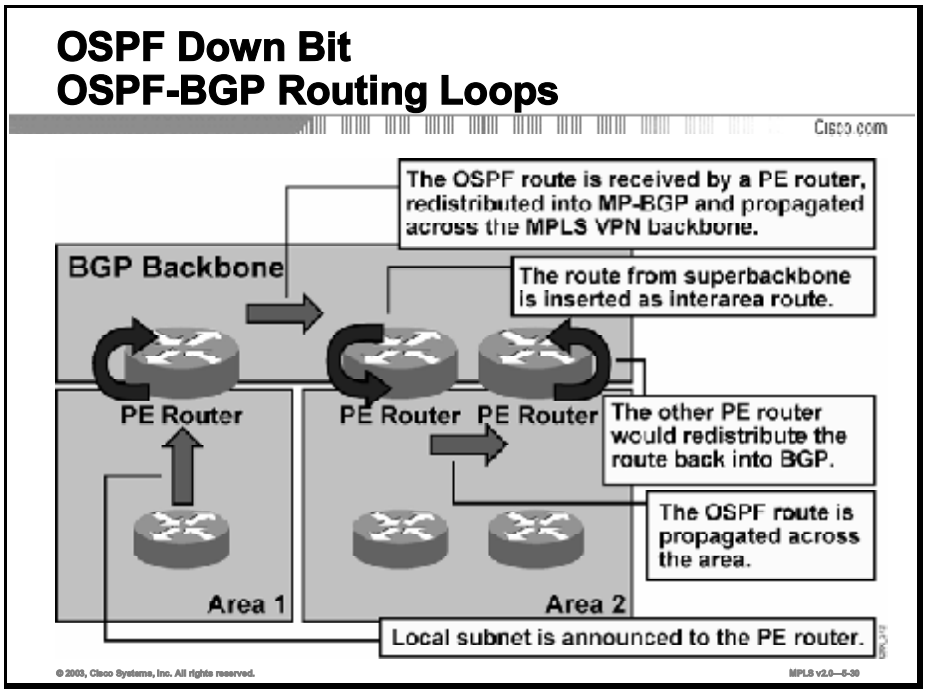

OSPF developers took many precautions to avoid routing loops between OSPF areas—for example, intra-area routes are always preferred over interarea routes. These rules do not work when the superbackbone is introduced. Consider, for example, a network in the figure here, where the receiving OSPF area has two PE routers attached to it.

The following table indicates the process steps that could produce a routing loop.

Step	Action
1	The sending PE router receives an intra-area OSPF route.
2	The intra-area OSPF route is redistributed into MP-BGP. An OSPF community is attached to the route to indicate that it was an OSPF route before being redistributed.
3	The receiving PE router redistributes the MP-BGP route into OSPF as an internal interarea summary route.
4	The summary route is propagated across the OSPF area and received by the other PE router attached to the same area.
5	The administrative distance of the OSPF route is better than the administrative distance of the MP-BGP route; therefore, the PE router selects the OSPF route and redistributes the route back into the MP-BGP process, potentially resulting in a routing loop.

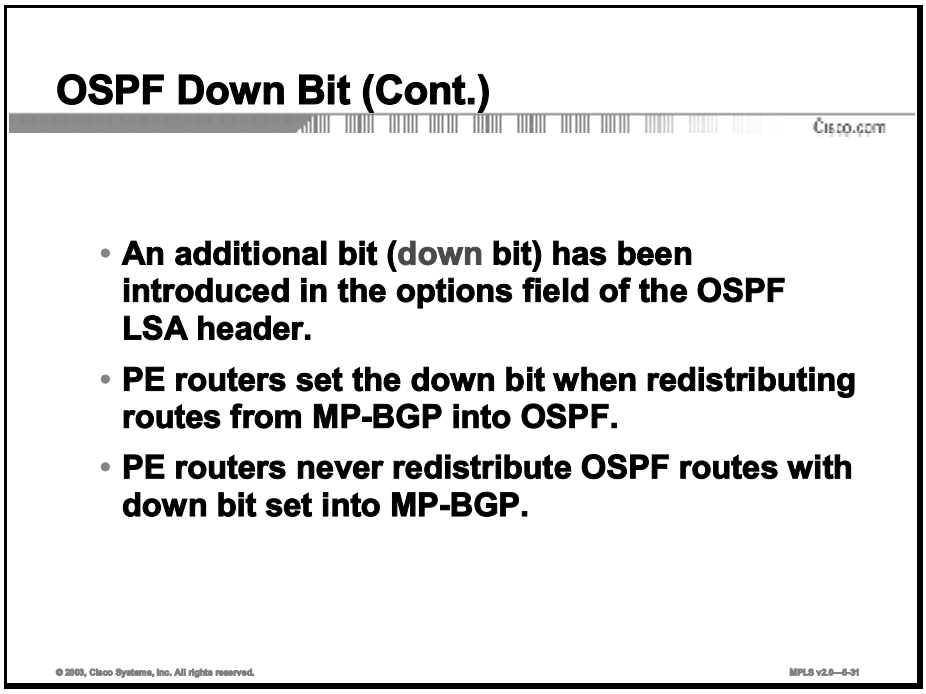

Two mechanisms were introduced to prevent route redistribution loops between OSPF (running between PE and CE routers) and MP-BGP running between PE routers:

- BGP Site of Origin (SOO), which is covered in the "MPLS VPN Routing Model" lesson of the "MPLS Virtual Private Networks Technology" module and detailed further in the "Configuring BGP as the Routing Protocol Between PE and CE Routers" lesson of the "MPLS VPN Implementation" module
- The down bit in the options field of the OSPF LSA header

The down bit is used between the PE routers to indicate which routes were inserted into the OSPF topology database from the MPLS VPN superbackbone and thus shall not be redistributed back in the MPLS VPN superbackbone. The PE router that redistributes the MP-BGP route as OSPF route into the OSPF topology database sets the down bit. Other PE routers use the down bit to prevent this route from being redistributed back into MP-BGP.

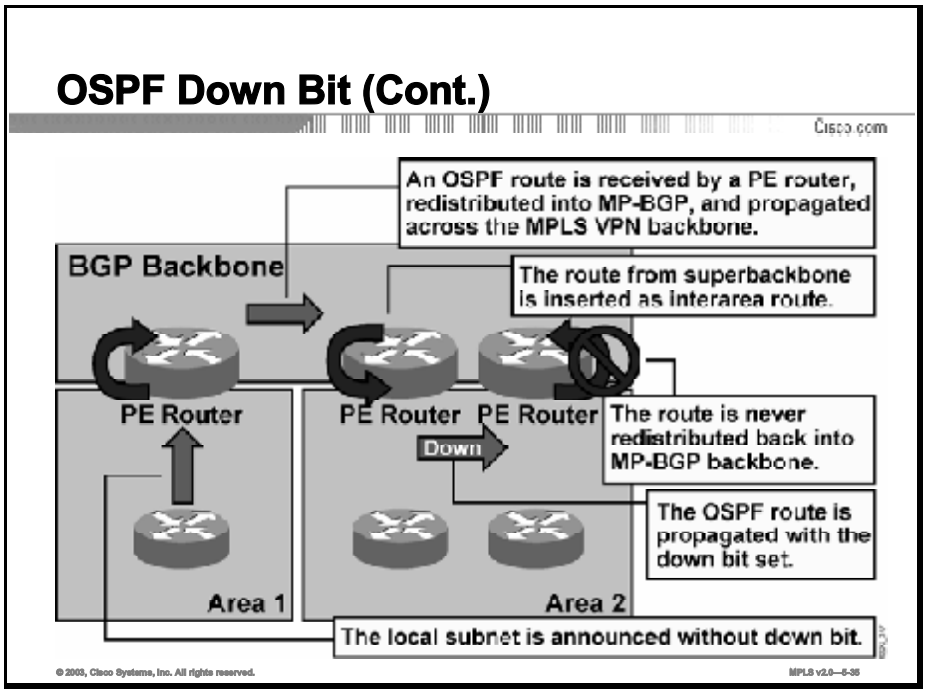

The typical usage of the down bit is shown in the figure, and the process steps that show how it prevents routing loops is detailed in the following table.

Step	Action
1	PE router receives an OSPF route.
2	PE router redistributes OSPF route into MP-BGP. The MP-BGP route is propagated to other PE routers.
3	The MP-BGP route is inserted as an interarea route into an OSPF area by the receiving PE router. The receiving PE router sets the down bit in the summary (type 3) LSA.
4	When the other PE routers receive the summary LSA with the down bit set, they do not redistribute the route back into MP-BGP.

Optimizing of Packet Forwarding Across the MPLS VPN Backbone

This topic describes how the optimizing of packet forwarding is accomplished across the MPLS VPN backbone.

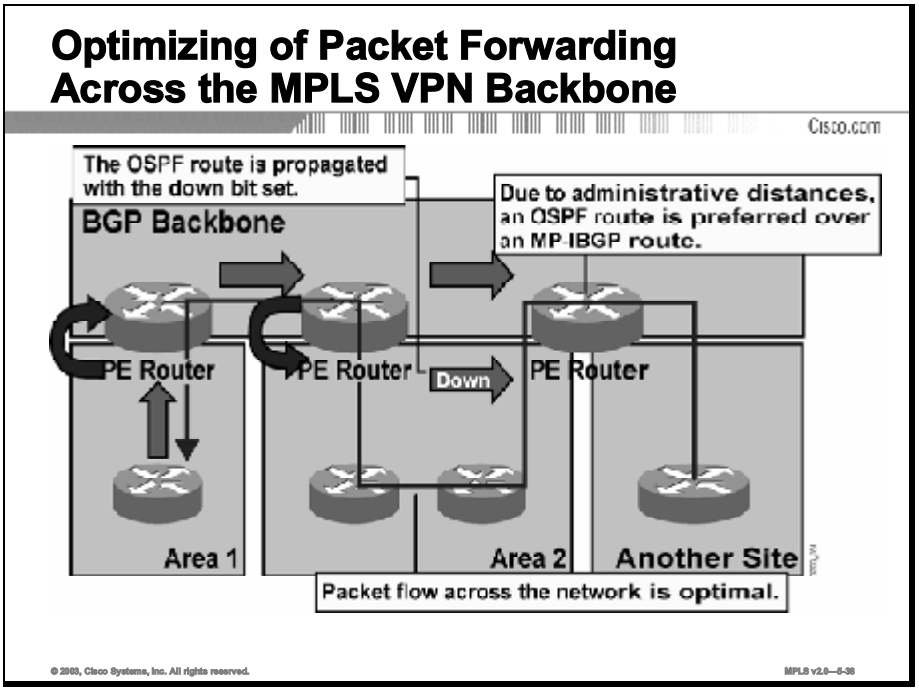

The OSPF superbackbone implementation with MP-BGP has other implications beyond the potential for routing loops between OSPF and BGP. Consider, for example, the network in the figure here.

The following table indicates a typical flow for routing updates.

Step	Action
1	The PE router redistributes the OSPF route into MP-BGP. The route is propagated to other PE routers as an MP-BGP route. It is also redistributed into other OSPF areas.
2	The redistributed OSPF route is propagated across the OSPF area with the down bit set.
3	The ingress PE router receives an MP-IBGP route with an administrative distance of 200 and an OSPF route with an administrative distance of 110. The OSPF route is preferred over the MP-IBGP route, and the data packets flow across customer sites, not directly over the MPLS VPN backbone.

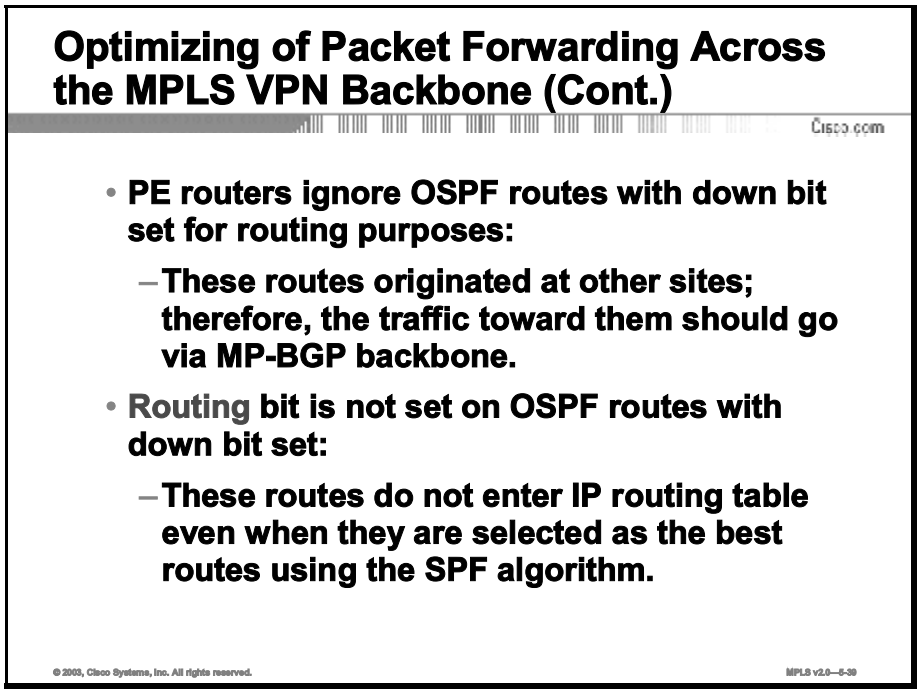

To prevent the customer sites from acting as transit parts of the MPLS VPN network, the OSPF route selection rules in PE routers need to be changed. The PE routers have to ignore all OSPF routes with the down bit set, because these routes originated in the MP-BGP backbone and the MP-BGP route should be used as the optimum route toward the destination.

This rule is implemented with the *routing* bit in the OSPF LSA. For routes with the down bit set, the routing bit is cleared and these routes never enter the IP routing table, even if they are selected as the best routes by the shortest path first (SPF) algorithm.

Note	The routing bit is the Cisco extension to OSPF and is used only internally in the router. It is never propagated between routers in LSA updates.

Optimizing of Packet Forwarding Across the MPLS VPN Backbone (Cont.)

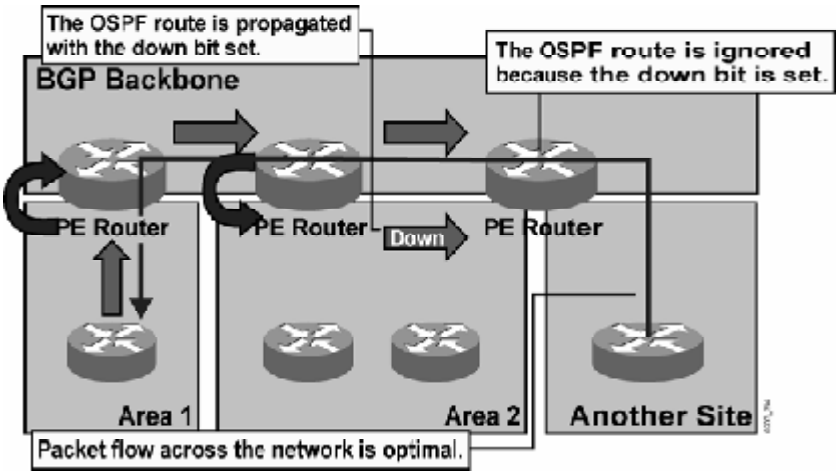

With the new route OSPF selection rules in place, the packet forwarding in the network shown here follows the desired path.

Step	Action
1	The OSPF route is redistributed into MP-BGP by a PE router and propagated to other PE routers.
2	The receiving PE routers redistribute the MP-BGP route into OSPF.

Other PE routers might receive the MP-BGP and OSPF routes but will ignore the OSPF route for routing purposes because it has the down bit set. The data packets will flow across the MPLS VPN backbone, following only the MP-BGP routes, not the OSPF routes derived from the MP-BGP routes.

OSPF Tag Field

This topic describes the route loop issue that may result from advertising routes across multiple OSPF domains in an MPLS VPN environment and how the OSPF tag field is used to address this issue.

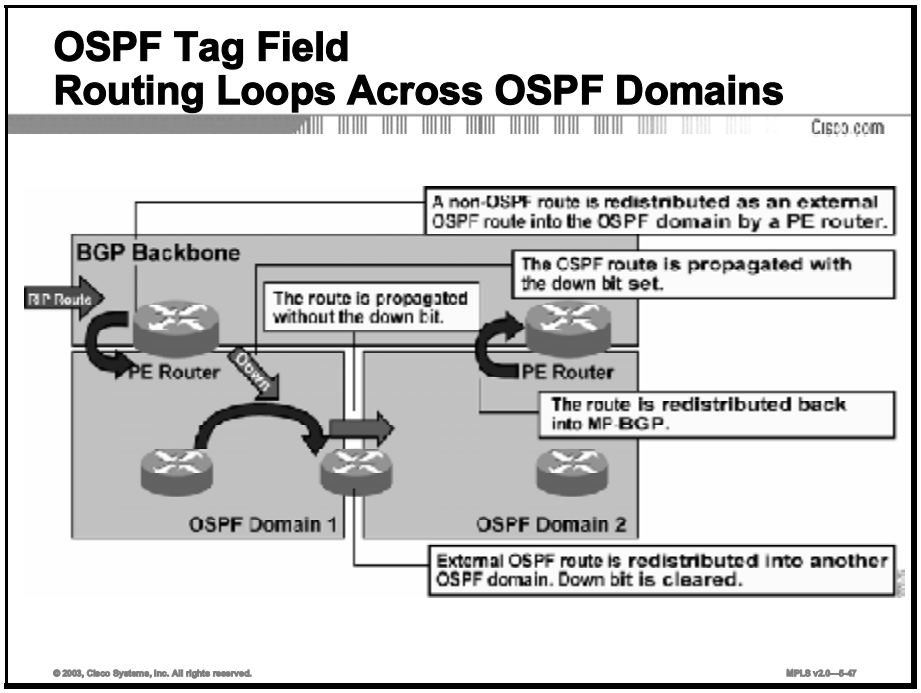

The down bit stops the routing loops between MP-BGP and OSPF. It cannot, however, stop the routing loops when redistribution between multiple OSPF domains is involved, as is the case in the network in the figure here.

The routing loop in this network occurs as part of the steps outlined in the table.

Step	Action
1	The PE router redistributes a non-OSPF route into an OSPF domain as an external route. The down bit is set because the route should not be redistributed back into MP-BGP.
2	A CE router redistributes the OSPF route into another OSPF domain. The down bit is lost if the CE router does not understand this OSPF extension.
3	The OSPF route is propagated through the other OSPF domain with the down bit cleared.
4	A PE-router receives the OSPF route; the down bit is not set, so the route is redistributed back into the MP-BGP backbone, resulting in a routing loop.

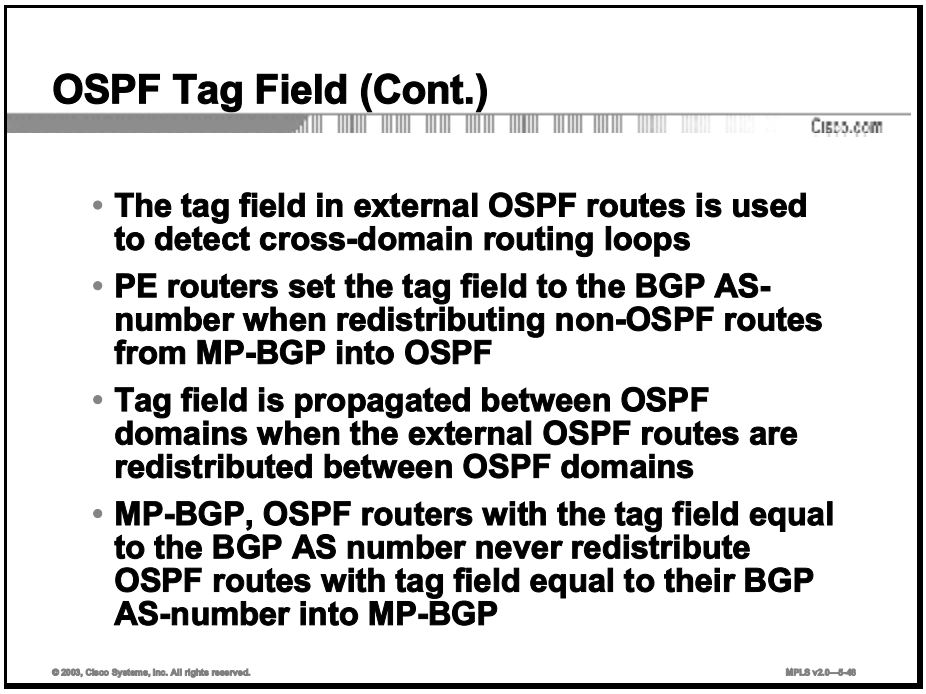

The routing loops introduced by route redistribution between OSPF domains can be solved with the help of the tag field, using standard –OSPF-BGP redistribution rules.

In standard –OSPF-BGP or OSPF-OSPF redistribution, the following rules apply:

- Whenever a router redistributes a BGP route into OSPF, the tag field in the type 5 (or type 7) LSA is set to the AS number of the redistributing router.
- The tag field from an external OSPF route is propagated across OSPF domains when the external OSPF route is redistributed into another OSPF domain.
- In addition to these standard mechanisms, PE routers filter external OSPF routes based on their tag field and do not redistribute, into MP-BGP, routes with a tag field equal to the BGP AS number.

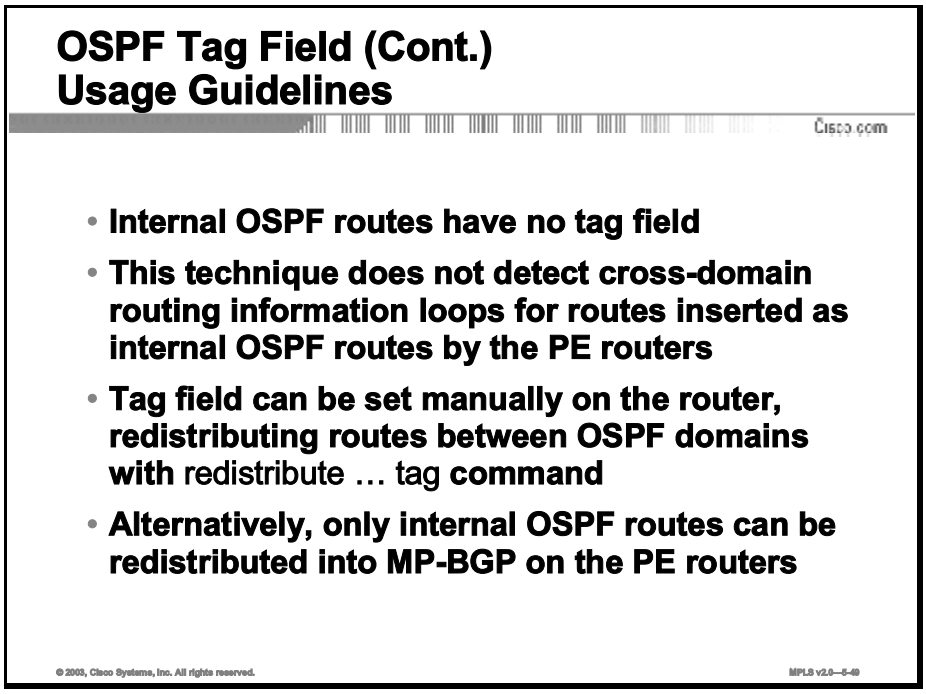

The OSPF tag field is present only in the external OSPF routes (type 5 LSA or type 7 LSA). This technique, therefore, cannot detect cross-domain loops involving internal OSPF routes. There are two manual methods that you can use to overcome this OSPF limitation:

- You can set the tag field manually on the router, redistributing routes between OSPF domains using the **redistribute ospf** *source-process-id* **tag** *value* command.
- The PE router can be configured to redistribute only internal OSPF routes into MP-BGP.

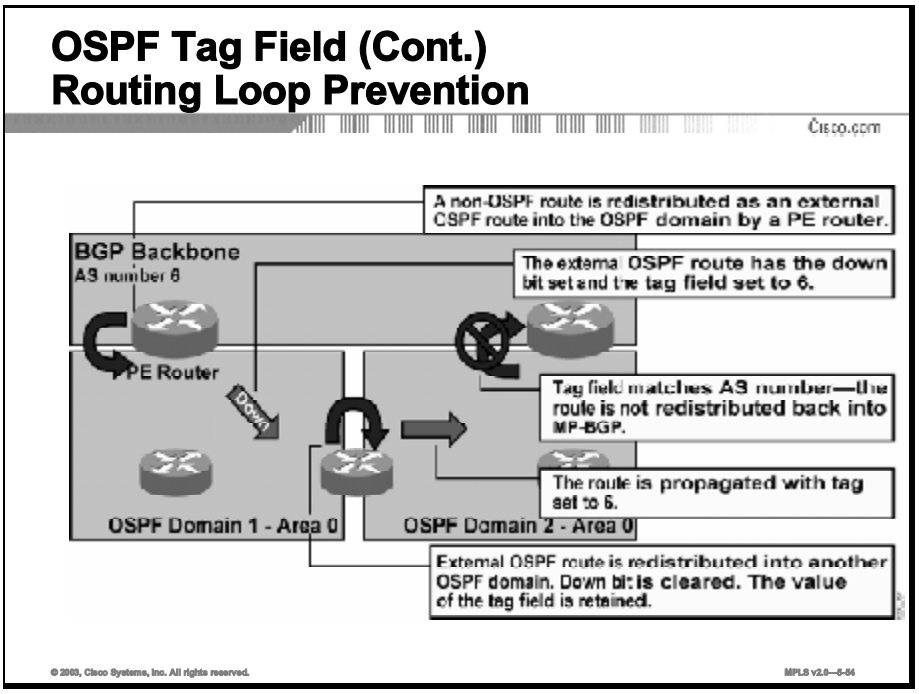

The figure illustrates how the OSPF tag field can be used to prevent routing loops when the redistribution is done between OSPF domains.

Step	Action
1	A non-OSPF route is redistributed as an external OSPF route by a PE router. The tag field is set to the BGP AS number and the down bit is set.
2	The redistributed route is propagated across the OSPF domain.
3	When the route is redistributed into another OSPF domain, the tag field is propagated, but the down bit is cleared.
4	Another PE router receives the external OSPF route and filters the route based on the tag field. The route is not redistributed into MP-BGP.

Sham Link

This topic describes the function of a sham link and how it is implemented.

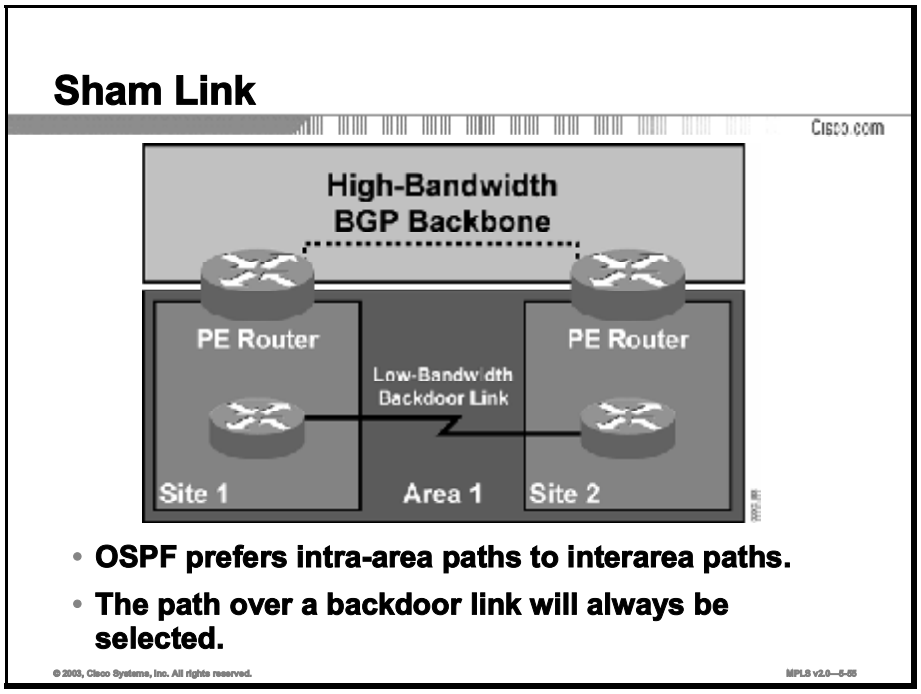

Although OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites (shown in the figure) may exist. If these sites belong to the same OSPF area, the path over a backdoor link will always be selected because OSPF prefers intra-area paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor links between VPN sites must be taken into account so that routing is performed based on policy.

Because each site runs OSPF within the same Area 1 configuration, all routing between the sites follows the intra-area path across the backdoor links, rather than over the MPLS VPN backbone.

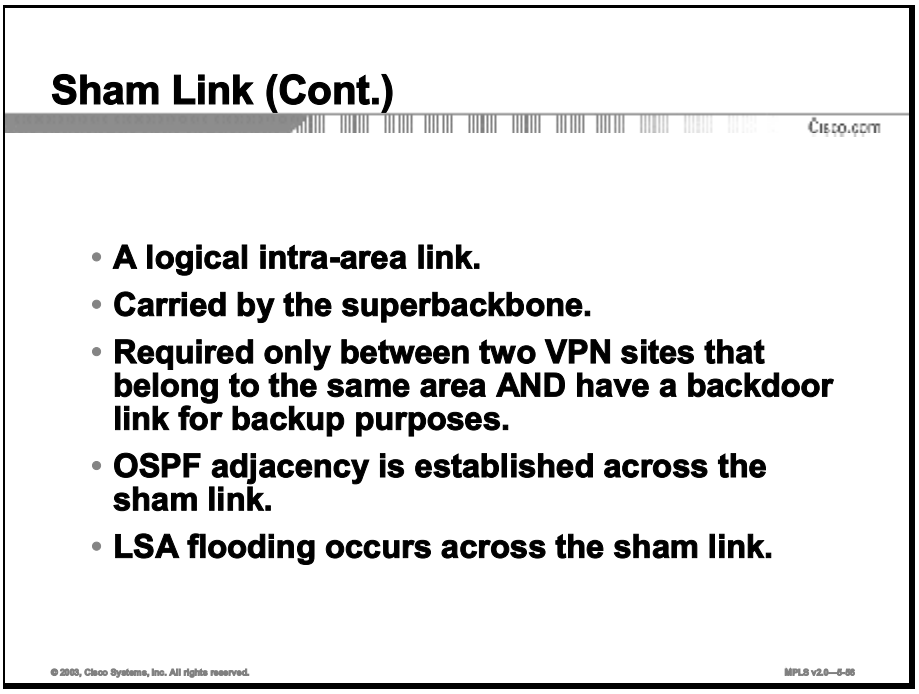

If the backdoor links between sites are used only for backup purposes and do not participate in the VPN service, then the default route selection shown in the preceding figure is not acceptable. To re-establish the desired path selection over the MPLS VPN backbone, you must create an additional OSPF intra-area (logical) link between ingress and egress VRFs on the relevant PE routers. This link is called a sham link.

A sham link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link. If no backdoor link exists between the sites, no sham link is required.

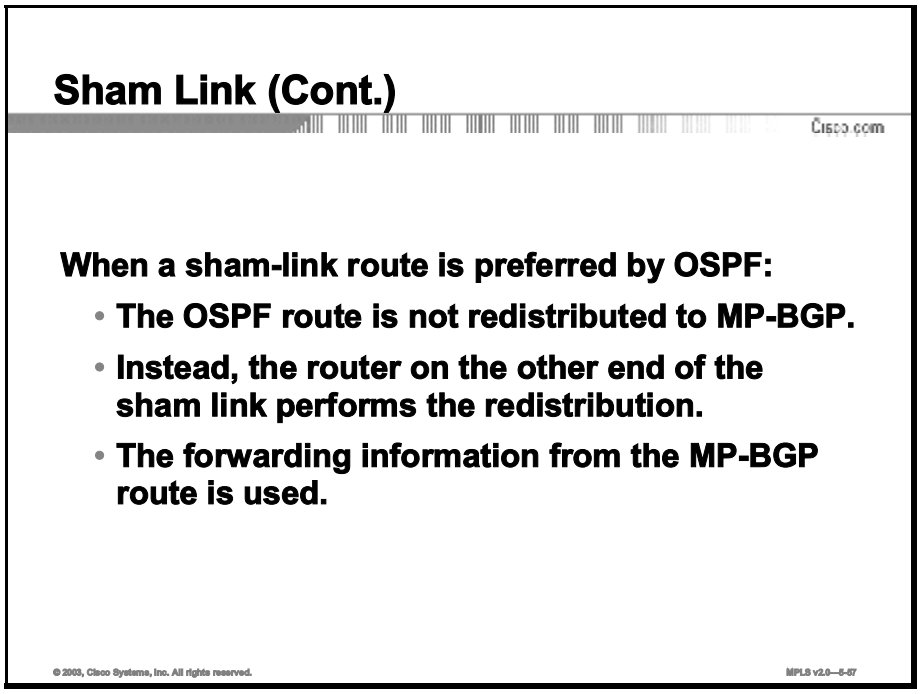

A cost is configured with each sham link and is used to decide whether traffic will be sent over the backdoor path or the sham-link path. When a sham link is configured between PE routers, the PEs can populate the VRF routing table with the OSPF routes learned over the sham link.

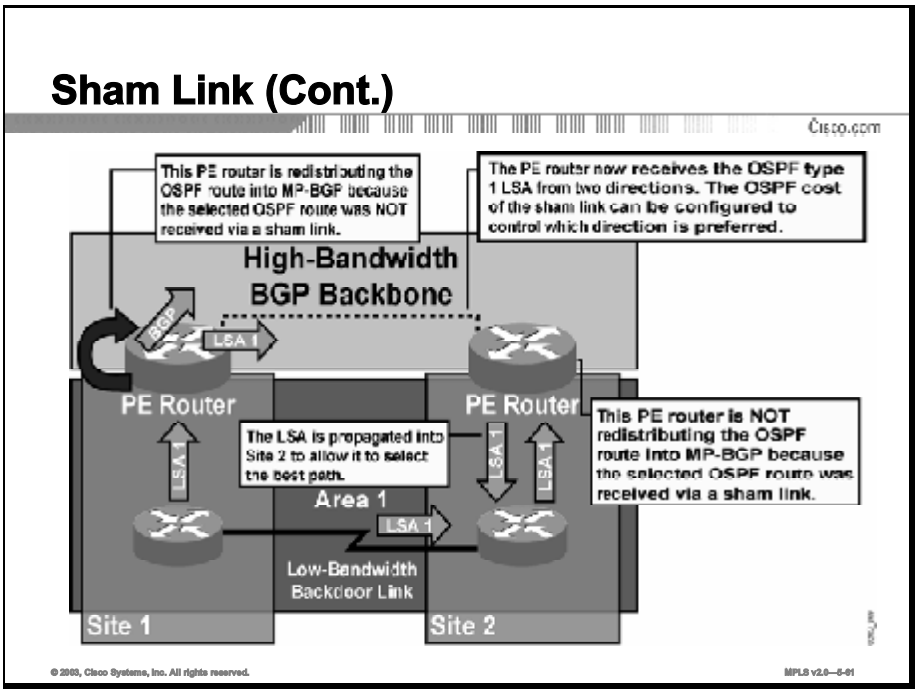

Because the sham link is seen as an intra-area link between PE routers, an OSPF adjacency is created and database exchange (for the particular OSPF process) occurs across the link. The PE router can then flood LSAs between sites from across the MPLS VPN backbone. As a result, the desired intra-area connectivity is created.

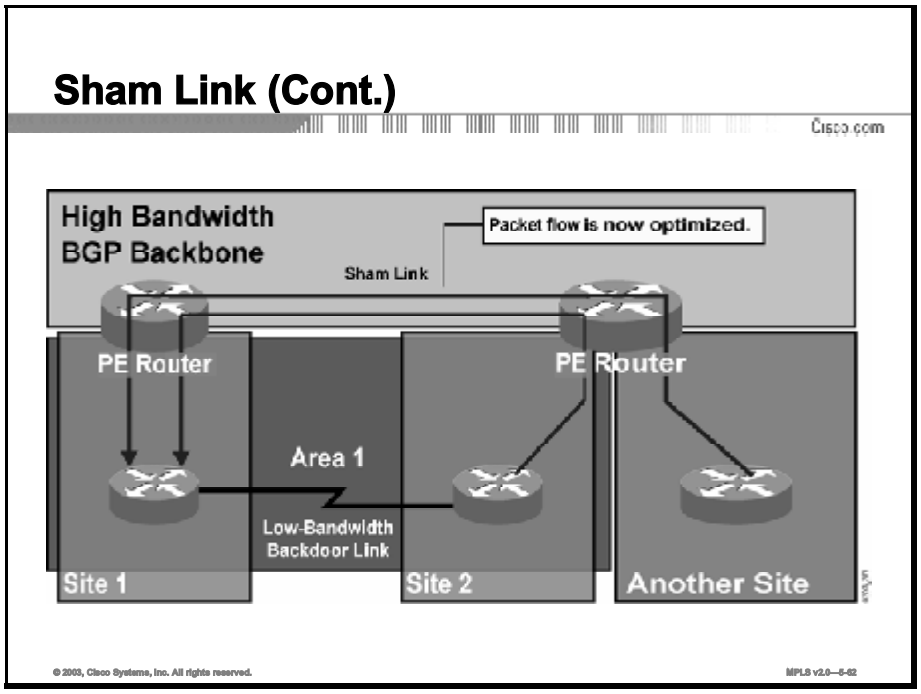

The implementation results in optimum packet flow.

Configuring a Sham Link

This topic describes how to implement a sham link.

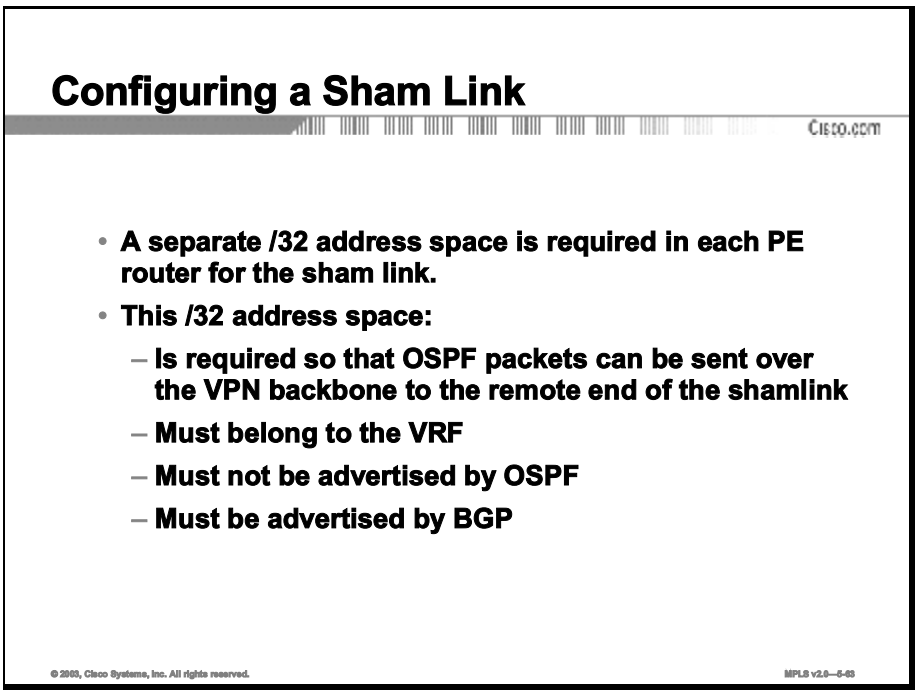

When you are implementing a sham link, a separate /32 address space is required in each PE router.

The following criteria apply to this /32 address space:

- Is required so that OSPF packets can be sent over the VPN backbone to the remote end of the sham link
- Must belong to the VRF
- Must not be advertised by OSPF
- Must be advertised by BGP

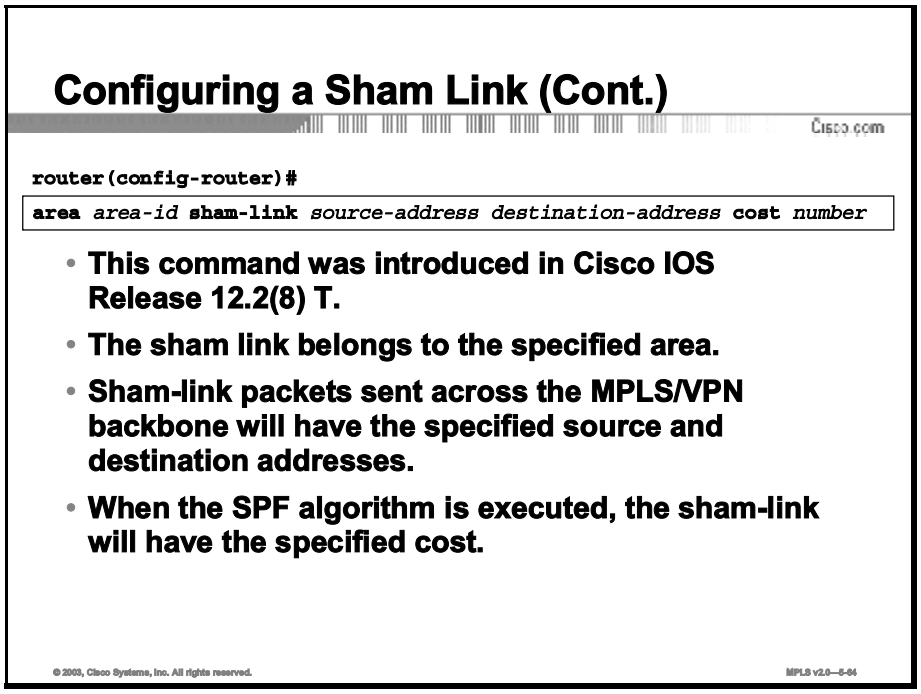

To configure a sham-link interface on a PE router in an MPLS VPN backbone, use the **area sham-link cost** command in global configuration mode. To remove the sham link, use the **no** form of this command:

- **area** *area-id* **sham-link** *source-address destination-address* **cost** *number*
- **no area** *area-id* **sham-link** *source-address destination-address* **cost** *number*

Syntax Description

Parameter	Description
area-id	ID number of the OSPF area assigned to the sham link. Valid values: numeric value or valid IP address. There is no default.
source-address	IP address of the source PE router in the format: *ip-address* [*mask*].
destination-address	IP address of the destination PE router in the format: *ip-address* [*mask*].
number	OSPF cost to send IP packets over the sham-link interface. Valid values are from 1 to 65535.

Defaults

There is no default behavior or values.

Command Modes

Global configuration

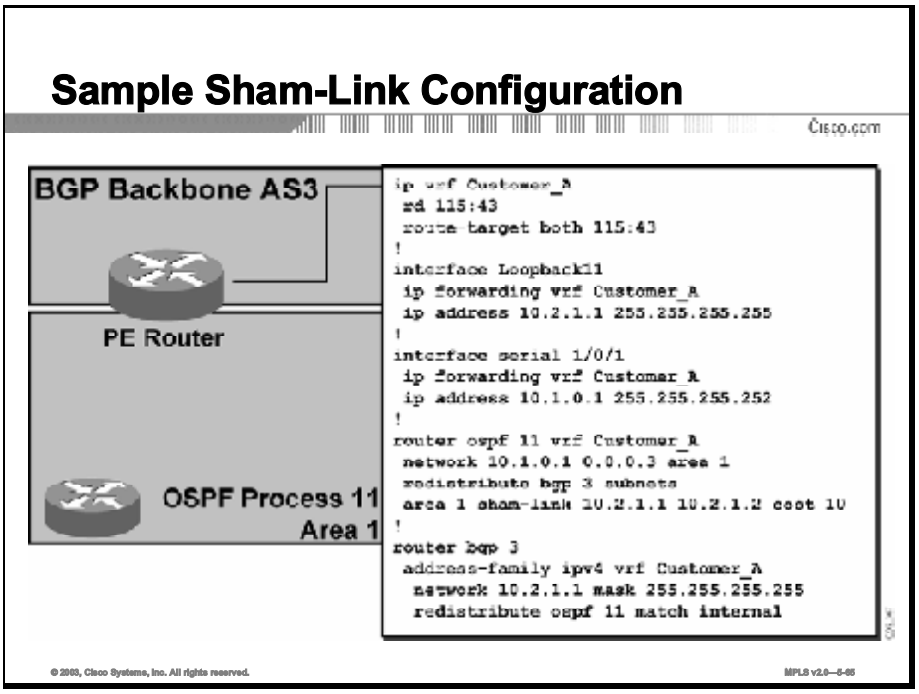

The example here shows how a sham link is used only to affect the OSPF intra-area path selection of the PE and CE routers. The PE router also uses the information received from MP-BGP to set the outgoing label stack of incoming packets, and to decide to which egress PE router to label-switch the packets.

The figure shows a sample MPLS VPN topology in which a sham-link configuration is necessary. A VPN client has two sites connected by a backdoor link. A sham link has been configured between the two PE routers.

Summary

This topic summarizes the key points discussed in this lesson.

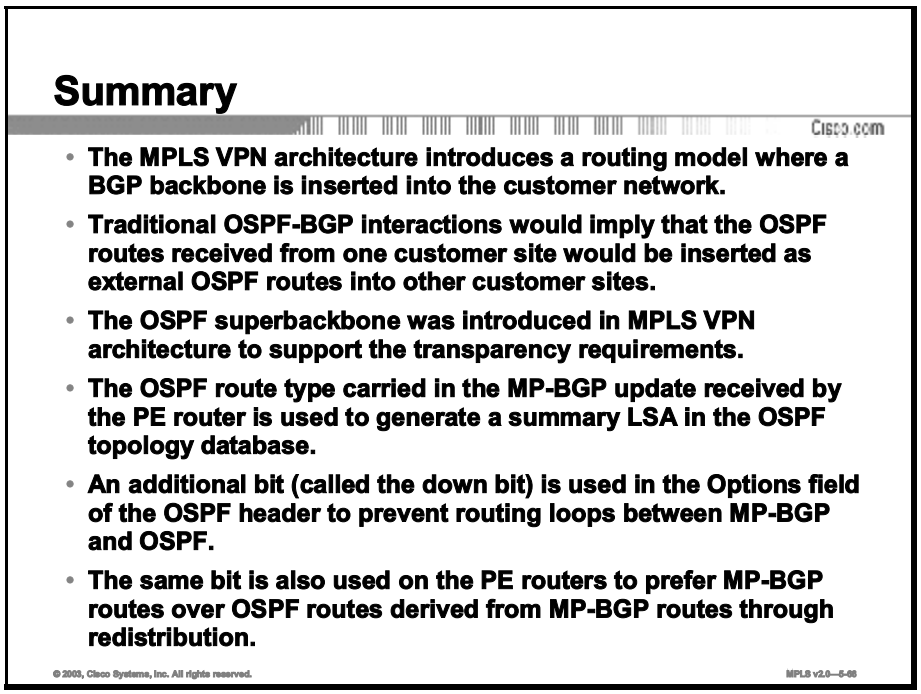

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Why is the OSPF superbackbone needed in MPLS VPN environments?

Q2) What is the interaction between Area 0 and a superbackbone?

Q3) What is the interaction between a superbackbone and other areas?

Q4) How are OSPF route attributes propagated across an MPLS VPN backbone?

Q5) What is the purpose of the down bit in an LSA header?

Quiz Answer Key

Q1) Because MPLS VPNs use BGP to propagate routes between sites, internal OSPF routers in one area will appear as external routes in another area unless the superbackbone makes the MPLS VPN backbone transparent to OSPF.

Relates to: OSPF Hierarchical Model; OSPF in an MPLS VPN Routing Model; OSPF Superbackbone

Q2) The superbackbone is transparent to Area 0.

Relates to: OSPF in an MPLS VPN Routing Model; OSPF Superbackbone

Q3) The superbackbone appears as Area 0 to non-Area 0 areas.

Relates to: OSPF in an MPLS VPN Routing Model; OSPF Superbackbone

Q4) The OSPF routes are propagated into BGP.
The OSPF metrics and LSA information are carried in the BGP community attribute.

Relates to: OSPF in an MPLS VPN Routing Model; OSPF Superbackbone

Q5) The down bit is used to prevent routing loops.

Relates to: OSPF Down Bit

Configuring BGP as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the PE-CE routing protocol configuration steps required when you are using BGP as the routing protocol between PE and CE routers.

Relevance

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need in order to configure MPLS VPN PE-CE routing exchange.

Objectives

The lesson identifies the command syntax that is required to configure PE-CE routing protocols.

Upon completing this lesson, you will be able to:

- Identify the command syntax that is used to configure a per-VRF BGP routing context
- Describe the reason for limiting the number of routes in a VRF and identify the command syntax that is required to enable this feature
- Describe the reason for limiting the number of prefixes received from a BGP neighbor and identify the command syntax that is required to enable this function
- Describe the reason for limiting the total number of VRF routes and identify the command syntax that is required to configure VRF route limits
- Describe the issues encountered when a customer wants to reuse the same AS number on several sites and describe how the AS-override feature is implemented to solve these issues

- Describe the issues encountered when a customer site links two VPNs and how the allowas-in feature is used to address these issues
- Describe how the BGP SOO attribute can be used as a loop prevention mechanism in an MPLS VPN environment

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Configuring Per-VRF BGP Routing Context
- Limiting the Number of Routes in a VRF
- Limiting the Number of Prefixes Received from a BGP Neighbor
- Limiting the Total Number of VRF Routes
- AS-Override
- Allowas-in
- Implementing SOO for Loop Prevention
- Summary
- Quiz

Configuring Per-VRF BGP Routing Context

This topic identifies the command syntax that is used to select the VRF routing context for BGP.

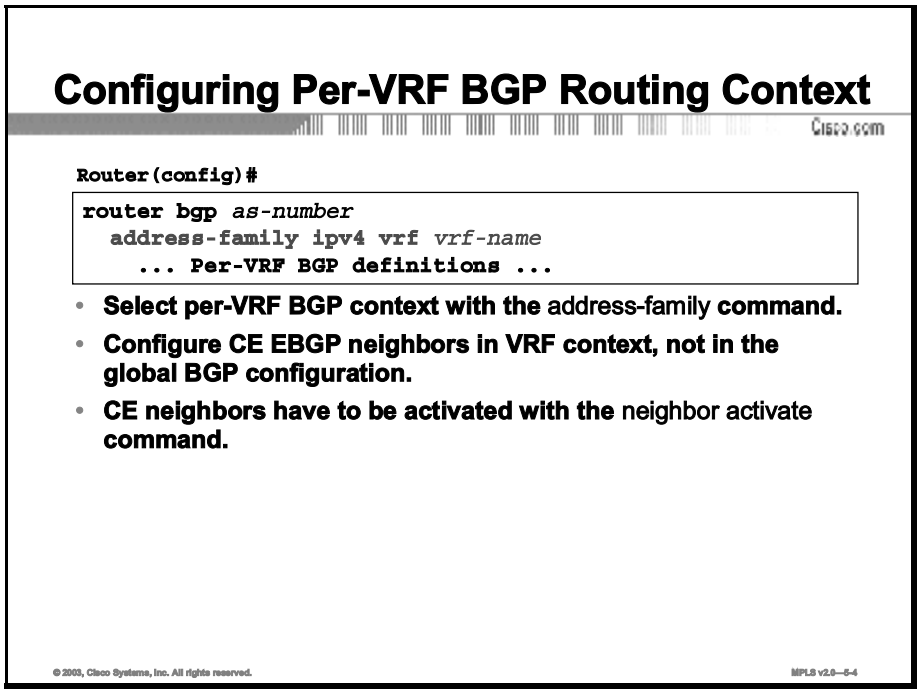

Select the VRF routing context with the **address-family ipv4 vrf** *vrf-name* command in the RIP and BGP routing processes. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

When you configure BGP as the PE-CE routing protocol, you must start with the per-VRF BGP configuration. Use the **address-family ipv4 vrf** *vrf-name* command in router configuration mode. Enter the address family configuration mode, and then define and activate the BGP neighbors. You also have to configure redistribution from all other per-VRF routing protocols into BGP.

Note	You always have to configure a BGP address family for each VRF and configure route redistribution into BGP for each VRF even if you do not use BGP as the PE-CE routing protocol.

Several BGP options have different default values when you configure the per-VRF BGP routing context:

- BGP synchronization is disabled (default = enabled).
- Autosummarization (automatic generation of classful networks out of subnets redistributed into BGP) is disabled (default = enabled), because the MPLS VPN backbone has to propagate customer subnets unchanged to facilitate transparent end-to-end routing between customer sites.

Redistribution of internal BGP routes into IGP is enabled (default = disabled).

Note	Common parameters defined in router configuration mode are inherited by all address families defined for this routing process and can be overridden for each individual address family.

address-family ipv4

To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the **address-family ipv4** command in router configuration mode. To disable address family configuration mode, use the **no** form of this command:

- **address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]
- **no address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]

Syntax Description

Parameter	Description
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vrf *vrf-name*	(Optional) Specifies the name of the VRF instance to associate with subsequent IPv4 address family configuration mode commands.

Defaults

IPv4 address prefixes are not enabled. Unicast address prefixes are the default when IPv4 address prefixes are configured.

Command Modes

Router configuration

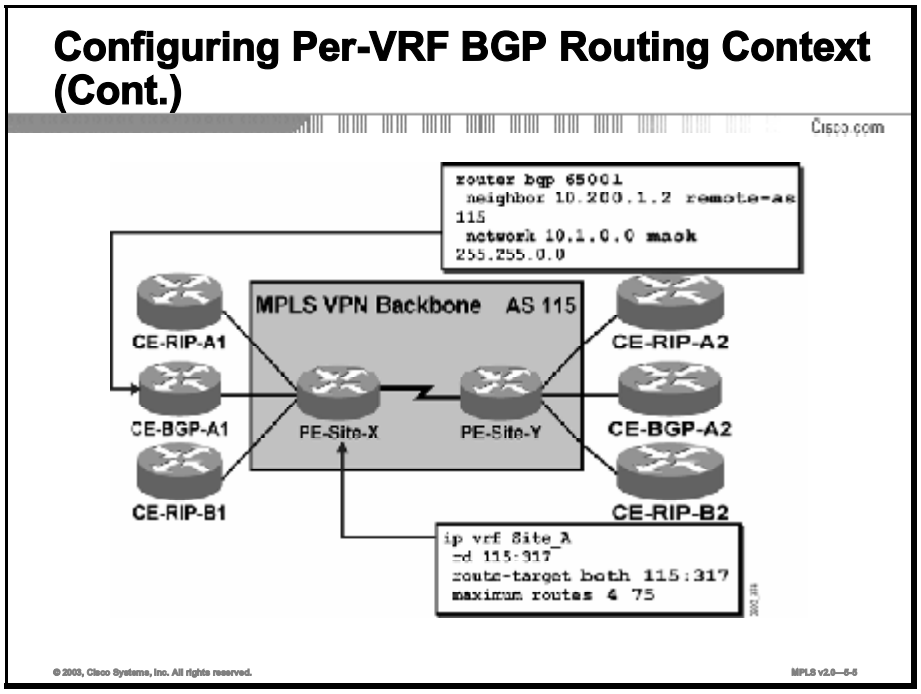

The figure here shows BGP that is activated on the CE router. The PE router is defined as a BGP neighbor. Similarly, the CE router is defined as a BGP neighbor and activated under **address-family ipv4 vrf** *Customer_A*.

Limiting the Number of Routes in a VRF

This topic presents the reason for limiting the number of routes in a VRF and identifies the command syntax that is required to enable this feature.

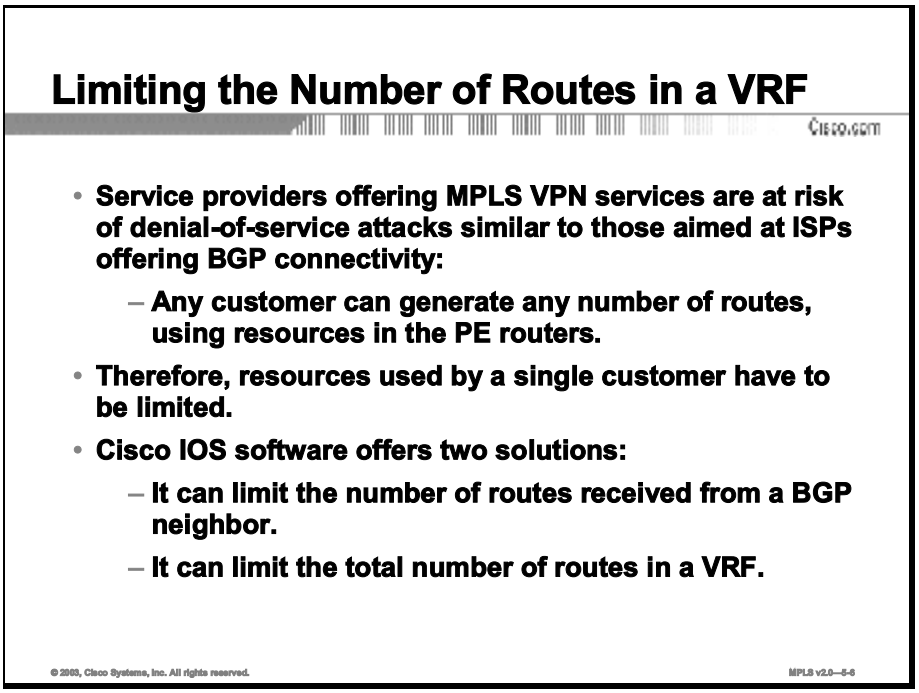

MPLS VPN architecture achieves a very tight coupling between the customer and the service provider network, resulting in a number of advantages. The tight coupling might also result in a few disadvantages, because the service provider network is exposed to design and configuration errors in customer networks, as well as a number of new denial-of-service attacks based on routing protocol behavior.

To limit the effect of configuration errors as well as malicious user behavior, Cisco IOS software offers two features that limit the number of routes and resource consumption that a VPN user can take advantage of at a PE router:

- The BGP maximum-prefix feature limits the number of routes that an individual BGP peer can send.
- The VRF route limit limits the total number of routes in a VRF regardless of whether they are received from CE routers or from other PE routers via MP-IBGP.

Limiting the Number of Prefixes Received from a BGP Neighbor

This topic presents the reason for limiting the number of prefixes received from a BGP neighbor and identifies the command syntax that is required to enable this function.

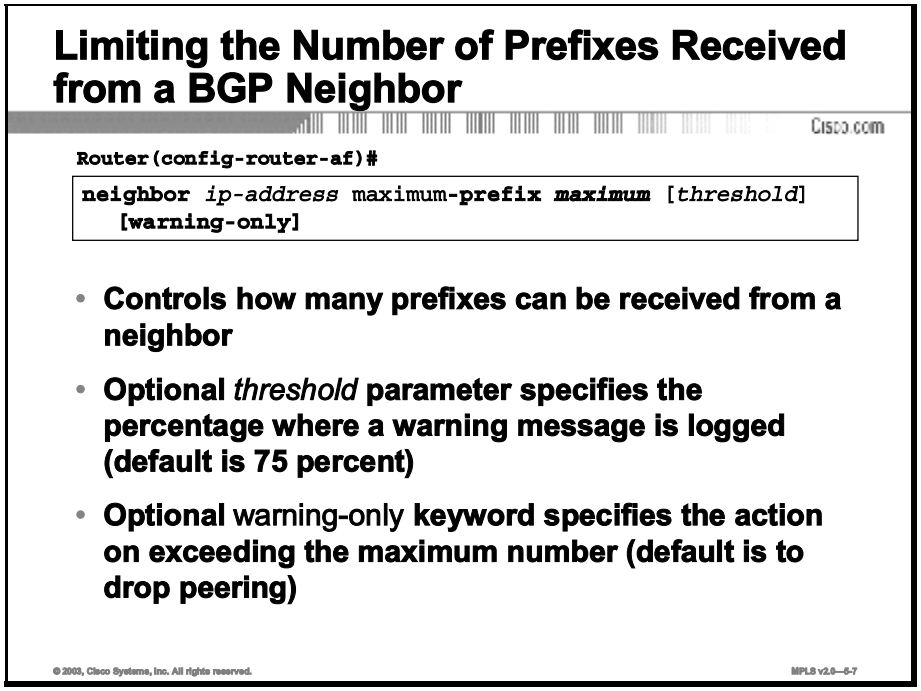

neighbor maximum-prefix

To control how many prefixes can be received from a neighbor, use the **neighbor maximum-prefix** command in router configuration mode. To disable this function, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **maximum-prefix** *maximum* [*threshold*] [**warning-only**]
- **no neighbor** {*ip-address* | *peer-group-name*} **maximum-prefix** *maximum* [*threshold*] [**warning-only**]

Syntax Description

Parameter	Description
ip-address	IP address of the neighbor.
peer-group-name	Name of a BGP peer group.
maximum	Maximum number of prefixes allowed from this neighbor.
threshold	(Optional) Integer specifying at what percentage of maximum the router starts to generate a warning message. The range is 1 to 100; the default is 75 (percent).
warning-only	(Optional) Allows the router to generate a log message when the maximum is exceeded instead of terminating the peering.

Defaults

Default is disabled; there is no limit on the number of prefixes.

Limiting the Total Number of VRF Routes

This topic describes the reason for limiting VRF routes and identifies the command syntax that is required to configure VRF route limits.

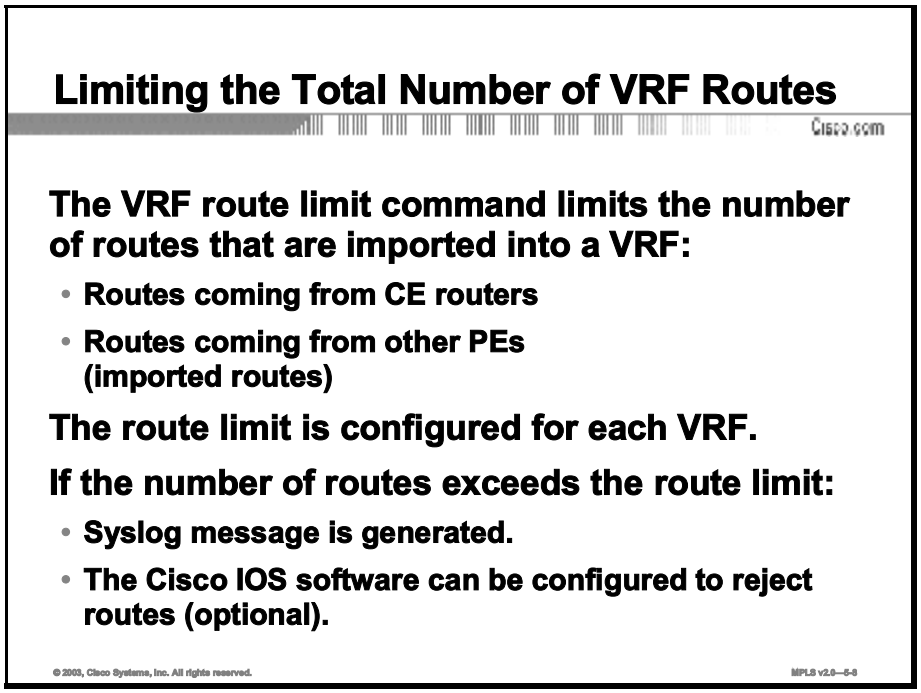

The VRF route limit, contrary to the BGP maximum-prefix limit, limits the overall number of routes in a VRF regardless of their origin. Similar to the BGP maximum-prefix feature, the network operator might be warned via a syslog message when the number of routes exceeds a certain threshold. Additionally, you can configure Cisco IOS software to ignore new VRF routes when the total number of routes exceeds the maximum configured limit.

The route limit is configured for each individual VRF, providing maximum design and configuration flexibility.

Note	The per-VRF limit could be used to implement add-on MPLS VPN services. A user desiring a higher level of service might be willing to pay to be able to insert more VPN routes into the network.

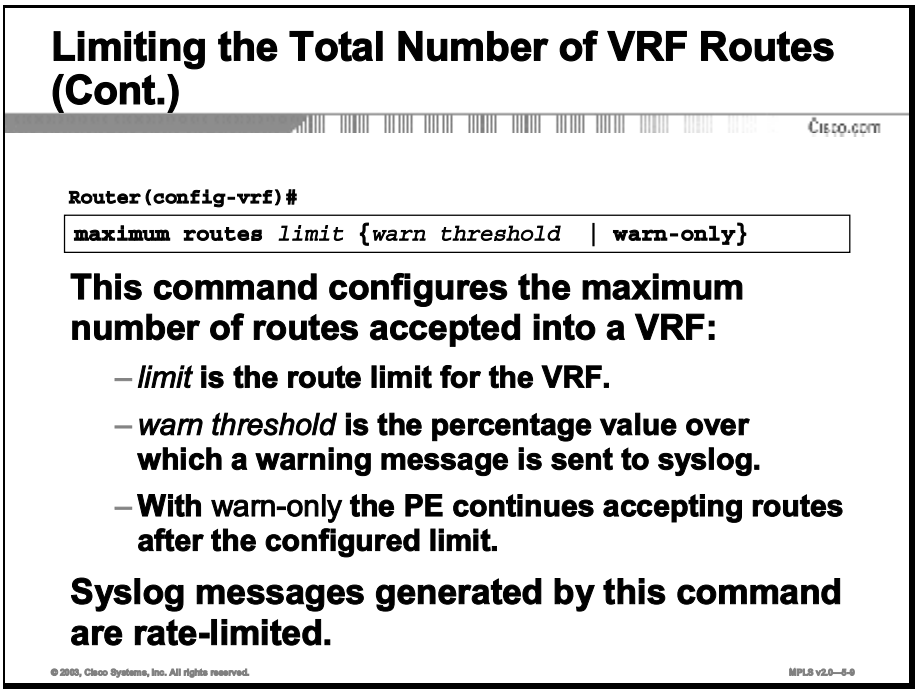

maximum routes

To limit the maximum number of routes in a VRF to prevent a PE router from importing too many routes, use the **maximum routes** command in VRF configuration submode. To remove the limit on the maximum number of routes allowed, use the **no** form of this command:

- **maximum routes** *limit* {*warn threshold* | **warn-only**}
- **no maximum routes**

Syntax Description

Parameter	Description
limit	Specifies the maximum number of routes allowed in a VRF. You may select from 1 to 4,294,967,295 routes to be allowed in a VRF.
warn threshold	Rejects routes when the threshold limit is reached. The threshold limit is a percentage of the limit specified, from 1 to 100.
warn-only	Issues a syslog error message when the maximum number of routes allowed for a VRF exceeds the threshold. However, additional routes are still allowed.

Defaults

This command has no default behavior or values.

Limiting the Total Number of VRF Routes (Cont.)

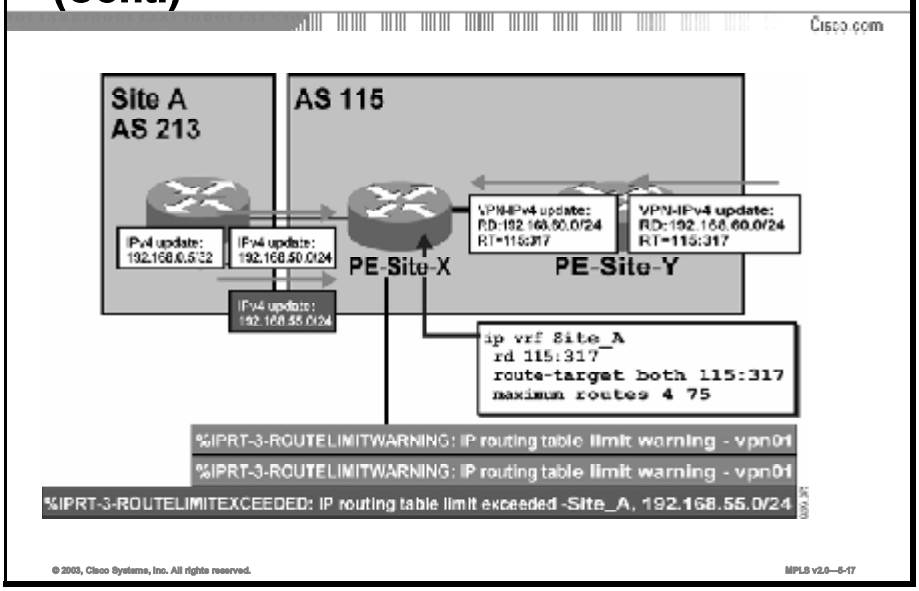

In this example, the network designer has decided to limit the number of routes in a VRF to four, with the warning threshold being set at 75 percent (or three routes).

When the first two routes are received and inserted into the VRF, the router accepts them. When the third route is received, a warning message is generated, and the message is repeated with insertion of the fourth route.

Note	The syslog messages are rate-limited to prevent indirect denial-of-service attacks on the network management station.

When the PE router receives the fifth route, the maximum route limit is exceeded and the route is ignored. The network operator is notified through another syslog message.

AS-Override

This topic uses a sample MPLS VPN network to identify the issues encountered when a customer wants to reuse the same AS number on several sites.

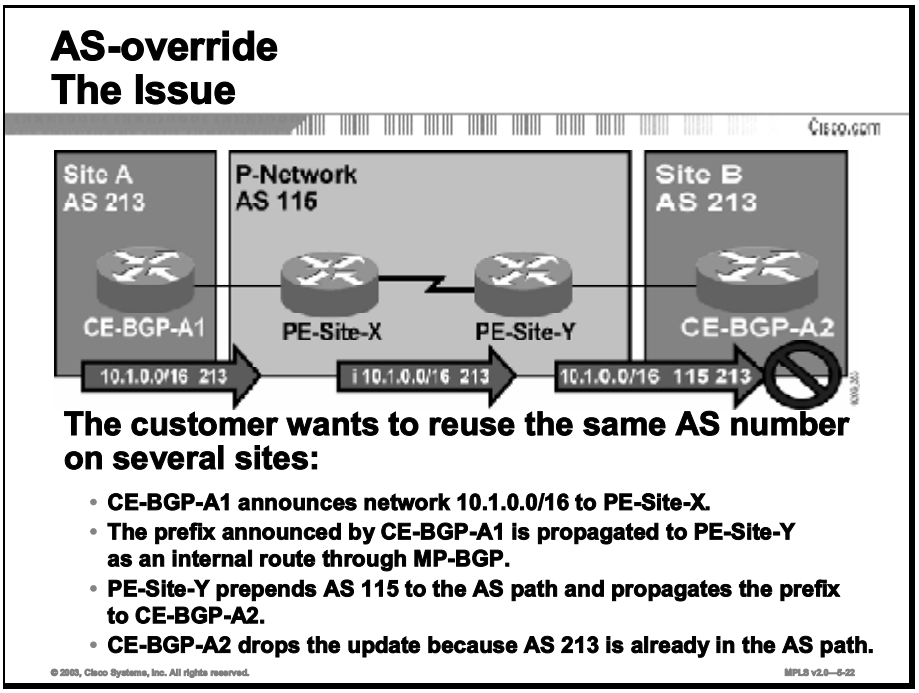

There are two ways that an MPLS VPN customer can deploy BGP as the routing protocol between PE and CE routers:

- If the customer has used any other routing protocol in the traditional overlay VPN network before, there are no limitations on the numbering of the customer autonomous systems. Every site can be a separate AS.
- If the customer has used BGP as the routing protocol before, there is a good chance that all the sites (or a subset of the sites) are using the same AS number.

BGP loop prevention rules disallow discontiguous autonomous systems. Two customer sites with the identical AS number cannot be linked by another AS. If such a setup happens (as in this example), the routing updates from one site are dropped when the other site receives them. There is no connectivity between the sites.

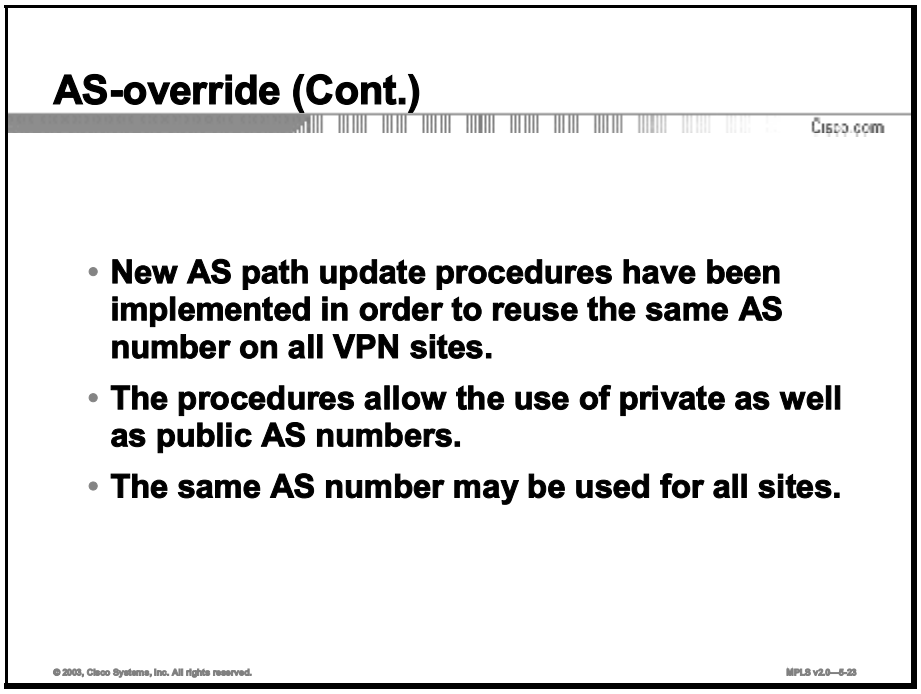

When you are migrating customers from traditional overlay VPNs to MPLS VPNs, it is not uncommon to encounter a customer topology that requires the same customer AS number to be used at more than one site. This requirement can cause issues with the loop prevention rules of BGP. However, the AS path update procedure in BGP has been modified to address this issue. The new AS path update procedure supports the use of one AS number at many sites (even between several overlapping VPNs) and does not rely on a distinction between private and public AS numbers.

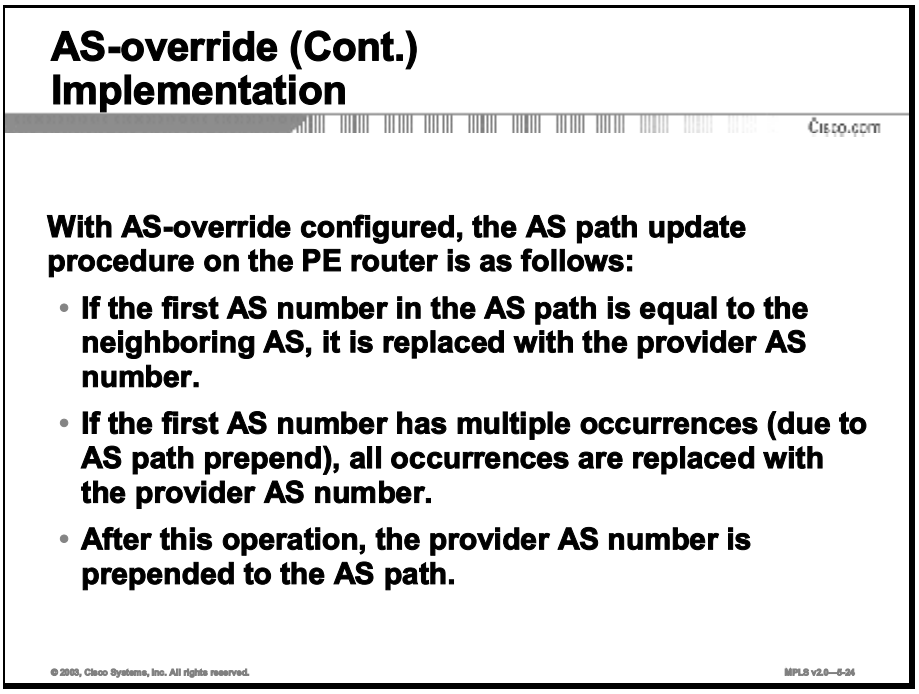

The modified AS path update procedure is called AS-override and is extremely simple:

- The procedure is used only if the first AS number in the AS path is equal to the AS number of the receiving BGP router.
- In this case, all leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Occurrences further down the AS path of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- An extra copy of the sending router AS number is prepended to the AS path. The standard AS number prepending procedure occurs on every EBGP update.

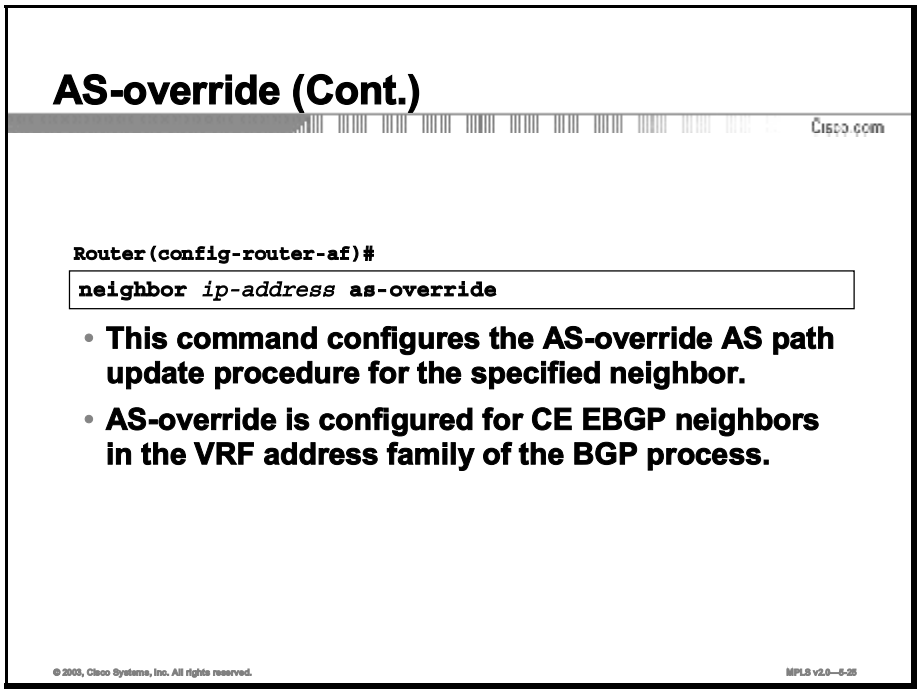

neighbor as-override

To configure a PE router to override a site AS number with a provider AS number, use the **neighbor as-override** command in router configuration mode. To remove VPNv4 prefixes from a specified router, use the **no** form of this command:

- **neighbor** *ip-address* **as-override**
- **no neighbor** *ip-address* **as-override**

Syntax Description

Parameter	Description
ip-address	Specifies the router IP address to override with the AS number provided.

Defaults

This command has no default behavior or values.

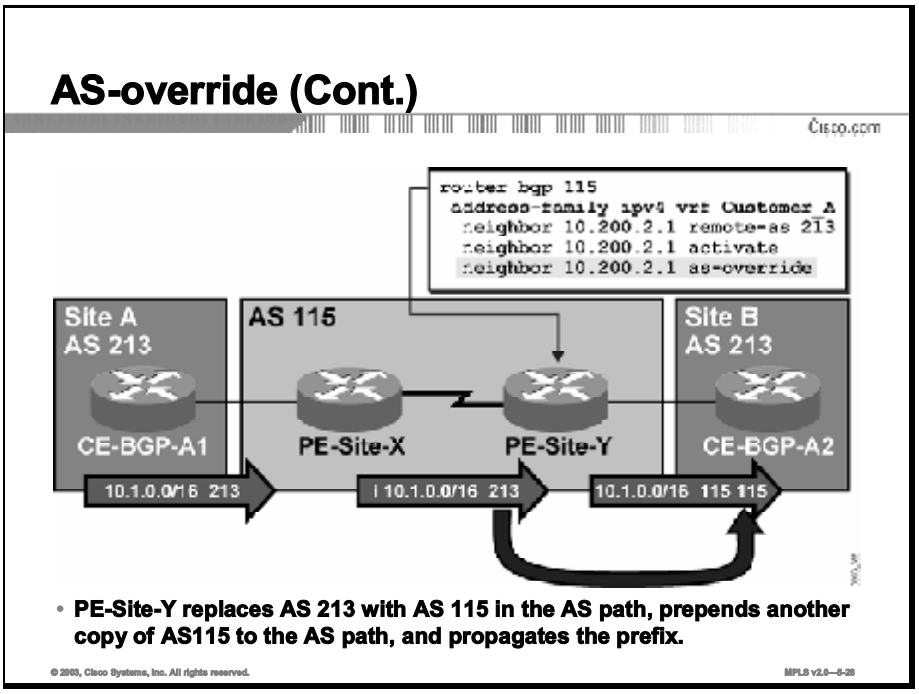

In the figure, customer sites A and B use BGP to communicate with the MPLS VPN backbone. Both sites use AS 213. Site B would drop the update sent by site A without the AS-override mechanism.

The AS-override mechanism, configured on the PE-Site-Y router, replaces the customer AS number (213) with the provider AS number (115) before sending the update to the customer site. An extra copy of the provider AS number is prepended to the AS path during the standard EBGP update process.

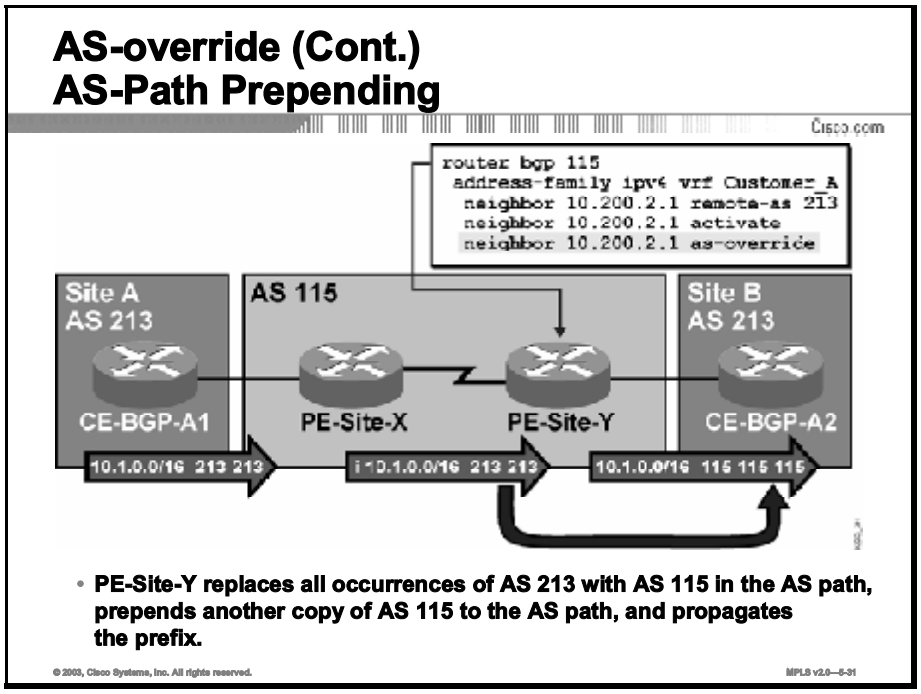

In this figure, the customer is using AS prepending to influence BGP path selection within the MPLS VPN backbone. The PE router has to send a route with an AS path containing multiple copies of the customer AS number to the CE router. In this case, all the leading copies of the customer AS number are replaced with the provider AS number (resulting in two occurrences of the provider AS number in the example), and the third occurrence of the provider AS number is prepended to the BGP update before it is sent to the CE router.

Allowas-in

This topic uses a sample MPLS VPN network to describe the issues encountered when a customer site links two VPNs.

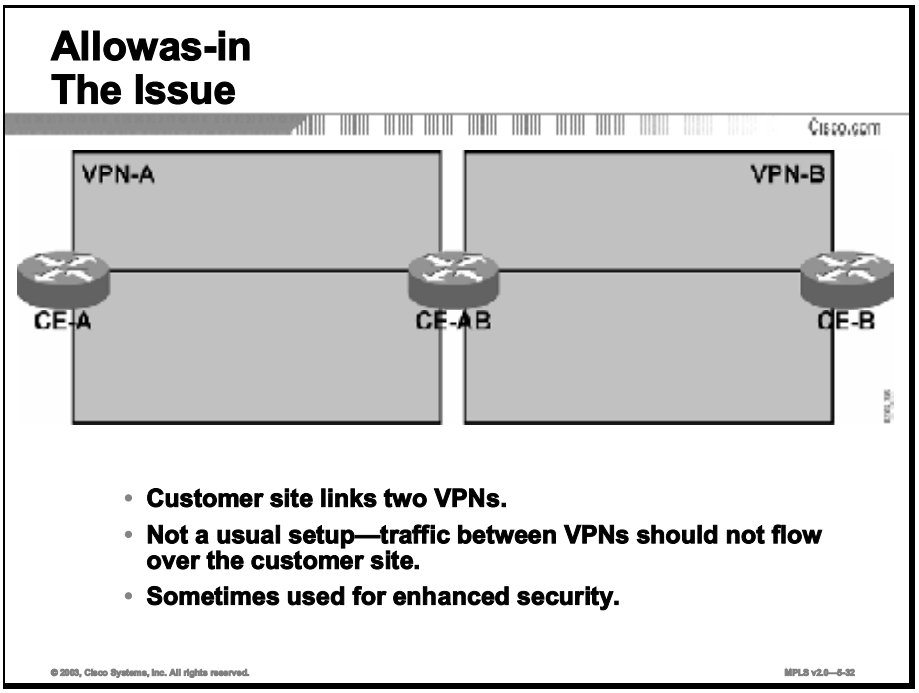

In some security-conscious implementations, customer VPNs are linked by a customer router that performs security functions such as access filtering or access logging.

Note	This setup is not usual because it deviates from the basic goal of MPLS VPN—replacing the hub-and-spoke routing of a traditional overlay VPN with optimum any-to-any routing.

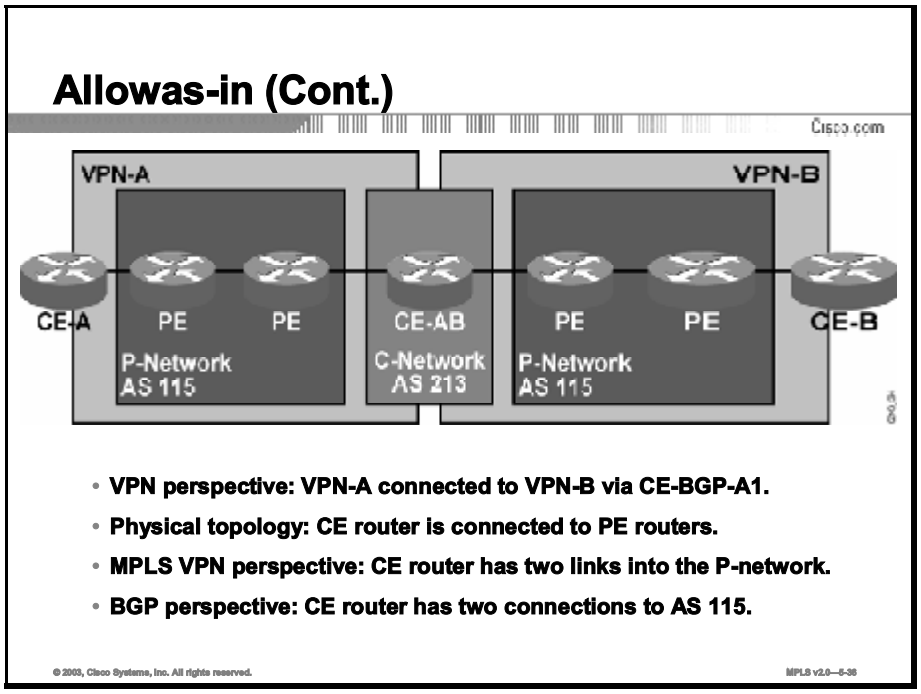

The setup in which a customer router links two VPNs in an MPLS VPN backbone can be viewed from several different perspectives:

- From the VPN perspective, a CE router links two VPNs.
- From the physical perspective, the CE router is connected through two separate links (physical or logical interface) to one or two PE routers.
- In MPLS VPN terms, the CE router has two links into the P-network.

There is no problem with the proposed customer setup if the setup is analyzed through these perspectives. They all represent valid connectivity or routing options. The problem occurs when the setup is analyzed through the BGP perspective, in which the CE router has to propagate routes between two PE routers, which are both in the same AS.

Allowas-in (Cont.)

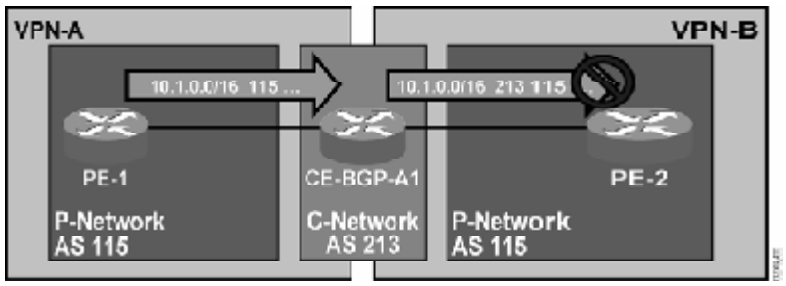

- **PE-1 announces network 10.1.0.0/16 to CE-BGP-A1.**
- **CE-BGP-A1 prepends its AS number to the AS path and propagates the prefix to PE-2.**
- **PE-2 drops the update because its AS number is already in the AS path.**
- **AS-override is needed on CE-BGP-A1, but that would require a Cisco IOS software upgrade on the CE router.**

This example is similar to the situation in which two customer sites use the same AS number. The BGP loop prevention rules prevent a PE router from accepting the routing update sent by the CE router if that routing update already contains the AS number of the MPLS VPN backbone (which it will if the CE router is propagating routes between two VPNs).

The solution to this BGP routing problem is identical to the previous one presented in this lesson. AS-override has to be used on the CE router. This solution requires a very recent version of Cisco IOS software (Cisco IOS Release 12.0 T or later) on the CE router. It is not enforceable in every customer situation.

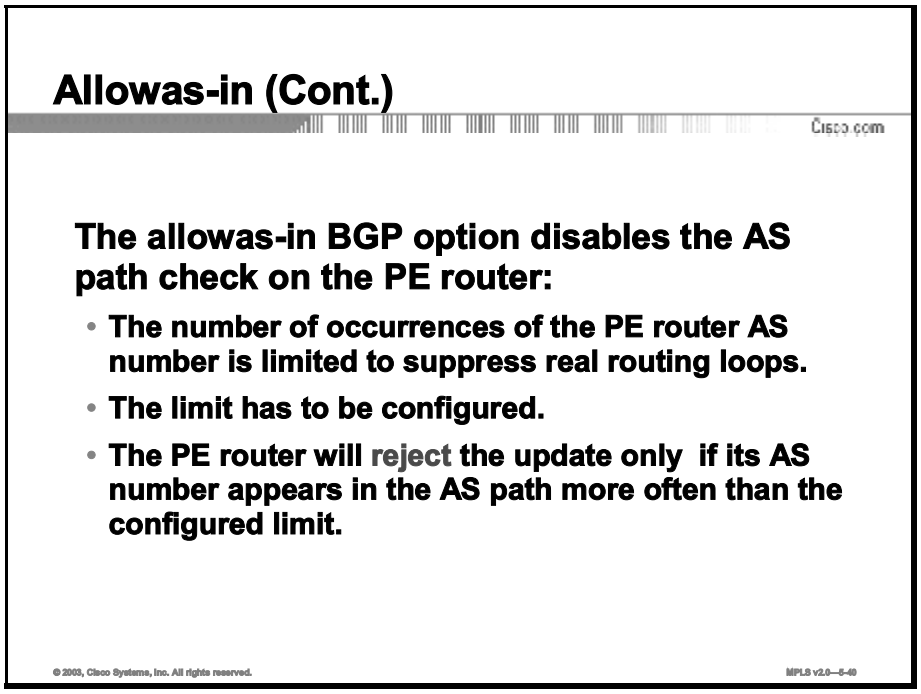

Networks need to support topologies in which a CE router with no AS-override support links two VPNs. A specific need exists to modify the BGP loop prevention mechanism on the PE routers. The allowas-in feature supports situations in which the PE router receives routes with its own AS number already in the AS path.

With this feature configured on a BGP neighbor of the PE router, the PE router would not drop incoming BGP updates with its AS number in the AS path if the updates are received from that neighbor. To prevent real BGP routing information loops, the number of occurrences of the MPLS VPN backbone AS number can be limited and incoming updates that exceed the limit can be dropped.

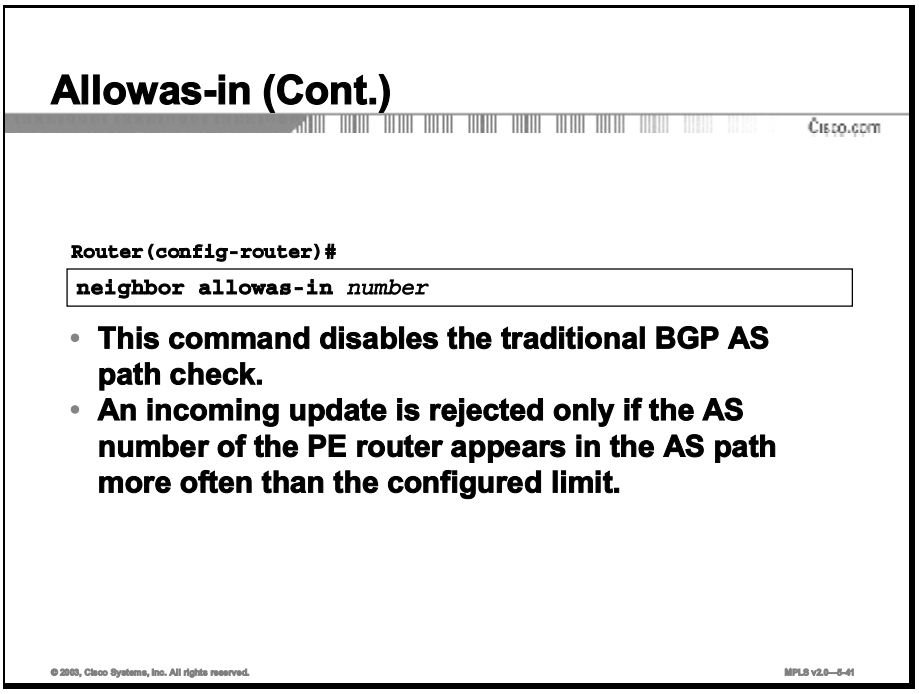

neighbor allowas-in

To configure PE routers to allow readvertisement of all prefixes containing duplicate AS numbers, use the **neighbor allowas-in** command in router configuration mode. To disable readvertisement of the AS number of a PE router, use the **no** form of this command:

- **neighbor allowas-in** *number*
- **no neighbor allowas-in** *number*

Syntax Description

Parameter	Description
number	Specifies the number of times to allow advertisement of the AS number of a PE router. Valid values are from 1 to 10 times.

Defaults

This command has no default behavior or values.

Implementing SOO for Loop Prevention

This topic describes the function of the Site of Origin (SOO) feature and identifies the command syntax to enable this feature.

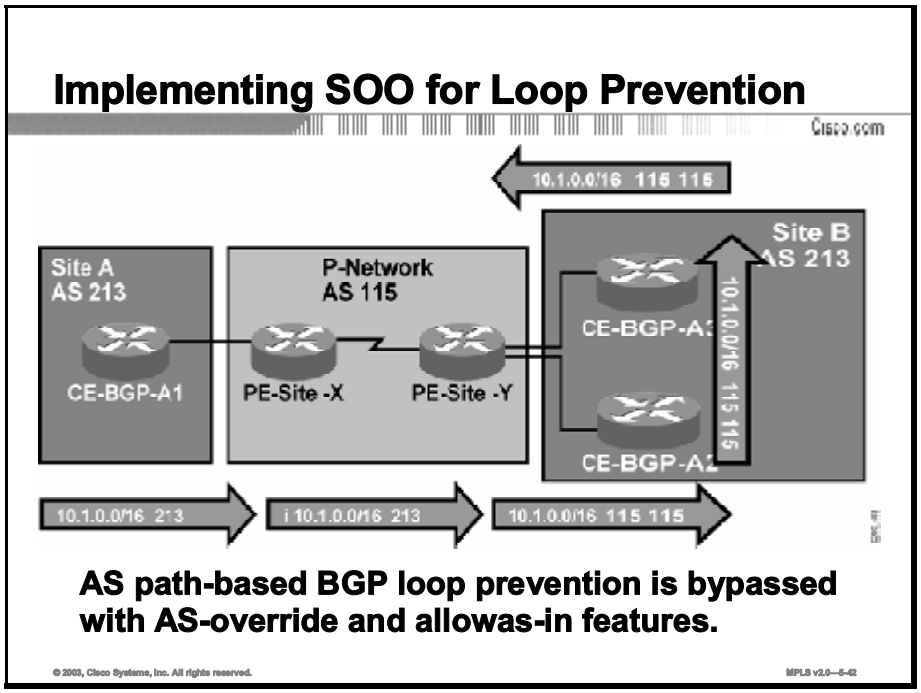

Most aspects of BGP loop prevention are bypassed when either the **AS-override** or the **allowas-in** feature is used. The routing information loops can still be detected by manually counting occurrences of an AS number in the AS path in an end-to-end BGP routing scenario then ensuring that the number field in the **neighbor allowas-in** command is set low enough to prevent loops.

This ability to still detect loops can present a particular problem when BGP is mixed with other PE-CE routing protocols. The SOO extended BGP community can be used as an additional loop prevention mechanism in these scenarios.

Note	SOO and any other loop prevention mechanisms are needed only for customer networks with multihomed sites. Loops can never occur in customer networks that have only stub sites.

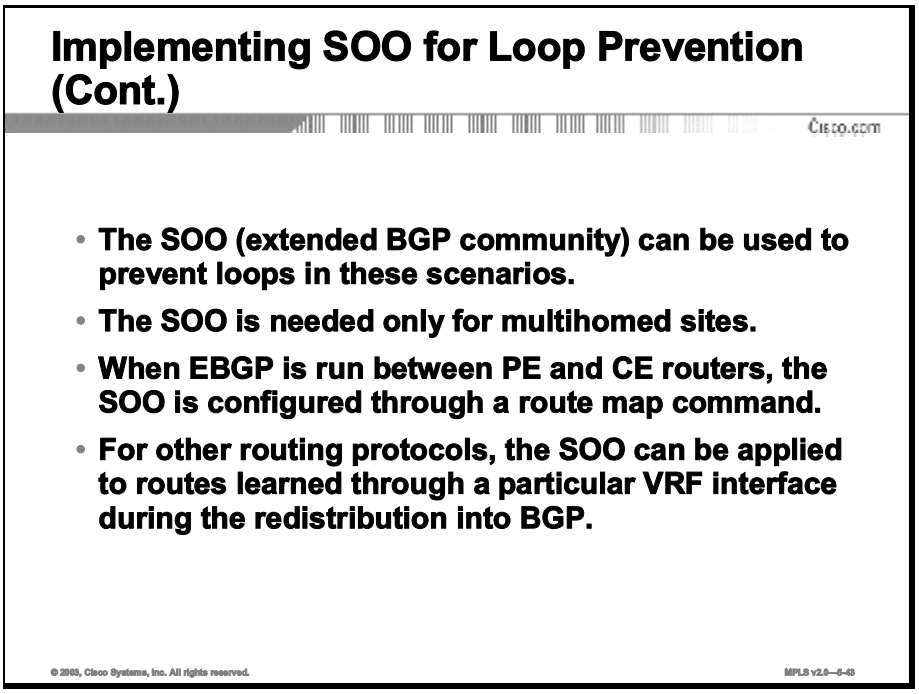

There are two ways to set the SOO attribute on a BGP route:

- For routes received from BGP-speaking CE routers, the SOO is configured by the incoming route map on the PE router.
- For all other routes, a route map setting the SOO is applied to the incoming interface. The SOO, as set by the route map, is attached to the BGP route when an IGP route received through that interface is redistributed into BGP.

Outgoing filters based on the SOO attribute also depend on the routing protocol used:

- Where EBGP is used as the PE-CE routing protocol, outbound route maps can be used on the PE router to deny routes matching particular SOO values.
- For all other routing protocols, filtering is performed on the basis of the SOO route map configured on the outgoing interface before the update is sent across that interface to the CE router.

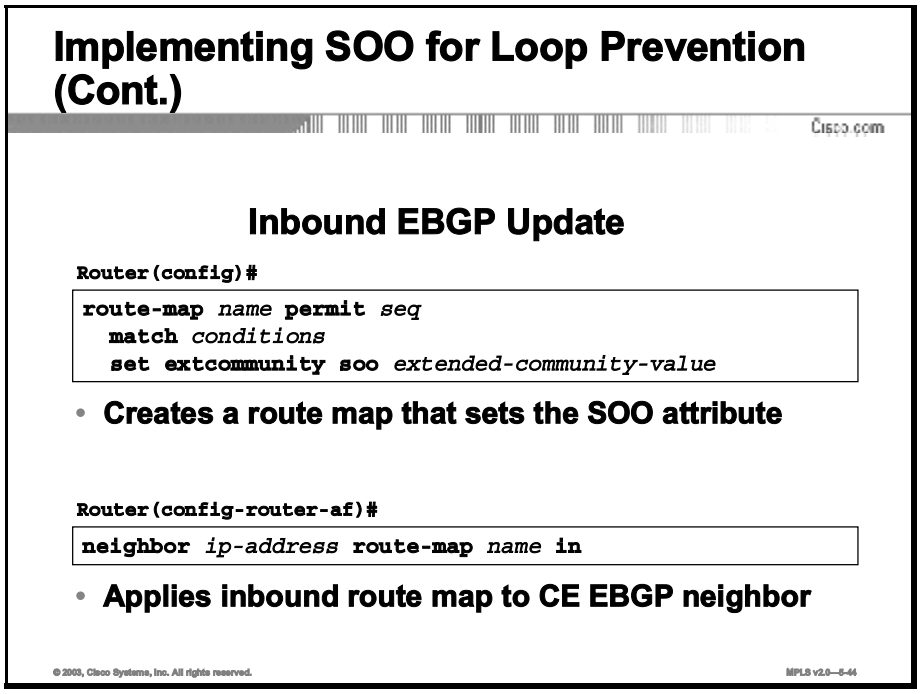

set extcommunity

To set the extended communities attribute, use the **set extcommunity** command in route map configuration mode. To delete the entry, use the **no** form of this command:

- **set extcommunity** {**rt** *extended-community-value* [**additive**] | **soo** *extended-community-value*}
- **no set extcommunity**
- **set extcommunity** *extcommunity-type community-number* [**additive**]
- **no set extcommunity** *extcommunity-type community-number* [**additive**]

Syntax Description

Parameter	Description
rt	Specifies the route target (RT) extended community attribute.
soo	Specifies the Site of Origin extended community attribute.
extended-community-value	Specifies the value to be set. The value can be one of the following combinations: ■ autonomous-system-number:network-number ■ *ip-address:network-number* The colon is used to separate the autonomous system number from the network number or the IP address from the network number.
additive	(Optional) Add space after the closing parenthesis. Adds the extended community to the already existing extended communities.

Defaults

No BGP extended community attributes are set by the route map.

neighbor route-map

To apply a route map to incoming or outgoing routes, use the **neighbor route-map** command in address family or router configuration mode. To remove a route map, use the **no** form of this command:

- **neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}
- **no neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}

Syntax Description

Parameter	Description
ip-address	IP address of the neighbor.
peer-group-name	Name of a BGP or MP-BGP peer group.
map-name	Name of a route map.
in	Applies route map to incoming routes.
out	Applies route map to outgoing routes.

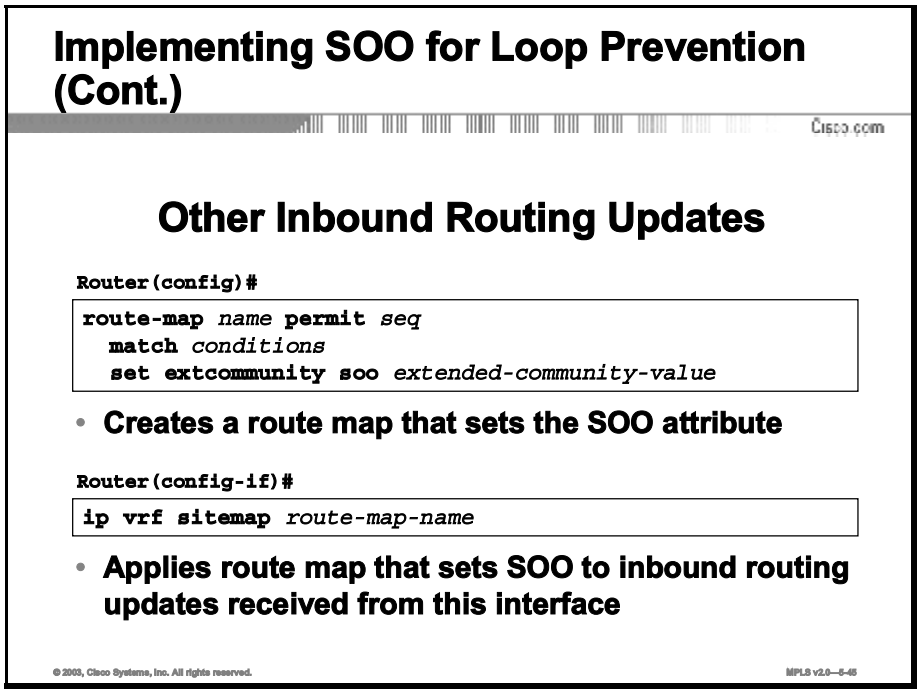

ip vrf sitemap

To set the SOO extended community attribute, use the **ip vrf sitemap** command in interface configuration mode. To delete the entry, use the **no** form of this command:

- **ip vrf sitemap** *route-map-name*
- **no ip vrf sitemap** *route-map-name*

Syntax Description

Parameter	Description
route-map-name	Sets the name of the route map to be used.

Defaults

No route map is used to set the SOO extended community.

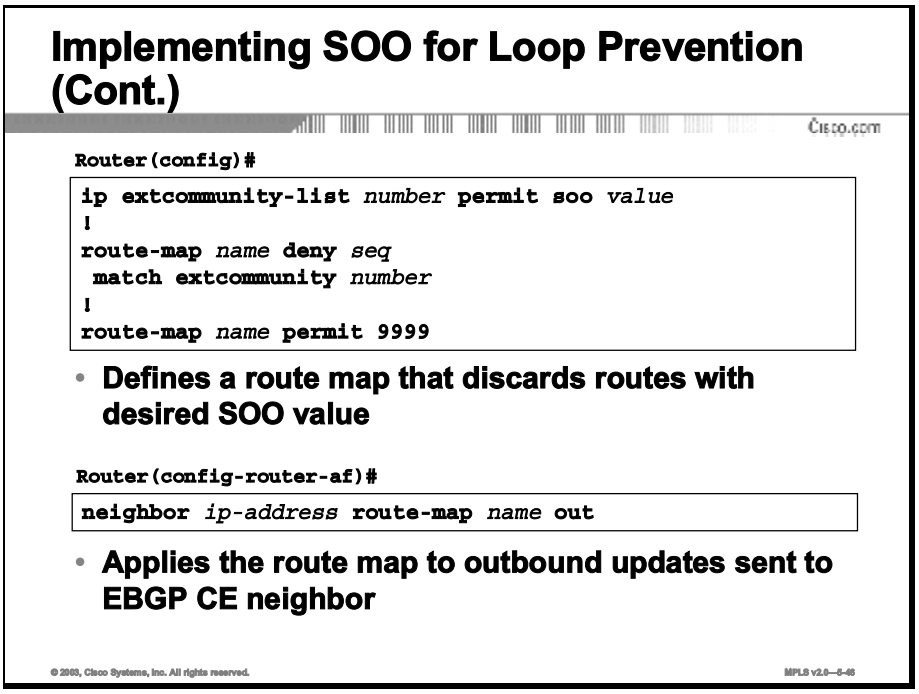

In this example, a route map matching a specific SOO value was defined using **ip extcommunity-list** to establish a SOO filter. The **route-map** command was used to define the route map based on the filter.

The newly defined route map is then applied to a BGP neighbor (CE router) on the PE router.

Summary

This topic summarizes the key points discussed in this lesson.

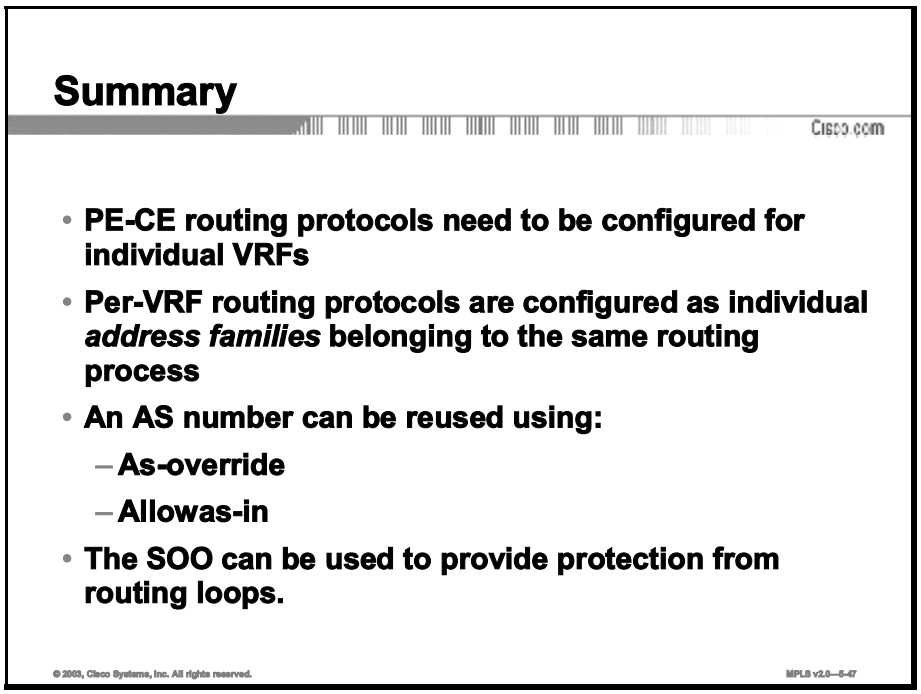

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Why do you need a VRF route limit command?

Q2) When would you need the AS-override feature?

Q3) How does the AS-override feature work?

Q4) When would you need the allowas-in feature?

Q5) When is it necessary to use the AS-override feature instead of the allowas-in feature?

Q6) How do you prevent BGP loops when using AS-override?

Q7) How do you prevent BGP loops when using allowas-in?

Q8) What is the Site of Origin?

Q9) When would you have to use the Site of Origin?

Q10) Where can you set the Site of Origin?

Quiz Answer Key

Q1) You need a VRF route limit command because of tight coupling of the customer and the service provider network in the MPLS VPN architecture. This tight coupling might also result in the service provider network being exposed to design and configuration errors in customer networks as well as to a number of new denial-of-service attacks based on routing protocol behavior.

Relates to: Limiting the Total Number of VRF Routes

Q2) When you need to connect two or more sites that use the same AS number via a VPN

Relates to: AS-Override

Q3) All leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Any other occurrences (farther down the AS path) of the AS number of the receiving router are not replaced because they indicate a real routing information loop.

Relates to: AS-Override

Q4) In some security-conscious implementations, customer VPNs are linked by a customer router that performs security functions such as access filtering or access logging.

Relates to: Allowas-in

Q5) In solutions where customer VPNs are linked by a customer router that do not support the AS-override feature

Relates to: Allowas-in

Q6) Only the leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Any other occurrences (farther down the AS path) of the AS number of the receiving router are not replaced because they indicate a real routing information loop.

Relates to: AS-Override

Q7) Allowas-in specifies the number of times to allow advertisement of an AS number of a PE router. Valid values are from 1 to 10 times using the *number* parameter of the **allowas-in** command.

Relates to: Allowas-in

Q8) The SOO is an extended BGP community that is used to indicate the site that has originated the routing update.

Relates to: Implementing SOO for Loop Prevention

Q9) It is used as an additional loop prevention mechanism in scenarios when the allowas-in feature is enabled.

Relates to: Implementing SOO for Loop Prevention

Q10) For routes received from BGP-speaking CE routers, the SOO is configured by the incoming route map on the PE router. For all other routes, a route map setting the SOO is applied to the incoming interface and the SOO is attached to the BGP route when an IGP route received through that interface is redistributed into BGP.

Relates to: Implementing SOO for Loop Prevention

Troubleshooting MPLS VPNs

Overview

This lesson explains the preliminary steps for troubleshooting an MPLS VPN. It also looks at routing information flow troubleshooting and VPN data flow troubleshooting.

Relevance

It is important to be able to determine what steps you should take when trying to solve a problem with your MPLS VPN network. This lesson looks at how to go about correcting MPLS VPN network problems.

Objectives

This lesson describes the process used to troubleshoot MPLS VPN implementation.

Upon completing this lesson, you will be able to:

- Identify the preliminary steps in MPLS VPN troubleshooting
- Identify the issues that you should consider when verifying the routing information flow in an MPLS VPN
- Describe the process used to validate CE-to-PE routing information flow
- Describe the process used to validate PE-to-PE routing information flow
- Describe the process used to validate PE-to-CE routing information flow
- Identify the issues that you should consider when verifying the data flow in an MPLS VPN
- Identify the issues that you should consider when validating CEF status
- Describe the process used to validate the end-to-end label switched path
- Describe the process used to validate the LFIB status

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementation, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Preliminary Steps in MPLS VPN Troubleshooting
- Verifying the Routing Information Flow
- Validating CE-to-PE Routing Information Flow
- Validating PE-to-PE Routing Information Flow
- Validating PE-to-CE Routing Information Flow
- Verifying the Data Flow
- Validating CEF Status
- Validating the End-to-End Label Switched Path
- Validating the LFIB Status
- Summary
- Quiz

Preliminary Steps in MPLS VPN Troubleshooting

This topic identifies the Preliminary steps in MPLS VPN troubleshooting.

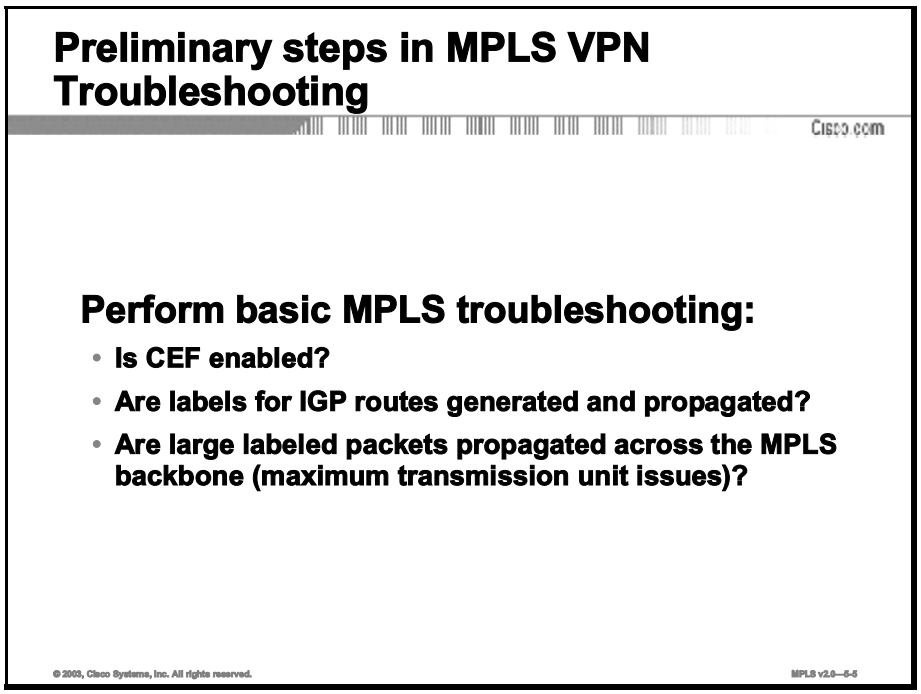

Before you start in-depth MPLS VPN troubleshooting, you should ask the following standard MPLS troubleshooting questions:

- Is CEF enabled on all routers in the transit path between the PE routers?
- Are labels for BGP next hops generated and propagated?
- Are there any maximum transmission unit (MTU) issues in the transit path (for example, LAN switches not supporting a jumbo Ethernet frame)?

MPLS VPN troubleshooting consists of two major steps:

- Verifying the routing information flow using the checks outlined in the figure
- Verifying the data flow, or packet forwarding

Verifying the Routing Information Flow

This topic identifies the issues that you should consider when verifying the routing information flow in an MPLS VPN.

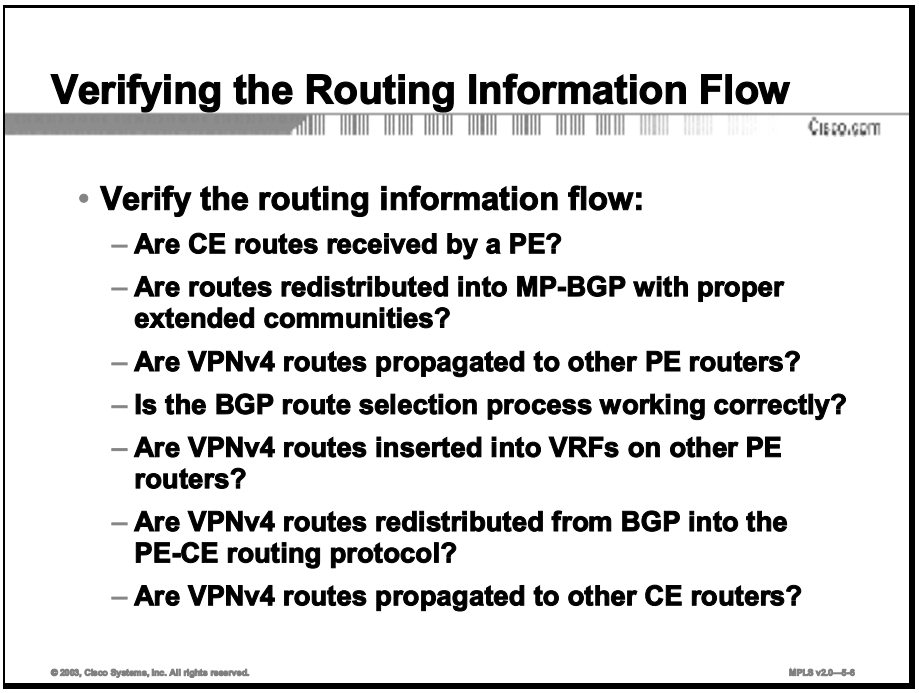

Verification of the routing information flow should be done systematically, starting at the ingress CE and moving to the egress CE.

Validating CE-to-PE Routing Information Flow

This topic describes the process that is used to validate CE-to-PE routing information flow.

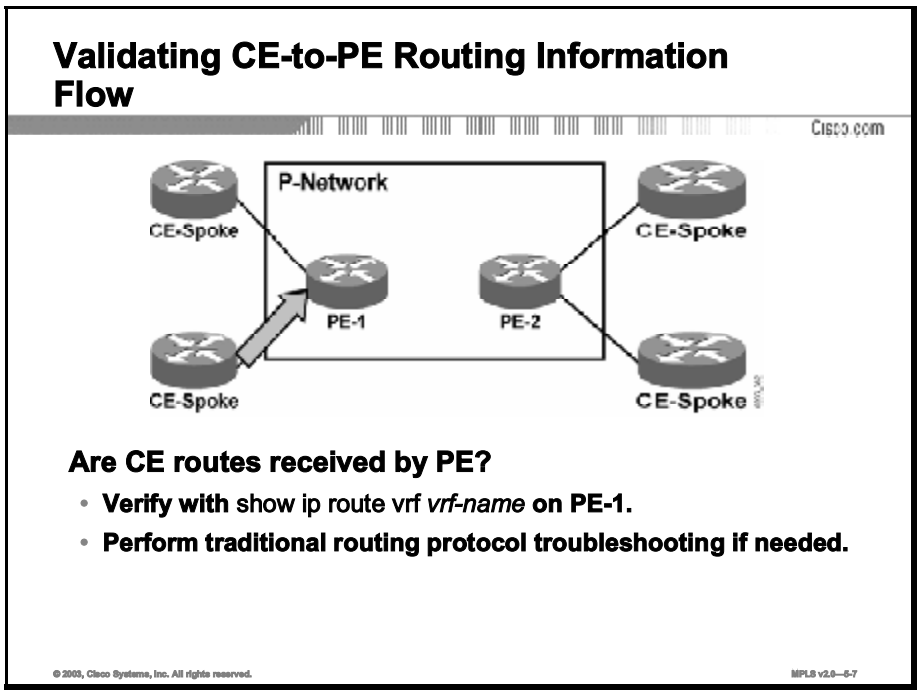

Routing information flow troubleshooting requires the verification of end-to-end routing information propagation between CE routers. The first step is to check the routing information exchange from CE routers to PE routers. Use the **show ip route vrf** *vrf-name* command to verify that the PE router receives customer routes from the CE router. Use traditional routing protocol troubleshooting if needed. Troubleshooting of standard enterprise routing protocols is described in the *Cisco Internetwork Troubleshooting* (CIT) course. BGP-specific troubleshooting is described in the individual modules of the *Configuring BGP on Cisco Routers* (BGP) course.

Validating PE-to-PE Routing Information Flow

Describe the process that is used to validate PE-to-PE routing information flow.

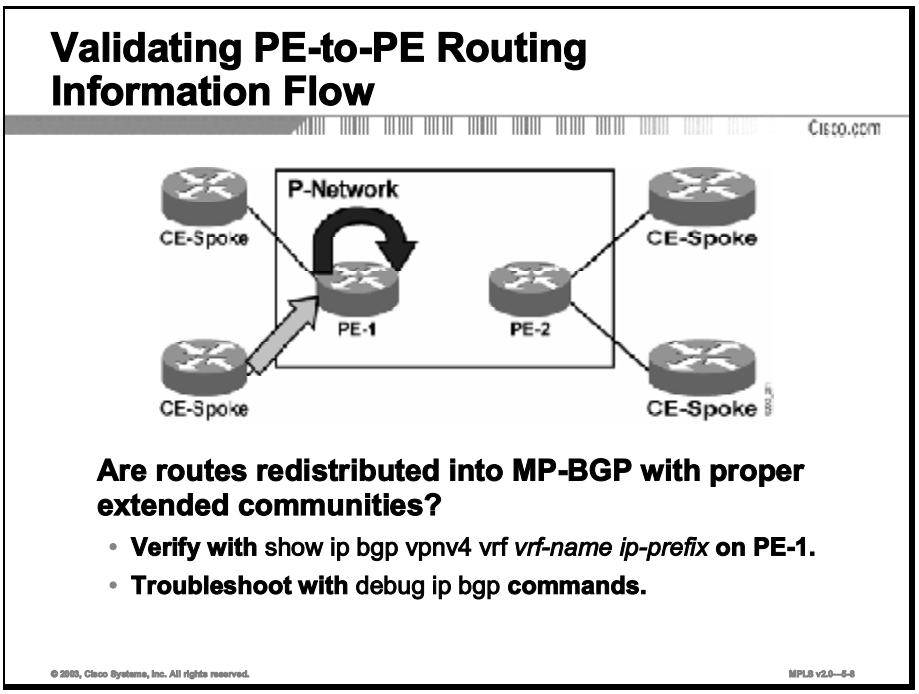

The CE routes received by the PE router need to be redistributed into MP-BGP; otherwise, they will not get propagated to other PE routers. Common configuration mistakes in this step include the following:

- Failing to configure redistribution between the PE-CE routing protocol and the per-VRF routing context of the BGP
- Using a route map on redistribution that filters CE routes

Proper redistribution of CE routes into a per-VRF instance of BGP can be verified with the **show ip bgp vpnv4 vrf** *vrf-name* command. The RD prepended to the IPv4 prefix and the RTs attached to the CE route can be verified with the **show ip bgp vpnv4 vrf** *vrf-name ip-prefix* command.

Validating PE-to-PE Routing Information Flow (Cont.)

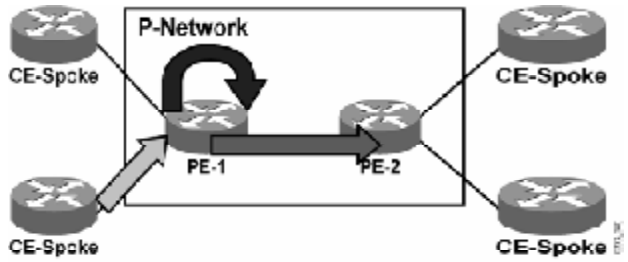

Are VPNv4 routes propagated to other PE routers?

- **Verify with** show ip bgp vpnv4 all *ip-prefix/length*.
- **Troubleshoot PE-to-PE connectivity with traditional BGP troubleshooting tools.**

MPLS v2.0—5-9

The CE routes redistributed into MP-BGP need to be propagated to other PE routers. Verify proper route propagation with the **show ip bgp vpnv4 all** *ip-prefix* command on the remote PE router.

Note	Routes sent by the originating PE router might not be received by a remote PE router because of automatic RT-based filters installed on the remote PE router.

Automatic route filters are based on RTs. Verify that the RTs attached to the CE route in the originating PE router match at least one of the RTs configured as import RTs in the VRF on the receiving PE router.

Validating PE-to-PE Routing Information Flow (Cont.)

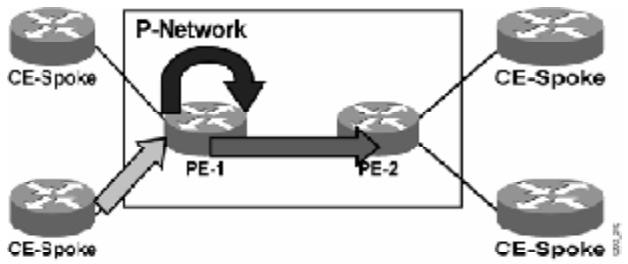

Is the BGP route selection process working correctly on PE-2?

- **Verify with** show ip bgp vpnv4 vrf *vrf-name ip-prefix*.
- **Change local preference or weight settings if needed.**
- **Do not change MED if you are using IGP-BGP redistribution on PE-2.**

In complex environments with multihomed customer sites, the BGP route selection process might affect proper MPLS VPN operation. Use standard BGP route selection tools (weights or local preference) to influence BGP route selection. The MED attribute should not be changed inside the MPLS VPN backbone if you plan to use two-way route redistribution between the PE-CE routing protocol and BGP.

Refer to the "BGP Filtering and Route Selection" module in the *Configuring BGP on Cisco Routers* (BGP) course for more information on BGP weights. Refer to the "Advanced BGP Configuration" module in the *Configuring BGP on Cisco Routers* (BGP) course for more information on BGP local preference and the MED attribute.

Validating PE-to-PE Routing Information Flow (Cont.)

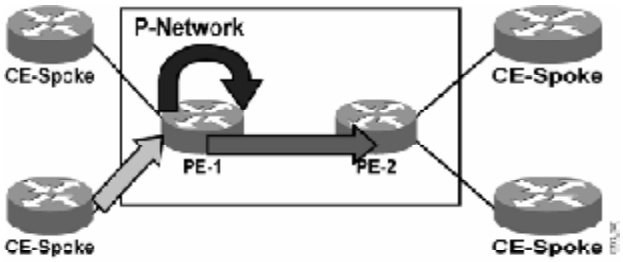

Are VPNv4 routes inserted into VRFs on PE-2?

- **Verify with** show ip route vrf.
- **Troubleshoot with** show ip bgp *ip-prefix* **and** show ip vrf detail.
- **Perform additional BGP troubleshooting if needed.**

The VPNv4 routes received by the PE router have to be inserted into the proper VRF. This insertion can be verified with the **show ip route vrf** command. Common configuration mistakes in this step include the following:

- The wrong import RTs are configured in the VRF.
- The route map configured as the import route map is rejecting the VPNv4 routes. Refer to the "Advanced VRF Import and Export Features" lesson in the "Complex MPLS VPNs" module for more information on import route maps.

The validity of the import RTs can be verified with the **show ip bgp vpnv4 all** *ip-prefix* command, which displays the RTs attached to a VPNv4 route, and with the **show ip vrf detail** command, which lists the import RTs for a VRF. At least one RT attached to the VPNv4 route needs to match at least one RT in the VRF.

Note	Be patient when troubleshooting this step. The import of VPNv4 routes into VRFs is not immediate and can take more than a minute in the worst circumstances.

Validating PE-to-PE Routing Information Flow (Cont.)

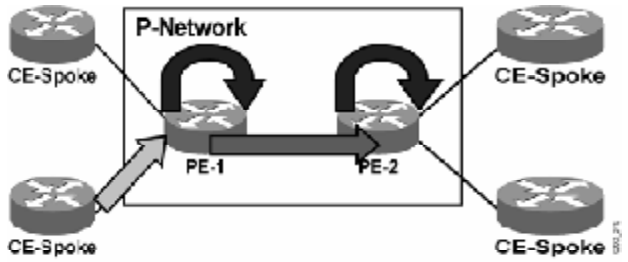

Are VPNv4 routes redistributed from BGP into the PE-CE routing protocol?

- **Verify redistribution configuration—is the IGP metric specified?**
- **Perform traditional routing protocol troubleshooting.**

Finally, the BGP routes received via MP-BGP and inserted into the VRF need to be redistributed into the PE-CE routing protocol. A number of common redistribution mistakes can occur here, starting with missing redistribution metrics.

Refer to the *Building Scalable Cisco Networks* (BSCN) and *Cisco Internetwork Troubleshooting* (CIT) courses for more information on route redistribution troubleshooting.

Validating PE-to-CE Routing Information Flow

This topic describes the process used to validate PE-to-CE routing information flow.

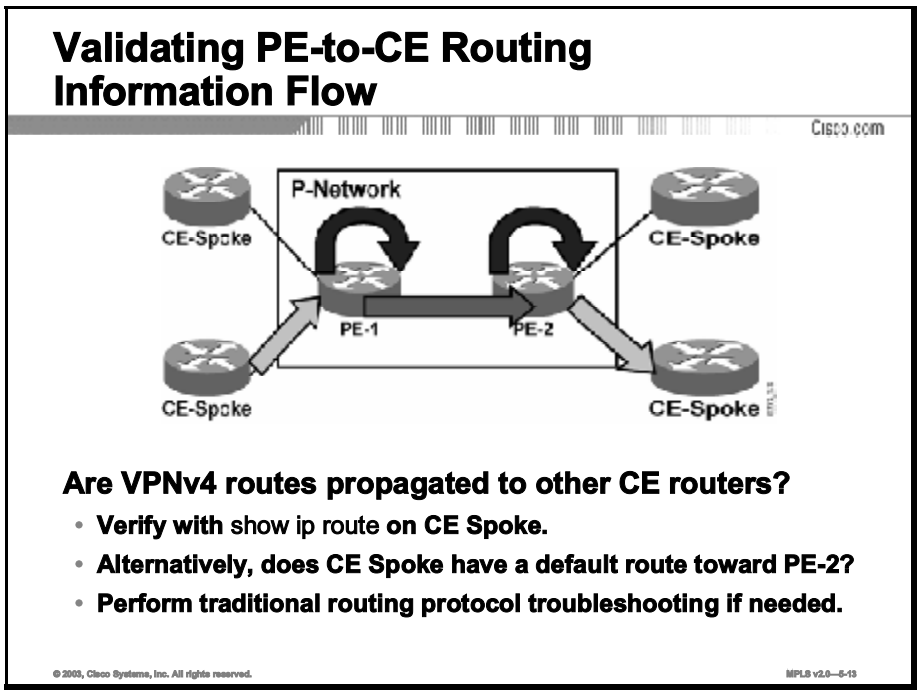

Last but not least, the routes redistributed into the PE-CE routing protocol have to be propagated to CE routers. You may also configure the CE routers with a default route toward the PE routers (see note). Use standard routing protocol troubleshooting techniques in this step.

Note	When using a default route on the CE routers, verify that the CE routers use classless routing configured with the **ip classless** command.

Verifying the Data Flow

This topic identifies the issues that you should consider when verifying the data flow in an MPLS VPN.

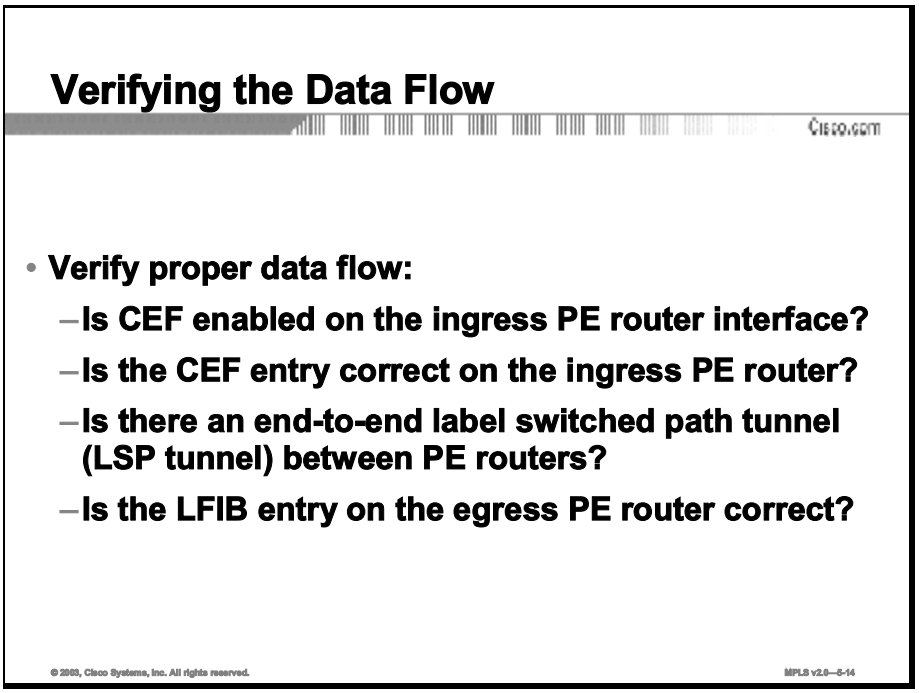

After you have verified proper route exchange, start MPLS VPN data flow troubleshooting using the checks listed in the following figures.

Validating CEF Status

This topic identifies the issues that you should consider when validating CEF status.

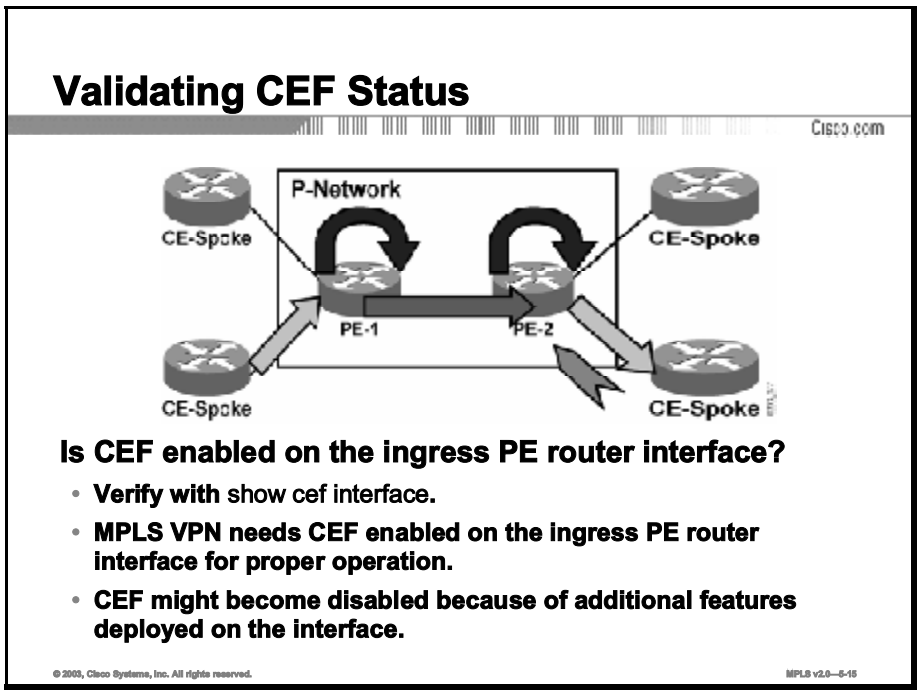

One of the most common configuration mistakes related to data flow is the failure to enable CEF in the ingress PE router interface. The presence of CEF can be verified with the **show cef interface** command. CEF is the only switching method that can perform per-VRF lookup and thus support MPLS VPN architecture.

Assuming that CEF is enabled on the router, there are three common reasons for CEF configuration mistakes:

- CEF is manually disabled on an interface.
- The interface is using an encapsulation method that is not supported by CEF, such as X.25 or Multilink PPP (MLP) with interleaving.
- Another feature has been configured on the interface that disables CEF (for example, IP precedence accounting).

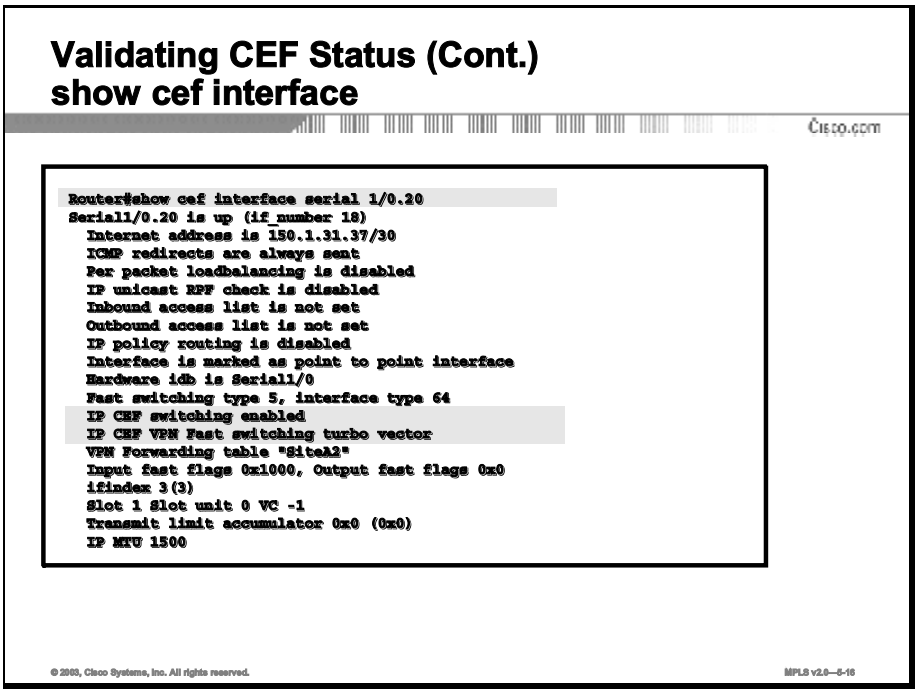

show cef interface

To display detailed CEF information for all interfaces, use the **show cef interface** command in EXEC mode:

- **show cef interface** *type number* [**statistics**][**detail**]

Syntax Description

Parameter	Description
type number	Interface type and number for displaying CEF information.
statistics	(Optional) Displays switching statistics for the line card.
detail	(Optional) Displays detailed CEF information for the specified interface type and 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.

The following table describes the fields shown in the output.

Parameter	Description
`interface type number is {up | down}`	Indicates status of the interface.
`Internet address`	Internet address of the interface.
`ICMP redirects are {always sent | never sent}`	Indicates how packet forwarding is configured.
`Per-packet load balancing`	Status of load balancing in use on the interface (enabled or disabled).
`IP unicast RPF check`	Indicates status of IP unicast Reverse Path Forwarding (RPF) check on the interface.
`Inbound access list {# | Not set}`	Number of access lists defined for the interface.
`Outbound access list`	Number of access lists defined for the interface.
`IP policy routing`	Indicates the status of IP policy routing on the interface.
`Hardware idb is type number`	Interface type and number configured.
`Fast switching type`	Used for troubleshooting; indicates switching mode in use.
`IP CEF switching {enabled | disabled}`	Indicates the switching path used.
`Slot n Slot unit n`	Slot number.
`Hardware transmit queue`	Indicates the number of packets in the transmit queue.
`Transmit limit accumulator`	Indicates the maximum number of packets allowed in the transmit queue.
`IP MTU`	Value of the MTU size set on the interface.

Validating CEF Status (Cont.)

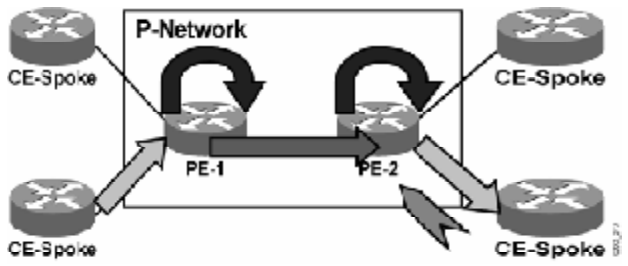

Is the CEF entry correct on the ingress PE router?

- **Display the CEF entry with** show ip cef vrf *vrf-name ip-prefix/length* detail**.**
- **Verify the label stack in the CEF entry.**

If CEF switching is enabled on the ingress interface, you can verify the validity of the CEF entry and the associated label stack with the **show ip cef vrf** *vrf-name ip-prefix* **detail** command. The top label in the stack should correspond to the BGP next-hop label as displayed by the **show tag forwarding** command on the ingress router. The second label in the stack should correspond to the label allocated by the egress router. You can verify this by using the **show tag forwarding** command on the egress router.

Validating the End-to-End Label Switched Path

This topic describes the process that is used to validate the end-to-end label switched path (LSP).

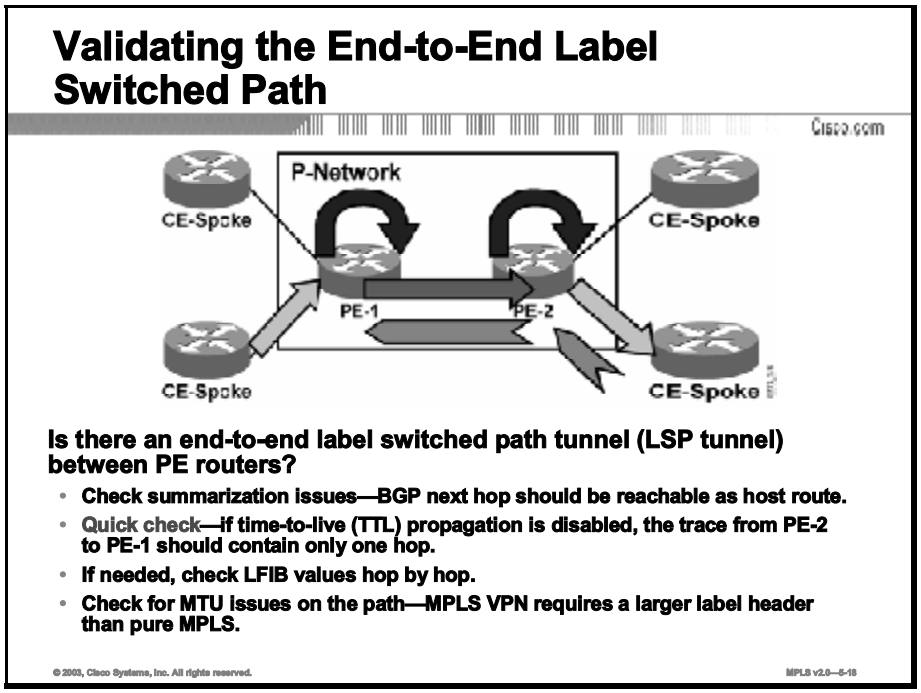

If CEF is enabled on the ingress interface and the CEF entry contains proper labels, the data flow problem might lie inside the MPLS core. Two common mistakes include summarization of BGP next hops inside the core IGP and MTU issues.

The quickest way to diagnose summarization problems is to disable IP time-to-live (TTL) propagation into the MPLS label header using the **no mpls ip ttl-propagate** command. The **traceroute** command toward the BGP next hop should display no intermediate hops when TTL propagation is disabled. If intermediate hops are displayed, the LSP tunnel between PE routers is broken at those hops and the VPN traffic cannot flow.

Validating the LFIB Status

This topic describes the process that is used to validate the LFIB status.

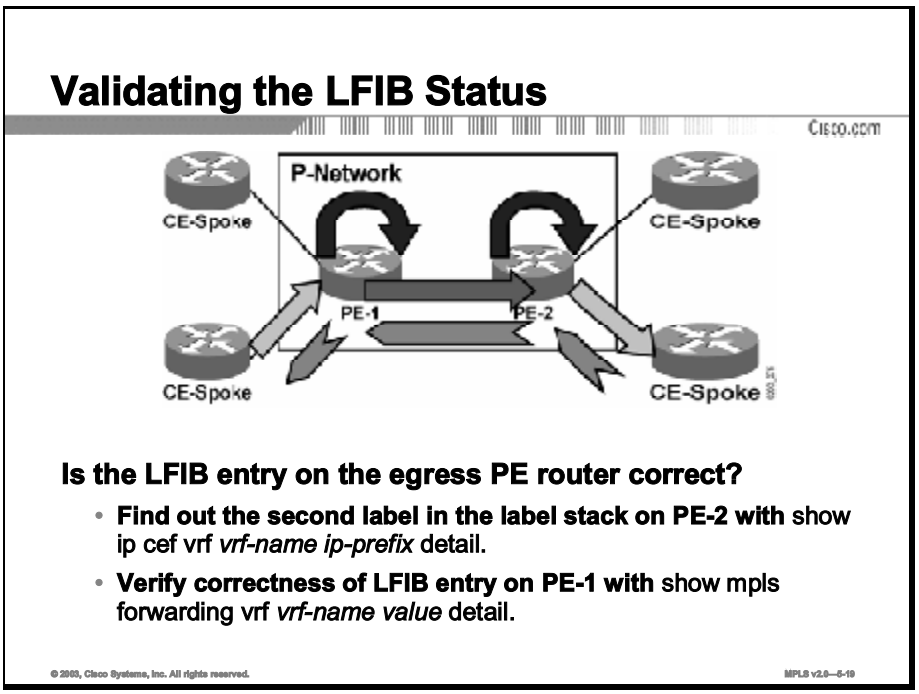

As a last troubleshooting measure (usually not needed), you can verify the contents of the LFIB on the egress PE router and compare them with the second label in the label stack on the ingress PE router. A mismatch indicates an internal Cisco IOS software error that you will need to report to the Cisco Technical Assistance Center (TAC).

Summary

This topic summarizes the key points discussed in this lesson.

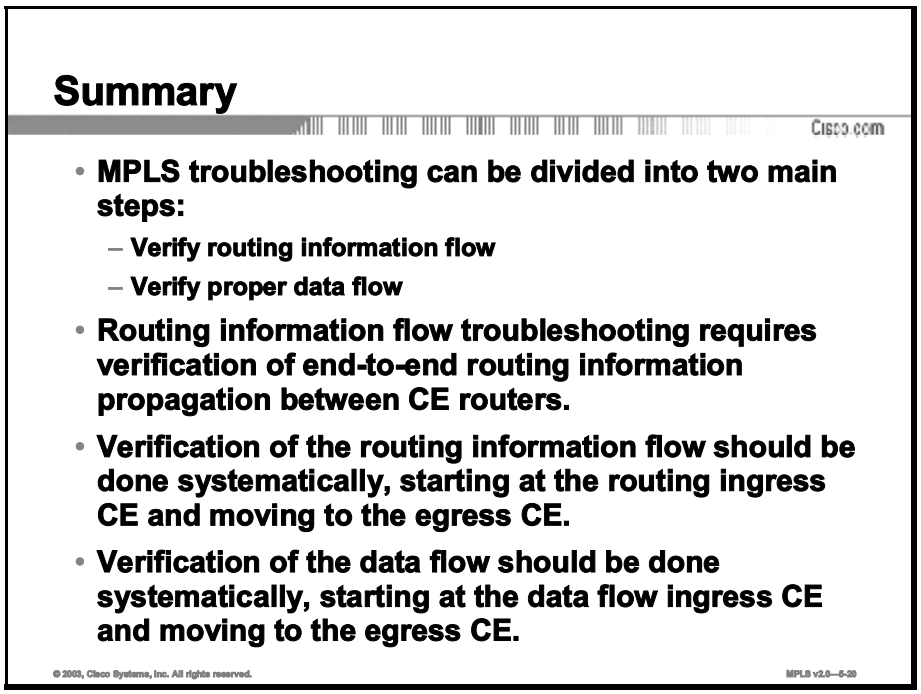

References

For additional information, refer to these resources:

- Access www.cisco.com for additional information about VPNs.
- "BGP Filtering and Route Selection" module in *the Configuring BGP on Cisco Routers* (BGP) course
- "Advanced BGP Configuration" module in the *Configuring BGP on Cisco Routers* (BGP) course
- *Building Scalable Cisco Networks* (BSCN)
- *Cisco Internetwork Troubleshooting* (CIT) course

Next Steps

For the associated lab exercise, refer to these sections of the course Lab Guide:

- Lab Exercise 5-1: Initial MPLS VPN Setup
- Lab Exercise 5-2: Running EIGRP Between PE and CE Routers
- Lab Exercise 5-3: Running OSPF Between PE and CE Routers
- Lab Exercise 5-4: Running BGP Between PE and CE Routers

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are the preliminary MPLS VPN troubleshooting steps?

Q2) Which command do you use to verify that the PE router is receiving customer routes from the CE router?

Q3) How do you verify routing information exchange between PE routers?

Q4) How do you verify redistribution of VPNv4 routes into the PE-CE routing protocol?

Q5) How do you test end-to-end data flow between PE routers?

Q6) How do you verify that the PE router ingress interface supports CEF switching?

Q7) How do you verify that there is an end-to-end LSP?

Q8) How do you verify that the LFIB entry on the egress PE router is correct?

Quiz Answer Key

Q1) Is CEF enabled?
Are labels for IGP routes generated and propagated?
Are large labeled packets propagated across the MPLS backbone (MTU issues)?

Relates to: Preliminary Steps in MPLS VPN Troubleshooting

Q2) **show ip route vrf** *vrf-name*

Relates to: Validating CE-to-PE Routing Information Flow

Q3) Use the **show ip bgp vpnv4 all** *ip-prefix/length* command to verify proper route propagation.

Relates to: Validating PE-to-PE Routing Information Flow

Q4) Use the **show ip bgp vrf** *vrf-name ip-prefix* command on the egress PE router or use **show ip route** on the egress CE router.

Relates to: Validating PE-to-CE Routing Information Flow

Q5) From the ingress PE router, use the **ping vrf** *vrf-name* command to ping the interface that supports the CE router.

Relates to: Verifying the Data Flow

Q6) Use the **show cef interface** command.

Relates to: Validating CEF Status

Q7) Check LFIB values hop by hop or use the **trace vrf** *vrf-name* command from the ingress PE router.

Relates to: Validating the End-to-End Label Switched Path

Q8) Find out the second label in the label stack on the ingress PE with the **show ip cef vrf** *vrf-name ip-prefix* **detail** command. Verify the correctness of the LFIB entry on the egress PE with the **show mpls forwarding vrf** *vrf-name value* **detail** command.

Relates to: Validating the LFIB Status

Module 6

Complex MPLS VPNs

Overview

This module discusses some advanced configuration features that can help increase the stability of the Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) backbone. The module also discusses various MPLS VPN features that a service provider might use to help meet service requirements, and looks at various types of VPN solutions and topologies.

Module Objectives

Upon completing this module, you will be able to use diagrams of a typical MPLS VPN solution to identify the Cisco IOS command syntax that is required to successfully configure VPN operations to support the overlapping model. You will also be able to describe how this model can be used to implement managed services and Internet access. This ability includes being able to do the following:

- Identify the command syntax that is required to configure advanced VRF import and export features
- Identify the characteristics of overlapping VPNs
- Identify the characteristics of the central services VPN solutions
- Identify the characteristics of the managed CE router service
- Identify and describe the Cisco MPLS VPN managed services

Module Outline

The module contains these lessons:

- Advanced VRF Import and Export Features
- Overlapping VPNs
- Central Services VPNs
- Managed CE Routers Service
- MPLS Managed Services

Advanced VRF Import and Export Features

Overview

Some virtual routing and forwarding (VRF) features allow you to be more selective with routes, either by specifying which routes will or will not be added. You may also limit the number of routes that a customer can insert into the virtual routing and forwarding instance (VRF). This lesson presents the command syntax that is used to limit each type of route, and shows configuration examples of these commands.

Relevance

It is important to understand how to fine-tune the MPLS VPN parameters that will enhance operation of the network. Customer service-level agreements (SLAs) should be adhered to so that they provide the best possible solutions for each specific customer. This lesson looks at some key areas regarding the use of VRF import and export features.

Objectives

This lesson identifies the command syntax that is required to configure advanced VRF import and export features.

Upon completing this lesson, you will be able to:

- Identify advanced VRF features and their usage
- Identify the command syntax that is required to configure selective VRF imports
- Identify the command syntax that is required to configure selective VRF exports

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementations, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Advanced VRF Features
- Configuring Selective VRF Import
- Configuring Selective VRF Export
- Summary
- Quiz

Advanced VRF Features

This topic identifies the advanced VRF features and their usage.

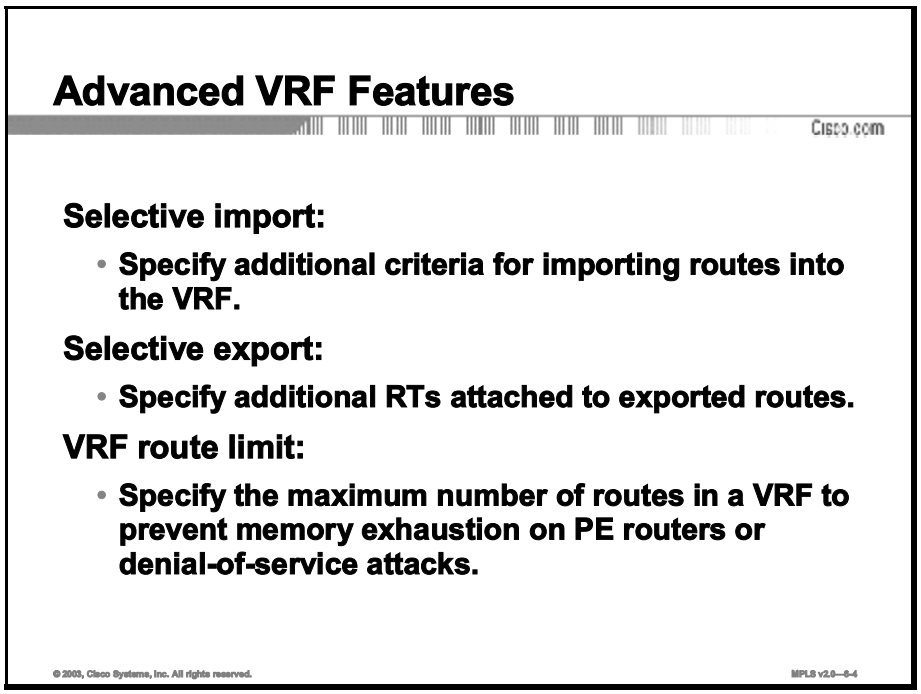

These advanced VRF features allow you to deploy advanced MPLS VPN topologies or increase the stability of the MPLS VPN backbone:

- The “selective import” feature allows you to select routes to be imported into a VRF based on criteria other than the RT of the VRF
- The “selective export” feature allows you to attach specific RTs to a subset of routes exported from a VRF. By default, the same RTs get attached to all exported routes.
- The “VRF route limit” feature allows you to limit the number of routes that the customer (or other PE routers) can insert in the VRF. This feature prevents undesirable consequences like configuration errors or denial-of-service attacks.

Configuring Selective VRF Import

This topic presents the command syntax that is required to configure selective VRF import.

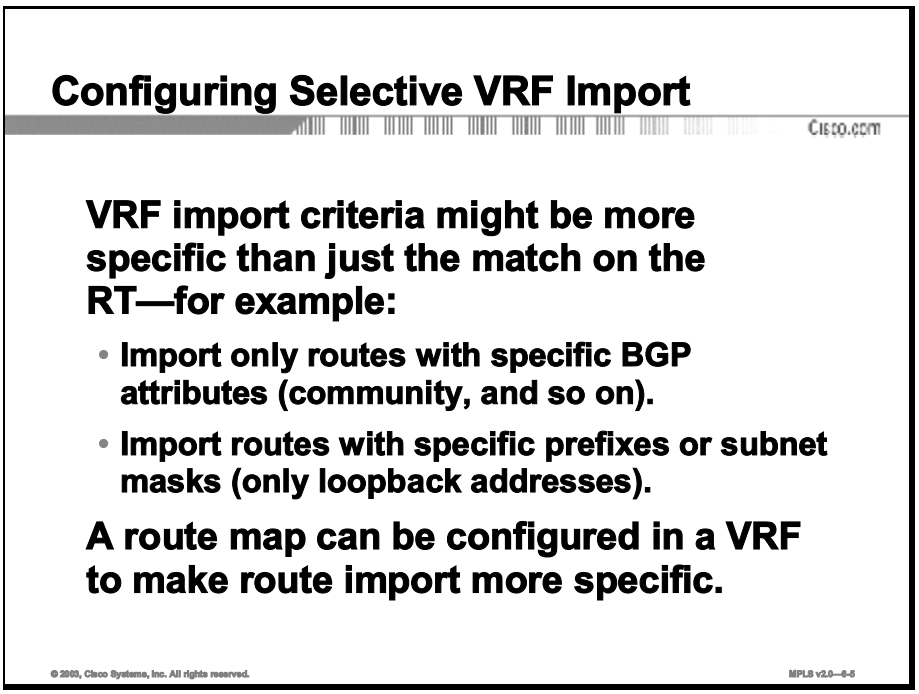

Selective route import into a VRF allows you to narrow the route import criteria. It uses a route map that can filter the routes selected by the route target (RT) import filter. The routes imported into a VRF are Border Gateway Protocol (BGP) routes, so you can use match conditions in a route map to match any BGP attribute of a route. These attributes include communities, local preference, multi-exit discriminator (MED), autonomous system (AS) path, and so on.

The import route map filter is combined with the RT import filter. A route has to pass the RT import filter first and then the import route map. The necessary conditions for a route to be imported into a VRF are as follows:

- At least one of the RTs attached to the route matches one of the import RTs configured in the VRF.
- The route is permitted by the import route map.

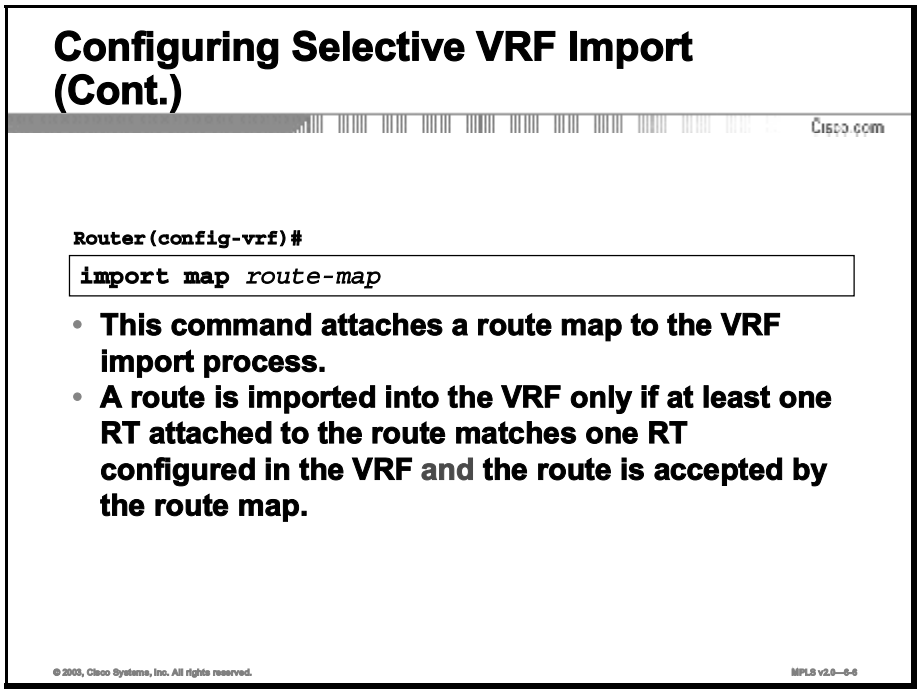

import map

To configure an import route map for a VRF, use the **import map** command in VRF configuration submode.

- **import map** *route-map*

Syntax Description

Parameter	Description
route-map	Specifies the route map to be used as an import route map for the VRF.

Defaults

There is no default. A VRF has no import route map unless one is configured using the **import map** command.

Configuring Selective VRF Import (Cont.)

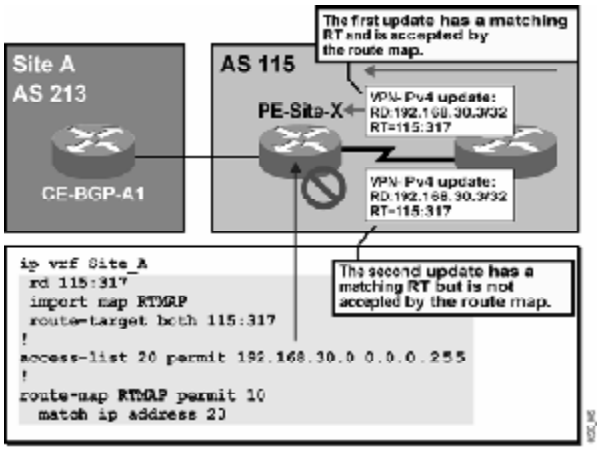

The figure shows an example in which an import route map is used to match the IP version 4 (IPv4) portion of incoming VPN IPv4 (VPNv4) routes and import into the VRF only routes matching a certain prefix. A configuration similar to this one could be used to accomplish the following:

- Deploy advanced MPLS VPN topologies (for example, a managed router services topology that is detailed later in this module)
- Increase the security of an extranet VPN by allowing only predefined subnets to be inserted into a VRF, thus preventing an extranet site from inserting unapproved subnets into the extranet

Note A similar function is usually not needed in an intranet scenario because all customer routers in an intranet are usually under common administration.

Configuring Selective VRF Export

This topic identifies the command syntax that is required to configure selective VRF export.

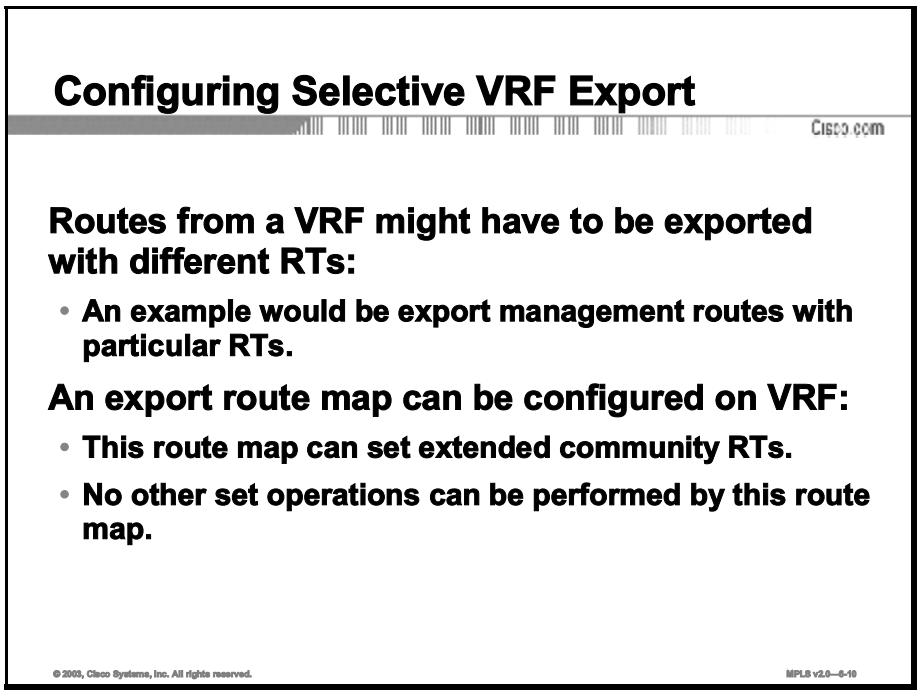

Some advanced MPLS VPN topologies are easiest to implement if you can attach a variety of RTs to routes exported from the same VRF. This capability allows only a subset of the routes exported from a VRF to be imported into another VRF. Most services in which customer routers need to connect to a common server (for example, network management stations, voice gateways, and application servers) fall into this category.

The export route map function provides exactly this functionality. A route map can be specified for each VRF to attach additional RTs to routes exported from that VRF. The export route map performs only the attachment of RTs. It does not perform any filtering function.

Attributes attached to a route with an export route map are combined with the export RT attributes. If you specify export RTs in a VRF and set RTs with an export route map, all specified RTs will be attached to the exported route.

Note	The export route map provides functionality almost identical to that of the import route map, but applied to a different VRF. Any requirement that can be implemented with an export route map can also be implemented with an import route map However, the implementation of export maps can be more complicated and harder to manage

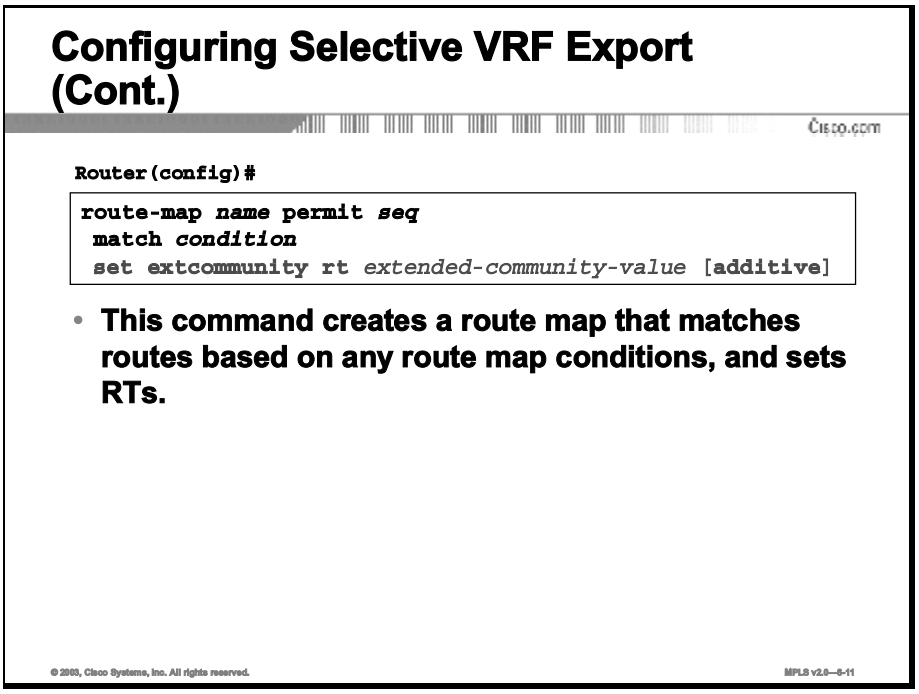

set extcommunity

To set the BGP extended communities attribute, use the **set extcommunity** command in route-map configuration mode. To delete the entry, use the **no** form of this command:

- **set extcommunity** {**rt** *extended-community-value* [**additive**] | **soo** *extended-community-value*}
- **no set extcommunity** {**rt** *extended-community-value* [**additive**] | **soo** *extended-community-value*}

Syntax Description

Parameter	Description
rt	Specifies the RT extended community attribute.
soo	Specifies the SOO extended community attribute.
extended-community-value	Specifies the value to be set. The value can be one of the following combinations: ■ autonomous-system-number: network-number ■ *ip-address: network-number* The colon is used to separate the AS number from the network number or the IP address from the network number.
additive	(Optional) Adds an RT to the existing RT list without replacing any RTs.

Defaults

No BGP extended community attributes are set by the route map.

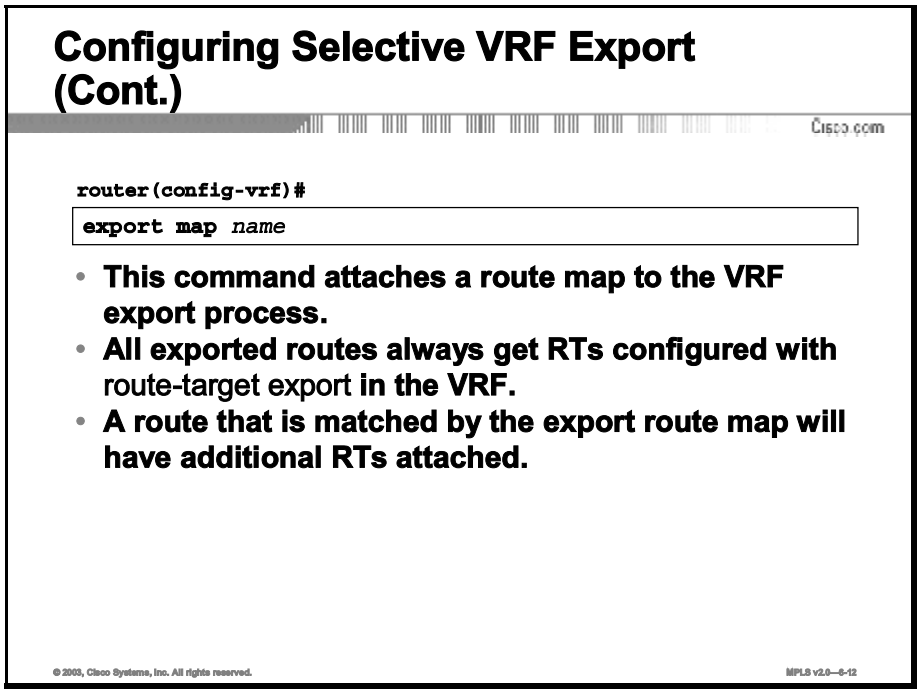

export map

To apply a route map in order to filter and modify exported routes, use the **export map** command in VRF configuration mode. To remove the route map from the VRF, use the **no** form of this command:

- **export map** *route-map*
- **no export map** *route-map*

Syntax Description

Parameter	Description
route-map	Specifies the name of the route map to be used.

Defaults

No route map is used.

Configuring Selective VRF Export (Cont.)

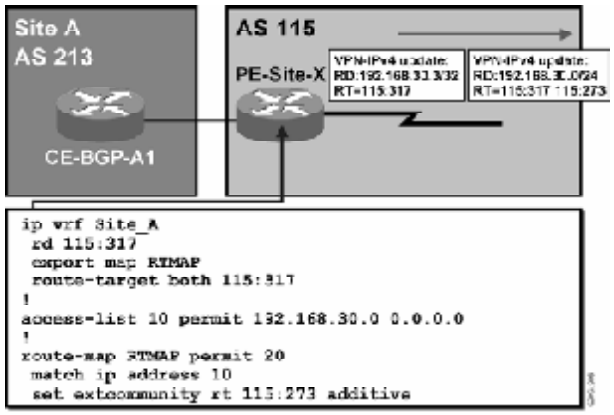

The example here mirrors the earlier example in this lesson. This time the configuration is implemented with an export route map. In the earlier example, selective import of routes into a VRF was achieved with an import route map in the receiving VRF that allowed only routes from a certain address block to be inserted into the VRF. In this example, routes from a certain address block are marked with an additional RT in the originating VRF and are automatically inserted into the receiving VRF on the basis of their RT.

The main difference between import and export route maps is therefore the deployment point:

- The import route map is deployed in the receiving VRF.
- The export route map is deployed in the originating VRF.
- Based on the network design, one or the other functionality might be preferred.

Note Import and export route maps can increase the number of routes processed by a router. The BGP maximum-prefix function can be used to ensure the number of routes does not exceed the network design. (See the "Configuring BGP as the Routing Protocol Between PE and CE Routers" lesson in the "MPLS VPN Implementation" module for further details.)

Summary

This topic summarizes the key points discussed in this lesson.

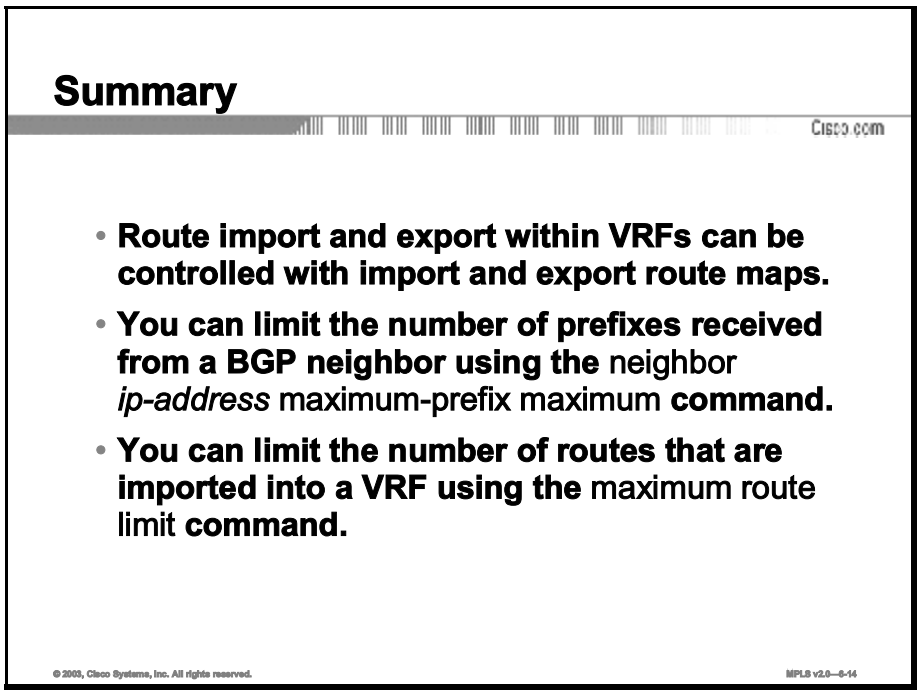

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Why do you need a selective VRF import command?

Q2) How does the import route map affect the VRF import process?

Q3) Why do you need a selective VRF export command?

Q4) How does the export route map affect the VRF export process?

Q5) Which BGP attributes can be set with an export route map?

Quiz Answer Key

Q1) It allows you to select routes to be imported into a VRF based on criteria other than the VRF RT.

Relates to: Configuring Selective VRF Import

Q2) The import route map filter is combined with the RT import filter—a route has to pass the RT import filter first and then the import route map to be imported into the VRF.

Relates to: Configuring Selective VRF Import

Q3) It allows you to attach specific RTs to a subset of routes exported from a VRF (by default, the same RTs get attached to all exported routes).

Relates to: Configuring Selective VRF Export

Q4) A route map can be specified for each VRF to attach additional RTs to routes exported from a VRF. The export route map performs only the attachment of RTs; it does not perform any filtering function, and you cannot change any other route attributes with it.

Relates to: Configuring Selective VRF Export

Q5) extended community RTs

Relates to: Configuring Selective VRF Export

Overlapping VPNs

Overview

Overlapping VPNs are usually used to connect parts of two separate VPNs. A third VPN is created within the MPLS VPN network that contains sites from both VPNs. A new RT extended community is used for networks originating in the sites that are also in the new VPN. This action may require a new VRF, resulting in a new route distinguisher (RD). Networks originating in these sites are exported with two RT extended communities: one for the original VPN and one for the overlapping VPN. This lesson looks at the requirements, usage, and solutions associated with overlapping VPNs.

Relevance

It is important to understand customer needs and how to best fit those needs. This lesson looks at an area that helps to clarify some solutions regarding multiple separate VPNs.

Objectives

This lesson identifies the characteristics of overlapping VPNs.

Upon completing this lesson, you will be able to:

- Identify the participants in overlay VPNs
- Identify typical overlapping VPN usages
- Describe the routing update flow in an overlapping VPN
- Describe the data flow in an overlapping VPN
- Identify the tasks that are required to enable an overlapping VPN
- Identify the command syntax that is required to enable overlapping VPN VRFs

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementations, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Overlapping VPNs
- Typical Overlapping VPN Usages
- Overlapping VPN Routing
- Overlapping VPN Data Flow
- Overlapping VPNs—Configuration Tasks
- Configuring Overlapping VPN VRFs
- Summary
- Quiz

Overlapping VPNs

This topic identifies the participants in overlapping VPNs.

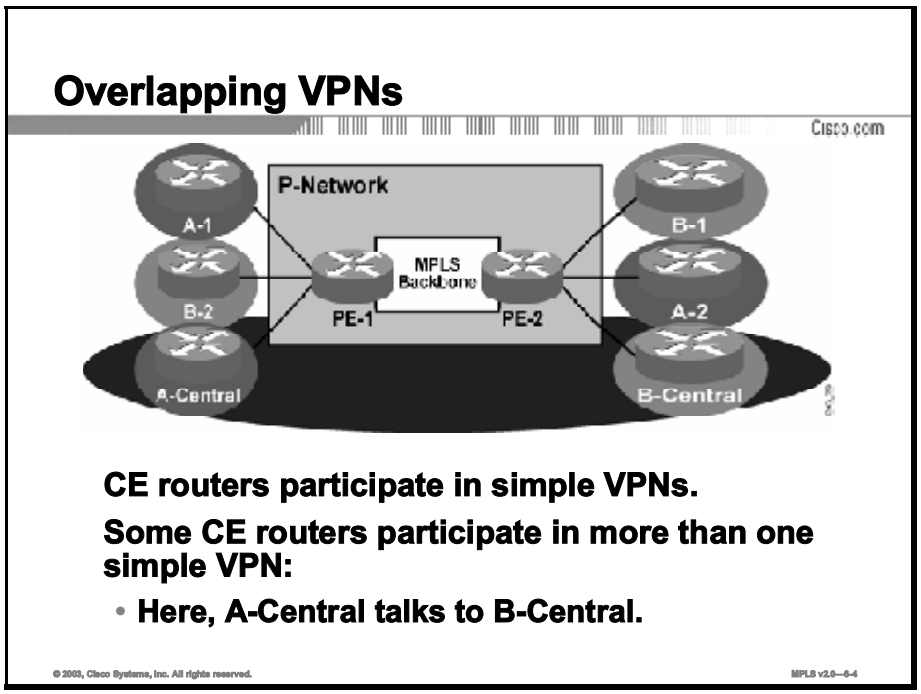

When two VPN customers want to share some information, they may decide to interconnect their central sites. To achieve this, two simple VPNs are created, each containing a customer central site and its remote sites. Then a third VPN that partially overlaps with the customer VPNs but connects only their central sites is created . The central sites can talk to each other. They can also talk to the remote sites in their simple VPN, but not to the remote sites belonging to the other customer simple VPN. The addresses used in the central sites, however, have to be unique in both VPNs.

Another option is to use dual Network Address Translation (NAT) with a registered address to be imported and exported between the two central sites.

Typical Overlapping VPN Usages

This topic identifies typical overlapping VPN usages.

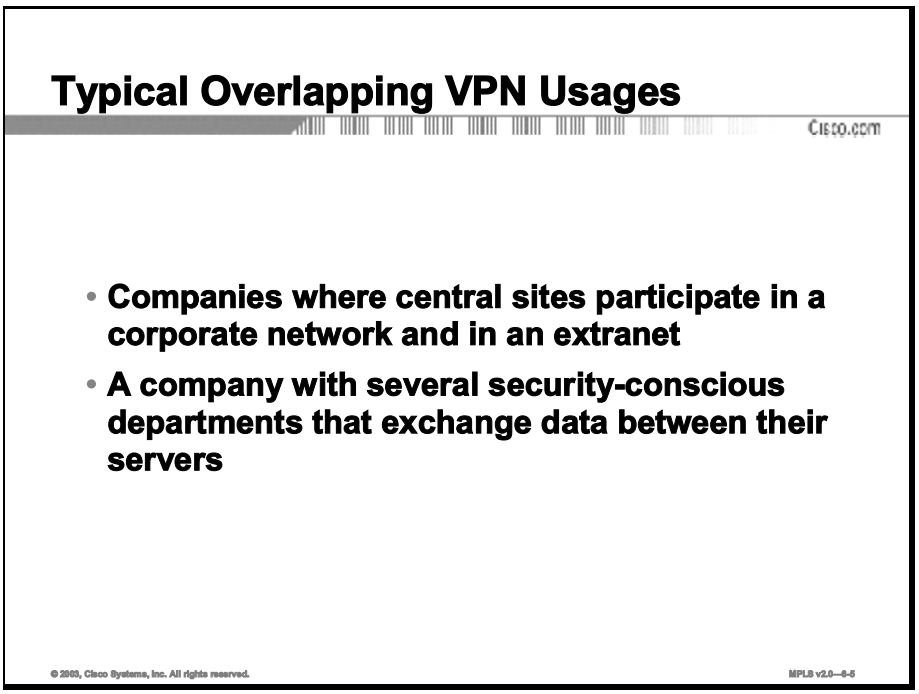

There are two typical uses for overlapping VPNs:

- Companies that use MPLS VPNs to implement both intranet and extranet services might use overlapping VPNs. In this scenario, each company participating in the extranet VPN would probably deploy a security mechanism on its customer edge (CE) routers to prevent other companies participating in the VPN from gaining access to other sites in the customer VPN.
- A security-conscious company might decide to limit visibility between different departments in the organization. Overlapping VPNs might be used as a solution in this case.

Note	Security issues might force an enterprise network to be migrated to an MPLS VPN even if it is not using MPLS VPN services from a service provider.

Overlapping VPN Routing

This topic describes the routing update flow in an overlapping VPN.

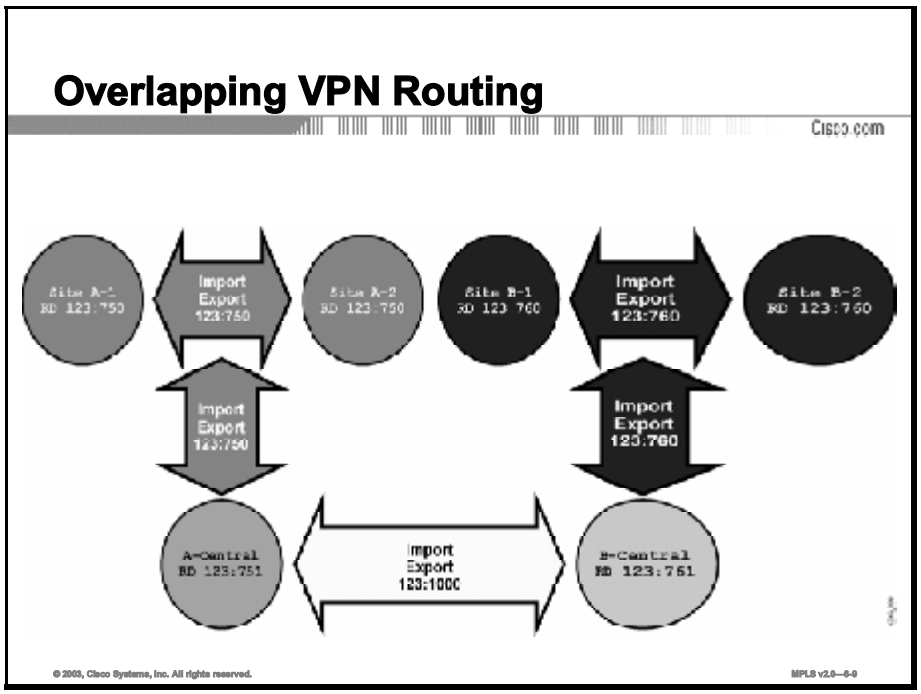

The figure shows how to implement overlapping VPNs:

- Each VPN has its own RT (123:750, 123:760) that the sites participating in the VPN import and export.
- Sites that participate in more than one VPN import routes with RTs from any VPN in which they participate and export routes with RTs for all VPNs in which they participate.

Site A-1 and Site A-2 (participating only in VPN-A):

- Export all networks with RT 123:750
- Import all networks that carry RT 123:750 (VPN-A)

Site B-1 and Site B-2 (participating only in VPN-B):

- Export all networks with RT 123:760
- Import all networks that carry RT 123:760 (VPN-B)

Site A-Central (participating in VPN-A and the overlapping VPN):

- Exports all networks with RTs 123:750 *and* 100:1000
- Imports all networks that carry RT 123:750 (VPN-A) *or* 100:1000 (overlapping VPN)

Site B-Central (participating in VPN-B and the overlapping VPN):

- Exports all networks with RTs 123:760 *and* 100:1000
- Imports all networks that carry RT 123:760 (VPN-B) *or* 100:1000 (Overlapping VPN)

Overlapping VPN Data Flow

This topic describes the data flow in an overlapping VPN.

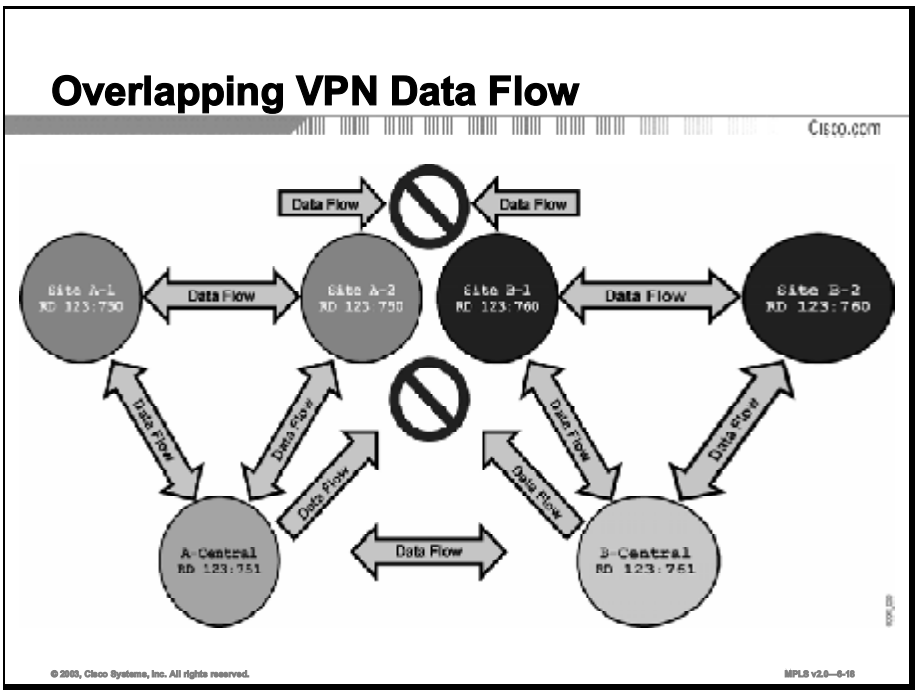

Because sites belonging to different VPNs do not share routing information, they cannot talk to each other.

- The simple VPN for customer A contains routes that originate from the following:
  - A-Central site
  - A remotes
- The simple VPN for customer A contains routes that originate from the following:
  - B-Central site
  - B remotes
- The overlapping VPN .contains routes that originate from the following:
  - A-Central site
  - B-Central site.
- All of the customer A sites can communicate with each other
- All of the customer B sites can communicate with each other
- A-Central and B-Central can communicate with each other
- The customer A remote site cannot communicate with the customer B remote sites

Note If a site participating in more than one VPN is propagating a default route to other sites, it can attract traffic from those sites and start acting as a transit site between VPNs, enabling sites that were not supposed to communicate to establish two-way communication.

Overlapping VPNs—Configuration Tasks

This topic identifies the tasks that are required to enable an overlapping VPN.

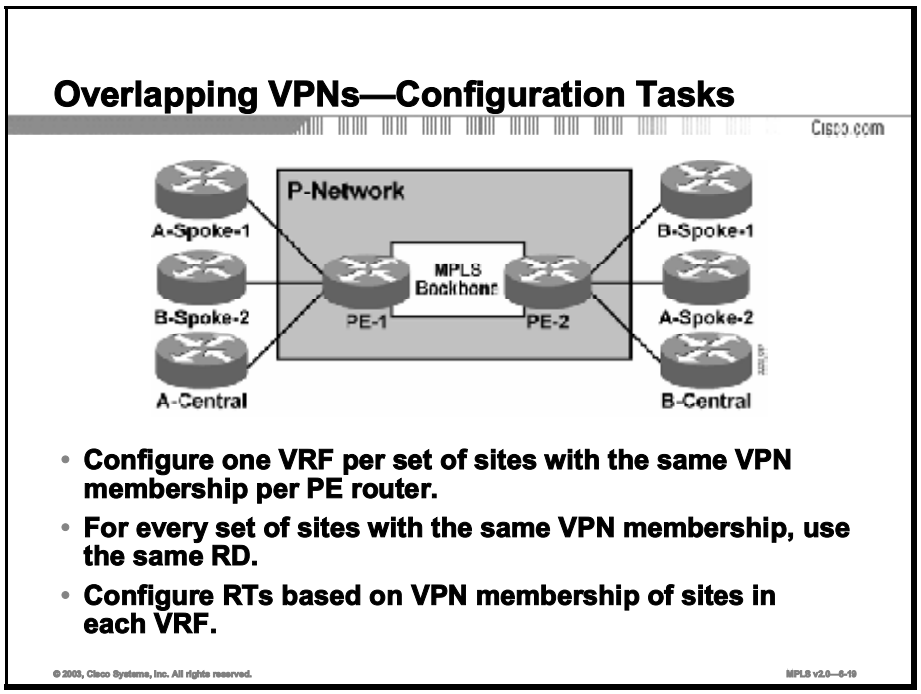

The figure shows four types of sites with different VPN memberships. This situation requires at least four VRFs:

- A-Spoke-1 and A-Spoke-2 are members of VPN-A only. (They need two VRFs because they are not connected to the same PE router; they can, however, use the same RD.)
- B-Spoke-1 and B-Spoke-2 are members of VPN-B only. (They need two VRFs because they are not connected to the same PE router; they can, however, use the same RD.)
- A-Central is a member of VPN-A and overlapping VPN-AB. (They need an additional RD.)
- B-Central is a member of VPN-B and overlapping VPN-AB. (They cannot use the same RD as A-Central because B-Central has different routing requirements from A-Central.)

The following table shows an RT and RD numbering scheme for PE-1.

Table 1: PE-1 RT and RD Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-A-Central	123:751	123:750 123:1001	123:750 123:1001

The following table shows an RT and RD numbering scheme for PE-2.

Table 2: PE-2 RT and RD Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-B-Central	123:761	123:760 123:1001	123:760 123:1001

Configuring Overlapping VPN VRFs

This topic identifies the command syntax that is required to enable overlapping VPN VRFs.

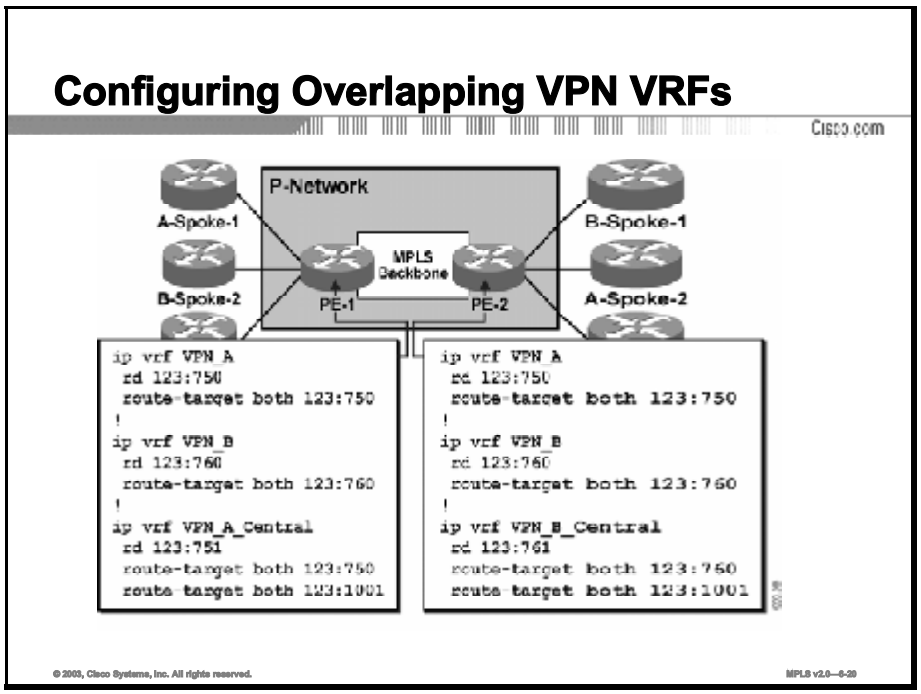

The Cisco IOS software configuration for PE-1 and PE-2 reflects the RT and RD numbering scheme from Tables 1 and 2. The example shows only VRF configuration and does not show VPN routing or Multiprotocol Border Gateway Protocol (MP-BGP) routing between the provider edge (PE) routers.

Summary

This topic summarizes the key points discussed in this lesson.

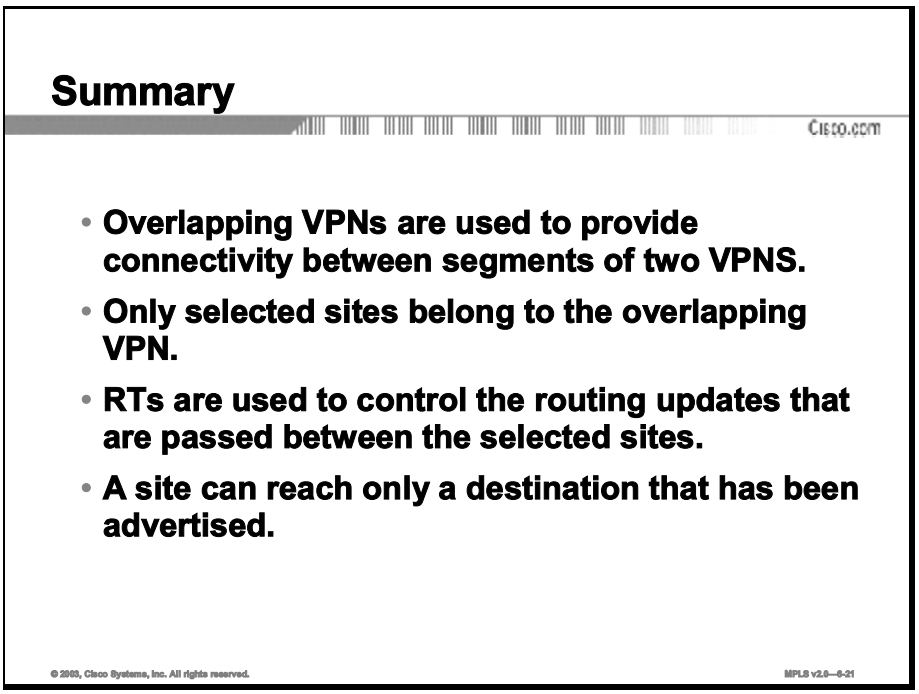

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Who are the typical users of overlapping VPNs?

Q2) What are the connectivity requirements for overlapping VPNs?

Q3) What is the expected data flow within an overlapping VPN?

Q4) How many VRFs do you need at most to implement three partially overlapping VPNs? How many route distinguishers? How many route targets?

Quiz Answer Key

Q1) companies that use MPLS VPNs to implement both intranet and extranet services, or a security-conscious company that wishes to limit visibility between different departments in the organization

Relates to: Typical Overlapping VPN Usages

Q2) Selected sites in a VPN can communicate only with sites within their VPN. Other selected sites can communicate with sites in their VPN and selected sites in a second VPN.

Relates to: Overlapping VPN Routing; Overlapping VPN Data Flow

Q3) Selected sites in a VPN can communicate only with sites within their VPN. Other selected sites can communicate with sites in their VPN and selected sites in a second VPN.

Relates to: Overlapping VPN Routing; Overlapping VPN Data Flow

Q4) VRFs – 3
Route distinguishers – 4
Route targets - 3

Relates to: Overlapping VPNs—Configuration Tasks; Configuring Overlapping VPN VRFs

Central Services VPNs

Overview

A central services VPN is used when more VPNs need to share a common set of servers. These servers reside in the central services VPN, and all other VPNs have access to this VPN. Those VPNs, however, are not able to see one another. The central services VPN is implemented using two RT extended communities, where one imports networks into the VPN and the other exports networks. The client sites do the opposite. Two RT extended communities are needed to prevent client sites from exchanging routing information. This lesson looks at central services VPN solution topologies and how routing updates within that topology would flow, and discusses the implications of combining a central services VPN with a simple customer VPN.

Relevance

It is important to fully understand the topologies that make the most sense for the customer, and to be able to configure or recommend alternative options. This lesson looks at one of those solution topologies.

Objectives

This lesson identifies the characteristics of the central services VPN solution.

Upon completing this lesson, you will be able to:

- Describe the access characteristics of a central services VPN
- Describe the routing characteristics of a central services VPN
- Describe the data flow within a central services VPN
- Identify the steps to configuring a central services VPN
- Identify the command syntax that is required to configure a central services VPN
- Identify the connectivity requirements when you are integrating a central services VPN with a simple VPN
- Identify the RD requirements when you are integrating a central services VPN with a simple VPN

- Identify the RT requirements when you are integrating a central services VPN with a simple VPN
- Identify the command syntax that is required when you are integrating a central services VPN with a simple VPN

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementations, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Central Services VPN
- Central Services VPN Routing
- Central Services VPN Data Flow Model
- Steps to Configuring a Central Services VPN
- Configuring a Central Services VPN
- Central Services VPN and Simple VPN Requirements
- Configuring RDs in a Central Services and Simple VPN
- Configuring RTs in a Central Services and Simple VPN
- Configuring VRFs in a Central Services and Simple VPN
- Summary
- Quiz

Central Services VPN

This topic describes the access characteristics of a central services VPN.

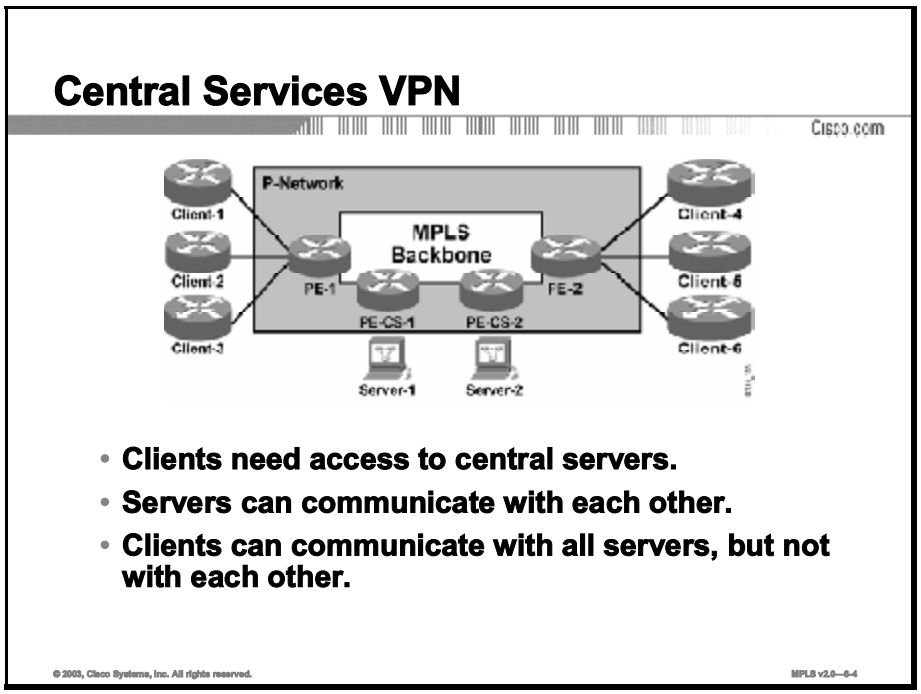

A central services VPN is a topology with the following characteristics:

- Some sites (“server sites”) can communicate with all other sites
- All the other sites (“client sites”) can communicate only with the server sites

This topology can be used in the following situations:

- The service provider offers services to all customers by allowing them access to a common VPN.
- Two (or more) companies want to exchange information by sharing a common set of servers.
- A security-conscious company separates its departments and allows them access only to common servers.

Central Services VPN Routing

This topic describes the routing characteristics of a central services VPN.

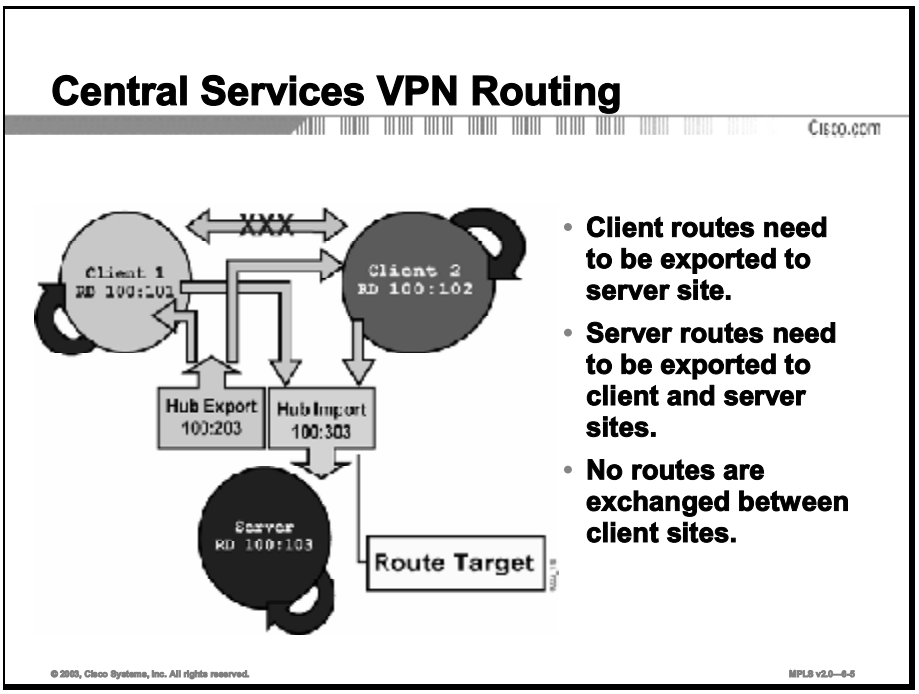

The figure illustrates the MPLS VPN routing model that is used to implement a central services VPN:

- Client 1 and Client 2 have their own RTs (100:101, 100:102) that they import and export; they also export networks with RT 100:303 and import networks with RT 100:203.

Note Client-specific RTs were introduced to comply with the implementation requirements of Cisco IOS Release 12.0 T, in which each VRF has to have at least one of its export RTs configured as its import RT.

- The central site imports and exports networks with the RT of its VPN, but it also imports networks with RT 100:303 and exports networks with RT 100:203.

Client 1:

- Exports all networks with RTs 100:101 *and* 100:303
- Imports all networks that carry RT 100:101 *or* 100:203

Client 2:

- Exports all networks with RTs 100:102 *and* 100:303
- Imports all networks that carry RT 100:102 *or* 100:203

Central site:

- Exports all networks with RT 100:203
- Imports all networks that carry RT 100:303

Central Services VPN Data Flow Model

This topic describes the data flow within a central services VPN.

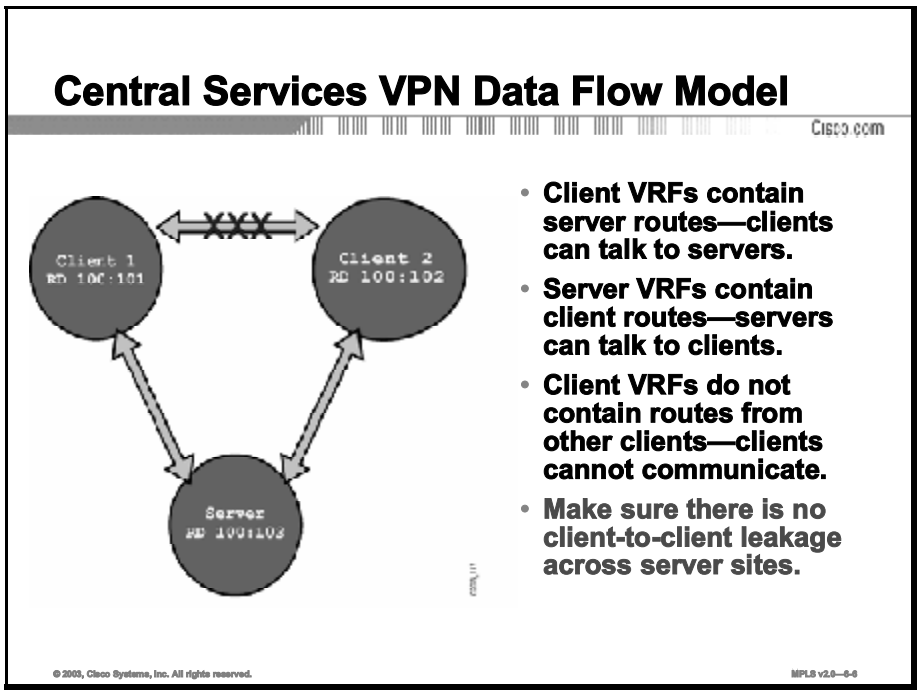

In the central services VPN topology, the client VRF contains only routes from the client site and routes from the server sites. This setup precludes the client sites from communicating with other client sites.

A server VRF in this topology contains routes from the site or sites attached to the VRF, as well as routes from all other client and server sites. Hosts in server sites can therefore communicate with hosts in all other sites.

Note	If the central site is propagating a default route to other sites, it can result in client sites seeing each other through the CE router in the central site.

Steps to Configuring a Central Services VPN

This topic identifies the steps to configuring a central services VPN.

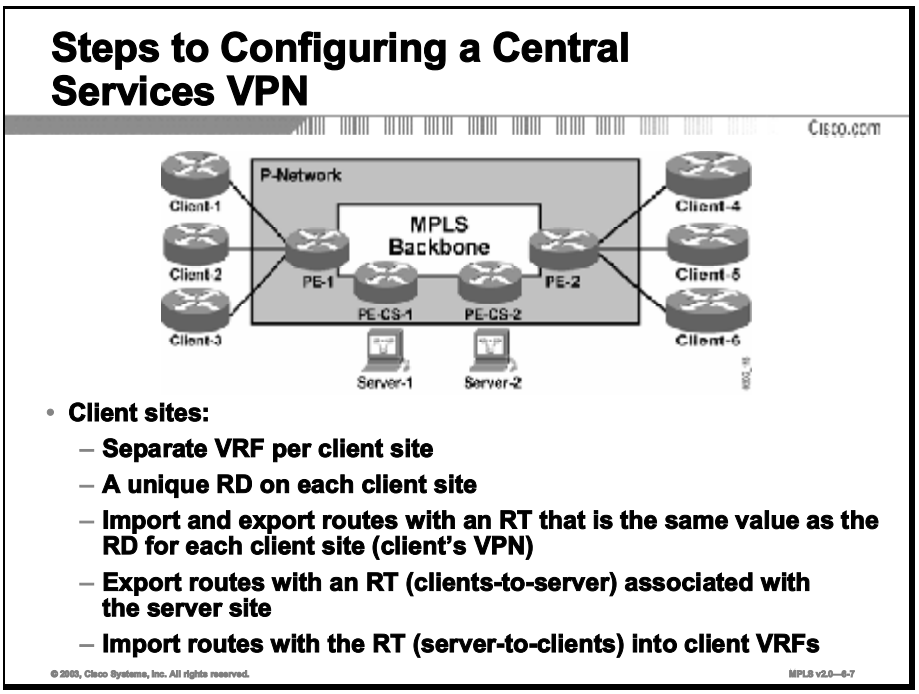

To configure a central services VPN, you need to address the following requirements:

- You need a separate VRF for each client.
- You need one VRF per PE router connecting a server site.
- You need a unique RD on each client site.
- You need a unique RD on each set of server sites.
- You need an import-export RT with the same value as the RD, for each client site.

Steps to Configuring a Central Services VPN (Cont.)

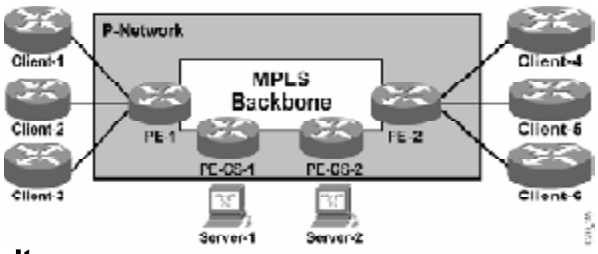

- **Server sites:**
  - **One VRF for each different service type**
  - **Unique RD on each different service type**
  - **Import and export routes with an RT that is the same value as the RD for each server site (server's VPN)**
  - **Export server site routes with an RT (server-to-client)**
  - **Import routes with the RT (clients-to-server) into the server VRFs**

This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Client-1	123:101	123:101	123:101
		123:203	123:303
Client-2	123:102	123:102	123:102
		123:203	123:303

This table shows an RD and RT numbering scheme for PE-2.

PE-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Client-4	123:111	123:111	123:111
		123:203	123:303
Client-5	123:112	123:112	123:112
		123:203	123:303

This table shows an RD and RT numbering scheme for PE-CS-1.

PE-CS-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:103	123:103	123:103
		123:303	123:203

This table shows an RD and RT numbering scheme for PE-CS-2.

PE-CS-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:103	123:103	123:103
		123:303	123:203

Configuring a Central Services VPN

This topic identifies the command syntax that is required to configure a central services VPN.

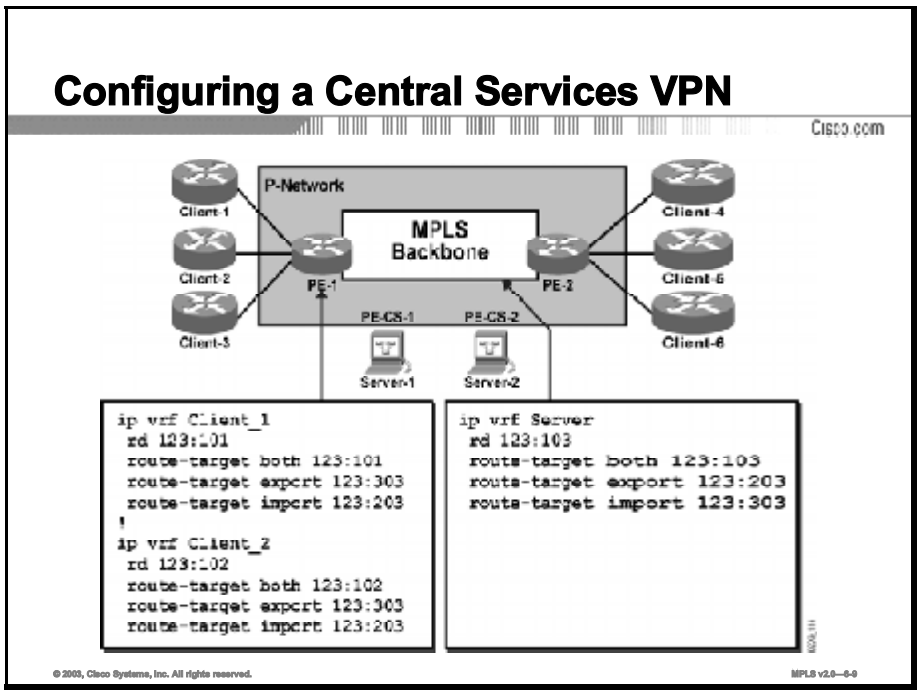

The figure shows a fraction of the configuration according to the RD and RT numbering scheme presented in the previous tables.

Central Services VPN and Simple VPN Requirements

This topic identifies the connectivity requirements when you are integrating a central services VPN with a simple VPN.

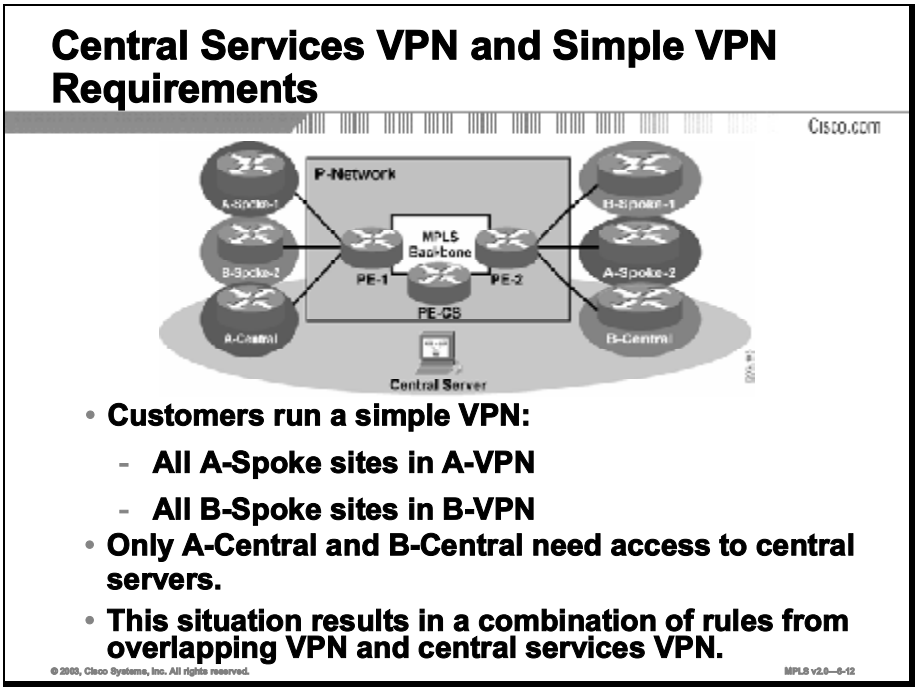

In this design, some of the customer sites need access to the central server. All other sites just need optimal intra-VPN access. The design is consequently a mixture of simple VPN topology and central services VPN topology.

Central Services VPN and Simple VPN Requirements (Cont.)

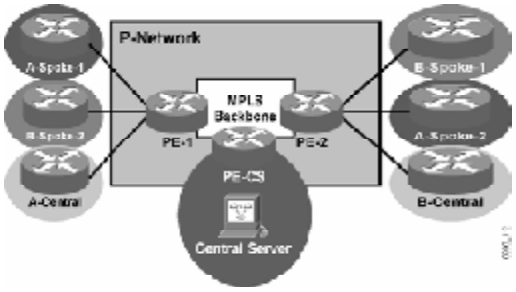

- **For all sites participating in a simple VPN, configure a separate VRF per set of sites participating in the same VPNs per PE router.**
- **For sites that are only clients of central servers, create a VRF per site.**
- **Create one VRF for central servers per PE router.**

When integrating a central services VPN with a simple VPN, you need one VRF per VPN for sites that have access to other sites in the customer VPN but no access to the central services VPN. You need one VRF per VPN for sites that have access to the central services VPN. Finally, you need one VRF for the central services VPN. This VPN is on another PE router in our example.

Configuring RDs in a Central Services and Simple VPN

This topic identifies the RD requirements when you are integrating a central services VPN with a simple VPN.

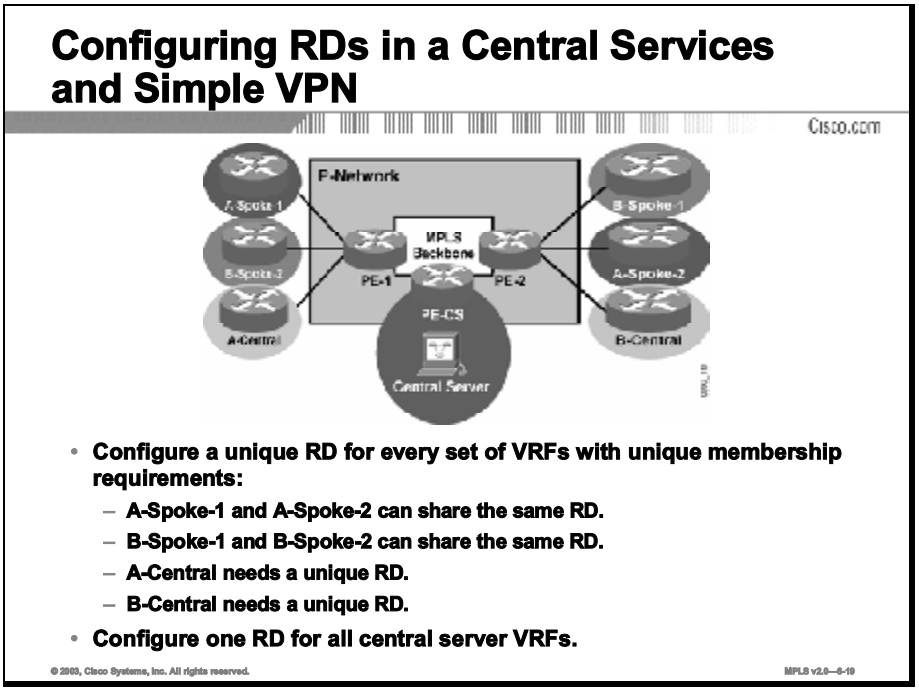

For this design, you need two RDs per VPN:

- One RD for simple VPN sites; the same value should also be used for import and export RT
- One RD for the central services VRFs

Configuring RTs in a Central Services and Simple VPN

This topic identifies the RT requirements when you are integrating a central services VPN with a simple VPN.

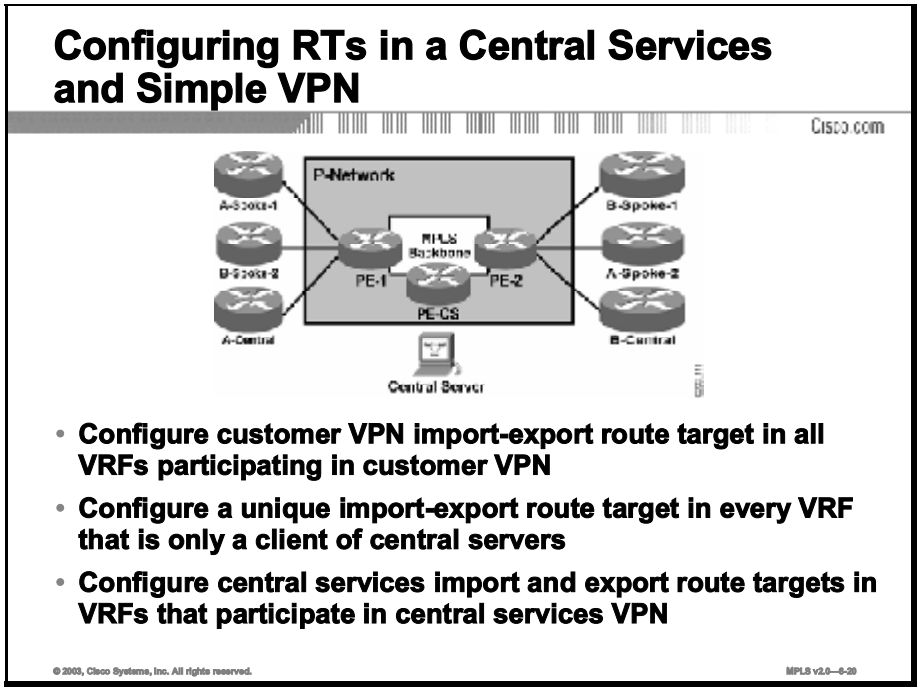

This table shows an RD and RT numbering scheme for PE-1:

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-A-Central	123:751	123:750 123:101	123:750 123:100

This table shows an RD and RT numbering scheme for PE-2.

PE-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-B-Central	123:761	123:760 123:101	123:760 123:100

This table shows an RD and RT numbering scheme for PE-CS.

PE-CS RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:101	123:101 123:100	123:101

Configuring VRFs in a Central Services and Simple VPN

This topic identifies the command syntax that is required when you are integrating a central services VPN with a simple VPN.

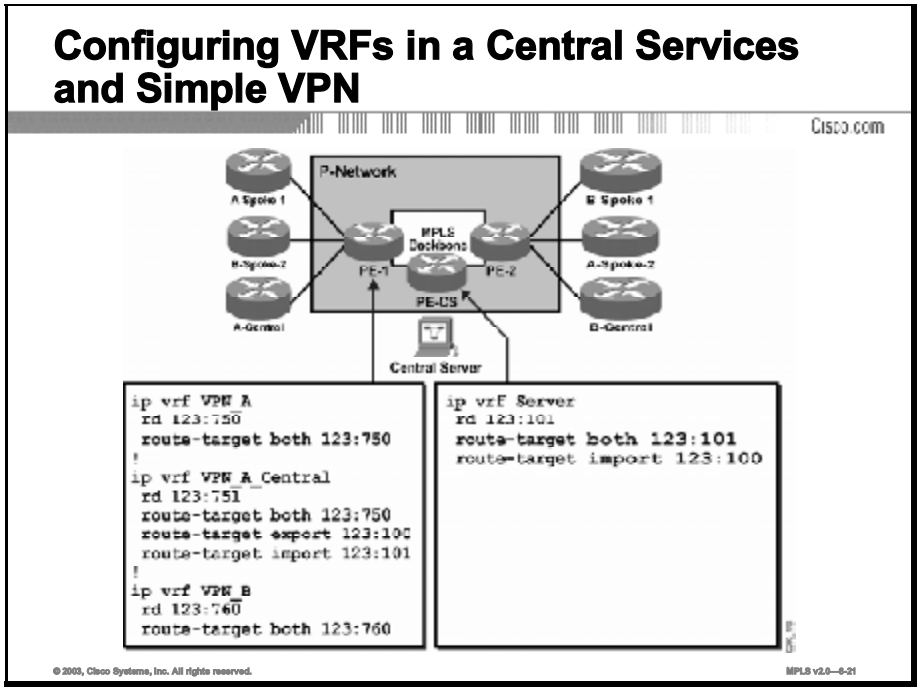

The example shows a fraction of the configuration according to the RD and RT numbering scheme presented in the previous tables.

Summary

This topic summarizes the key points discussed in this lesson.

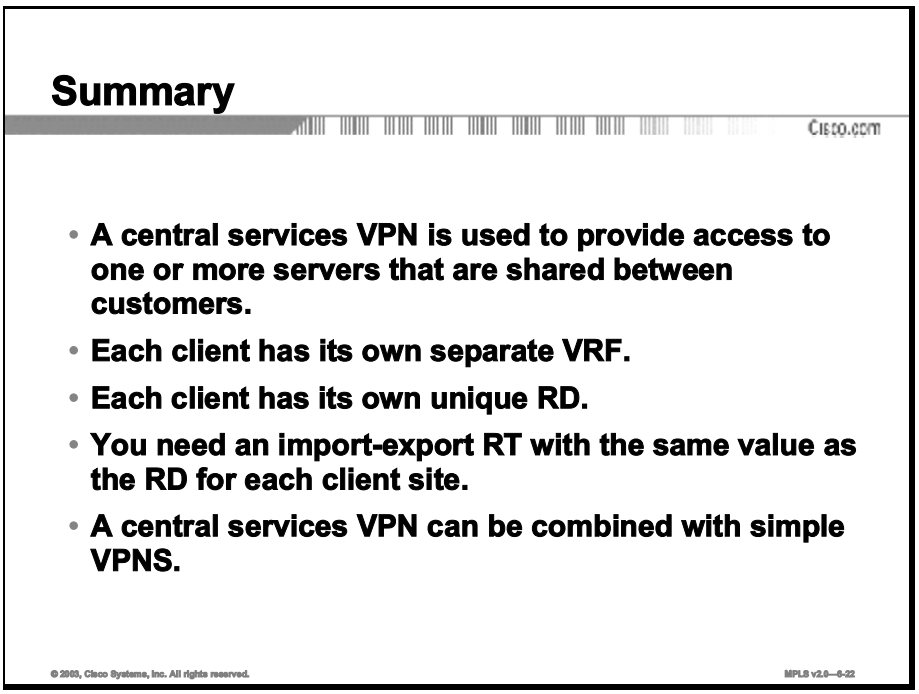

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are the typical usages for a central services VPN topology?

Q2) What is the connectivity model for a central services VPN topology?

Q3) What command syntax do you use to implement a central services VPN topology that supports two clients?

Client PE router	Server PE router

Client PE router	Server PE router

Q4) How many route distinguishers do you need for a central services VPN solution with 50 client sites and three server sites? How many route targets?

route distinguishers = _____ route targets = _____

Q5) How do you combine a central services VPN topology with a simple VPN topology?

_____

_____

_____

_____

_____

_____

Quiz Answer Key

Q1) in solutions where some sites (server sites) can communicate with all other sites, but all the other sites (client sites) can communicate only with the server sites

Relates to: Central Services VPN

Q2) The clients have their own RTs that they import and export; they also export networks with an RT that will be used by the services VRF and import networks with an RT that identifies the routes of the services site. The services site imports and exports networks with the RT of its VPN, but it also imports networks with RTs that identify the client sites.

Relates to: Central Services VPN Routing; Central Services VPN Data Flow Model

Q3)

Client PE router	Server PE router
ip vrf Client_1	ip vrf Server
rd 123:101	rd 123:103
route-target both 123:101	route-target both 123: 203
route-target export 123:303	route-target import 123:303
`route-target import 123:203`	
`ip vrf Client_2`	
`rd 123:102`	
`route-target both 123:102`	
`route-target export 123:303`	
`route-target import 123:203`	

Relates to: Steps to Configuring a Central Services VPN; Configuring a Central Services VPN

Q4) route distinguishers = 51 route targets = 52

You need one RD for each client (50) and one RD (1) shared by both server sites. You need one RT for each client (50) to export its routes to its VPN, one RT (1) for all of the clients to export their routes to the server, and one RT (1) that is shared by both server sites to export their routes to the clients.

Relates to: Configuring a Central Services VPN

Q5) Create a simple VPN that provides connectivity for all of the customer sites that do not need access to the central services.
Create a simple VPN that provides access between all of the server sites that are in the service.
Create an overlapping VPN that contains the sites that must have access to both the customer VPNs and the services VPN.

Relates to: Central Services VPN and Simple VPN Requirements;
Configuring RDs in a Central Services and Simple VPN

Managed CE Routers Service

Overview

A service provider can use a separate network management VPN to manage the CE routers of all the VPNs. A pair of RT extended communities is used to accomplish this goal. One RT exports CE router loopback addresses and is imported into the VRF of the network management VPN. The other RT exports the networks from the VRF associated with the network management VPN and imports them into all other VRFs. This lesson explains some of the requirements and the implementation solution for the managed CE routers service.

Relevance

It is important to be able to recognize the requirements of the network and to match them up with customer requests. This lesson looks at one such requirement and explains how to handle it.

Objectives

This lesson identifies the characteristics of the managed CE routers service.

Upon completing this lesson, you will be able to:

- Identify the overall requirements of a managed CE routers VPN
- Identify the VRF and RD requirements of a managed CE routers VPN
- Identify the RT requirements of a managed CE routers VPN
- Identify the command syntax that is required to configure a managed CE routers VPN

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of MPLS components and VPN implementations, and familiarity with Cisco IOS platforms

Outline

This lesson includes these topics:

- Overview
- Managed CE Routers
- VRF Creation and RD Overview
- Configuring Route Targets
- Configuring VRFs
- Summary
- Quiz

Managed CE Routers

This topic identifies the overall requirements of a managed CE routers VPN.

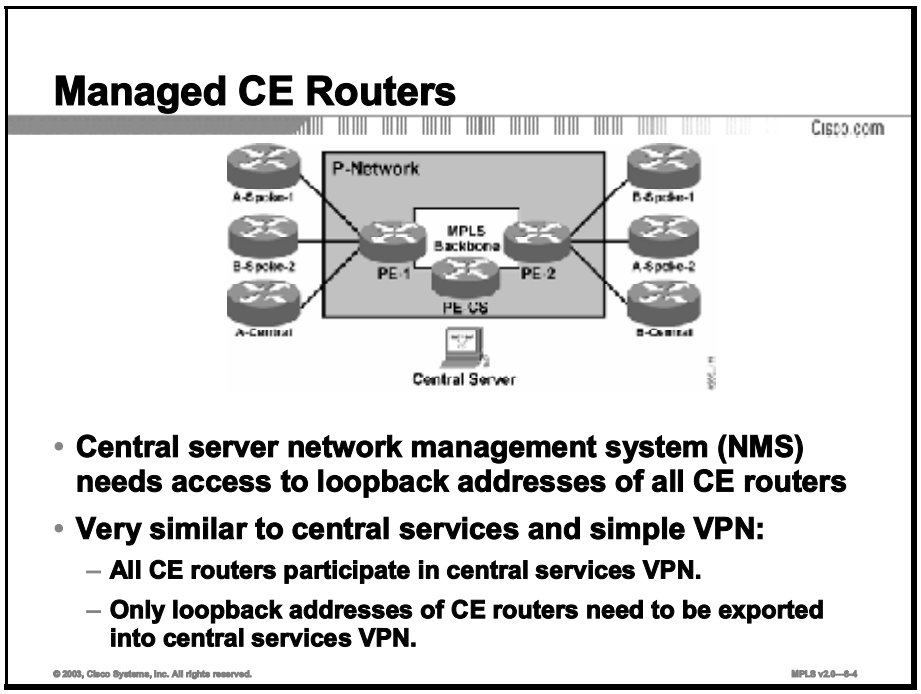

If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but not to the other destinations at customer sites. This requirement is usually implemented by deploying a separate VPN for management purposes. This VPN needs to see all the loopback interfaces of all the CE routers. All CE routers have to see the network management VPN. The design is very similar to that of the central services VPN; the only difference is that you mark only loopback addresses to be imported into the network management VPN.

Note The topology described in this lesson is sometimes referred to as a "gray" network management VPN implementation because all CE routers are accessed through a single link between the network management system (NMS) CE router and the network core. An alternative solution (a "rainbow" network management VPN), in which the NMS CE router has separate connections to each managed CE router, is usually used in combination with overlay VPNs (for example, Frame Relay networks).

VRF Creation and RD Overview

This topic identifies the VRF and RD requirements of a managed CE routers VPN.

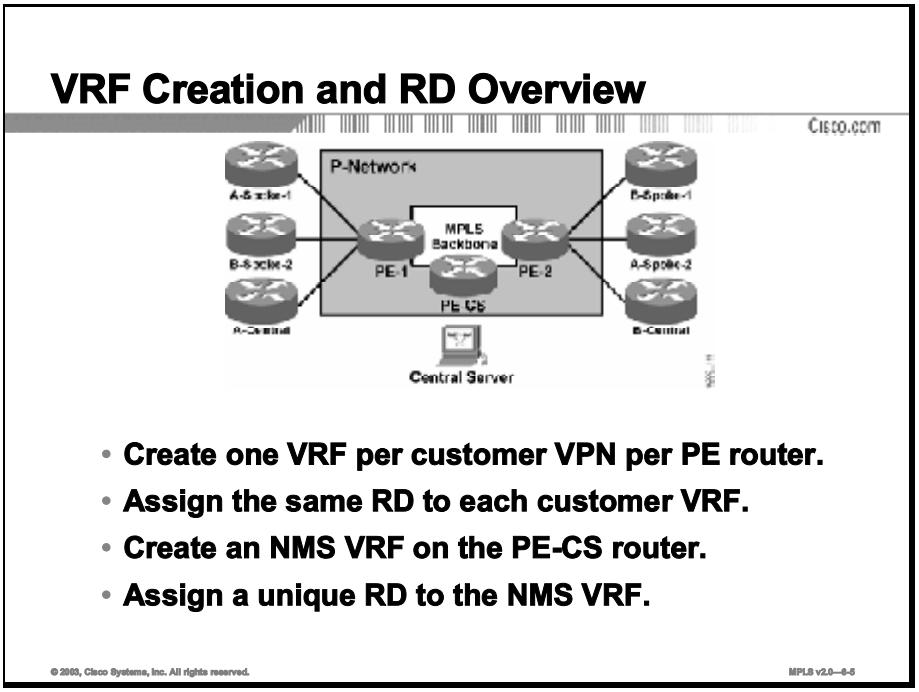

The VRF and RD design is the same as with central services VPNs. The only difference between this topology and the central services VPN topology combined with a simple VPN topology is the RT marking process during route export.

Configuring Route Targets

This topic identifies the RT requirements of a managed CE routers VPN.

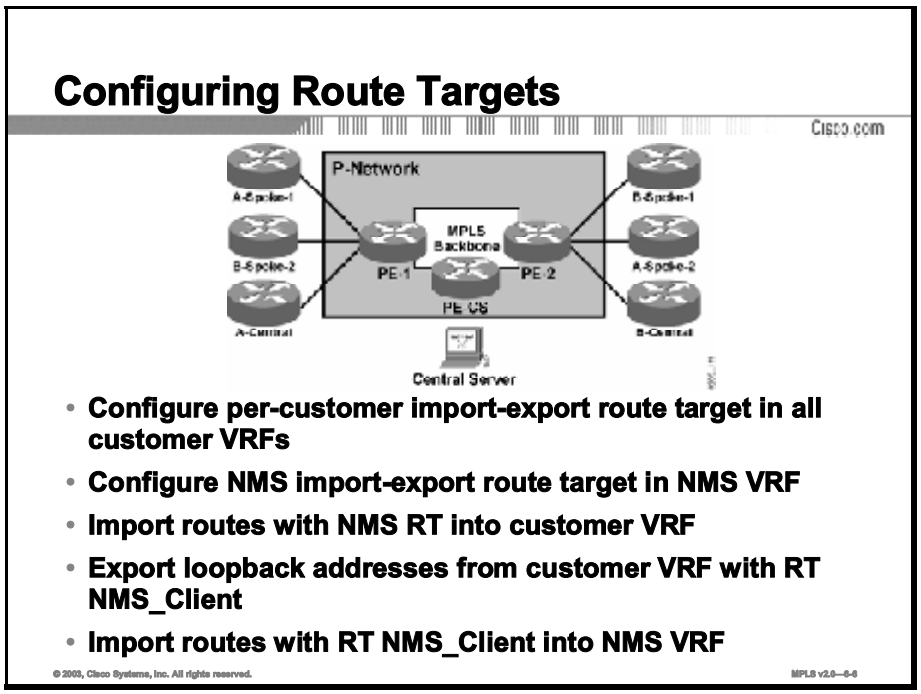

This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750 123:101	123:750 123:100 (NMS_Client)
VPN-B	123:760	123:760 123:101	123:760 123:100 (NMS_Client)

This table shows an RD and RT numbering scheme for PE-CS.

PE-CS RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
NMS	123:101	123:101 123:100 (NMS_Client)	123:101

Configuring VRFs

This topic identifies the command syntax that is required to configure a managed CE routers VPN.

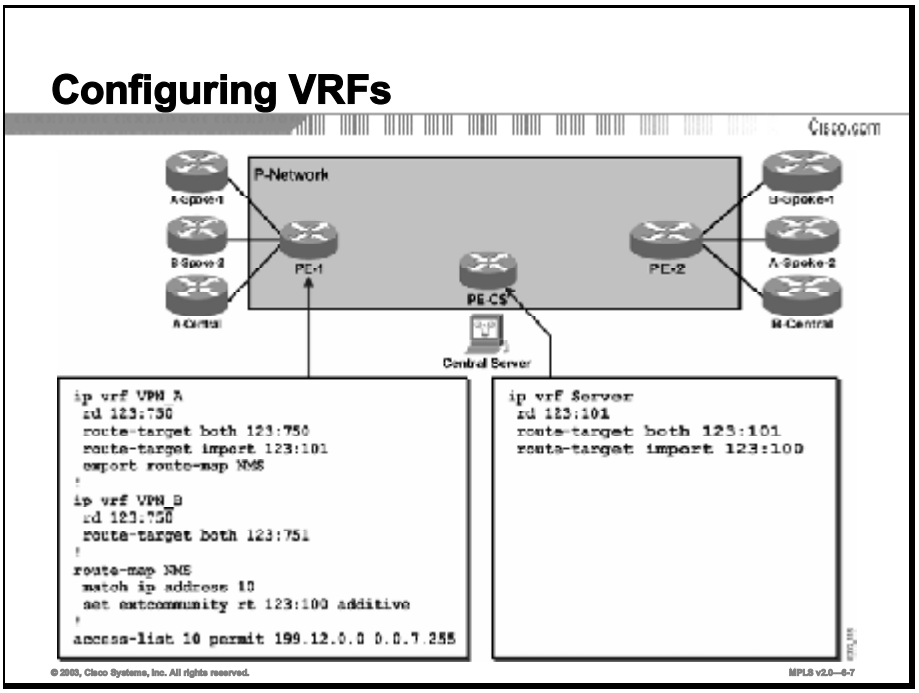

The figure shows a sample configuration for a customer VRF with differentiated RT export for loopback addresses according to the numbering scheme shown on the previous page. An export route map is used to match one part of the IP address space and attach an additional RT to the routes within this address space (CE router loopback addresses).

Note The routing protocol between PE and CE routers has to be secured (with distribute lists or prefix lists) to prevent customers from announcing routes in the address space dedicated to network management; otherwise, customers can gain two-way connectivity to the network management station.

The CE router loopback addresses are then imported into the server VPN based on the additional RT attached to them during the export process.

Note This design allows client sites to send packets to the network management VPN regardless of the source address. Special precautions should be taken to protect the network management VPN from potential threats and denial-of-service attacks coming from customer sites.

Summary

This topic summarizes the key points discussed in this lesson.

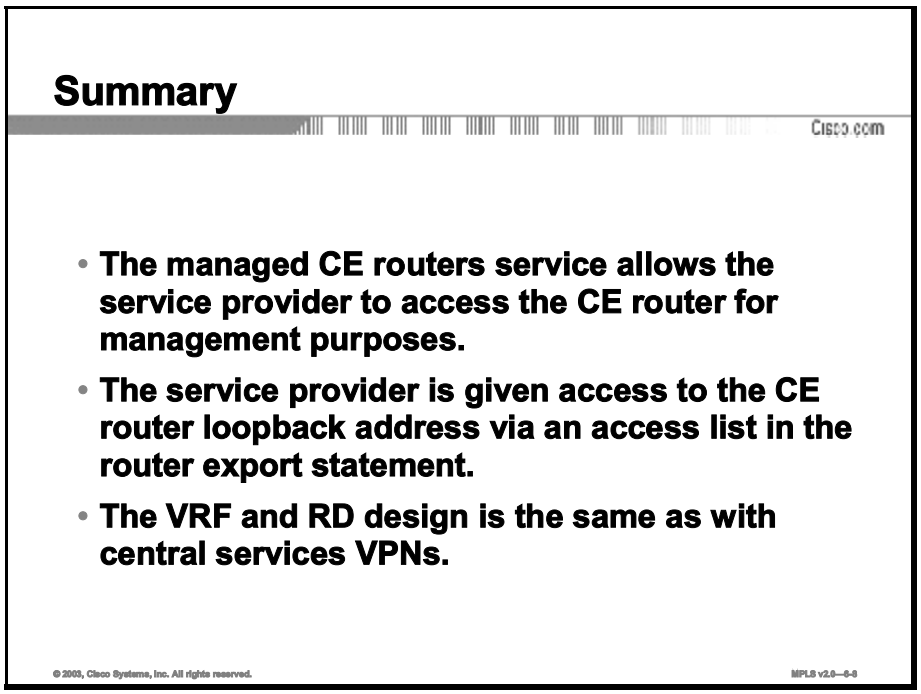

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about VPNs.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Why do you need the managed CE routers service?

Q2) What is the main difference between the managed CE routers service and the typical central services VPN topology?

Q3) What syntax would you use for an export statement that limits the export to the loopback address of 192.168.10.1?

Quiz Answer Key

Q1) If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but not to the other destinations at customer sites.

Relates to: Managed CE Routers

Q2) The VRF and RD design is similar to that of a central services VPN. The managed CE routers service combines a service VPN and simple VPN topology like the central services VPN. However, the route export statement uses an access list to limit the exported addresses to the loopback address of the managed routers.

Relates to: VRF Creation and RD Overview; Configuring Route Targets

Q3)
```
ip vrf VPN_A
export route-map NMS
route-map NMS
match ip access-list 10
set extcommunity rt 123:100 additive
access-list 10 permit 192.160.10.1 0.0.0.255
```

Relates to: VRF Creation and RD Overview; Configuring Route Targets; Configuring VRFs

MPLS Managed Services

Overview

Market forces are driving service providers to provide additional centralized services to their customers. In addition, these services need to be integrated with existing VPN service. To meet this need, Cisco has provided a set of VPN-aware managed services. This lesson discusses Cisco MPLS managed services, focusing on which service provider needs that they meet, and how they are implemented in an MPLS network.

Relevance

To successfully implement managed services, you need to understand the needs of the service provider, which kind of service can meet those needs, and how that service is implemented.

Objectives

This lesson identifies the features of Cisco managed services, focusing on where they can be applied and how they are implemented.

Upon completing this lesson, you will be able to:

- Identify the Cisco MPLS VPN managed services
- Describe the implementation of managed Network Address Translation
- Describe the implementation of managed DHCP relay
- Describe the implementation of managed on-demand address pools
- Describe the implementation of managed HSRP and VRRP
- Describe the implementation of managed multicast VPNs

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS VPN technology and configuration knowledge
- Basic NAT functionality and configuration knowledge
- Basic DHCP relay functionality and configuration knowledge
- Basic on-demand address pools (ODAP) functionality and configuration knowledge
- Basic Hot Standby Router Protocol (HSRP) functionality and configuration knowledge
- Basic Virtual Router Redundancy Protocol (VRRP) functionality and configuration knowledge

Outline

This lesson includes these topics:

- Overview
- Managed Services Overview
- Network Address Translation
- DHCP Relay
- On-Demand Address Pools
- HSRP and VRRP
- Multicast VPNs
- Summary
- Quiz

Managed Services Overview

This topic provides an overview of the Cisco MPLS VPN managed services.

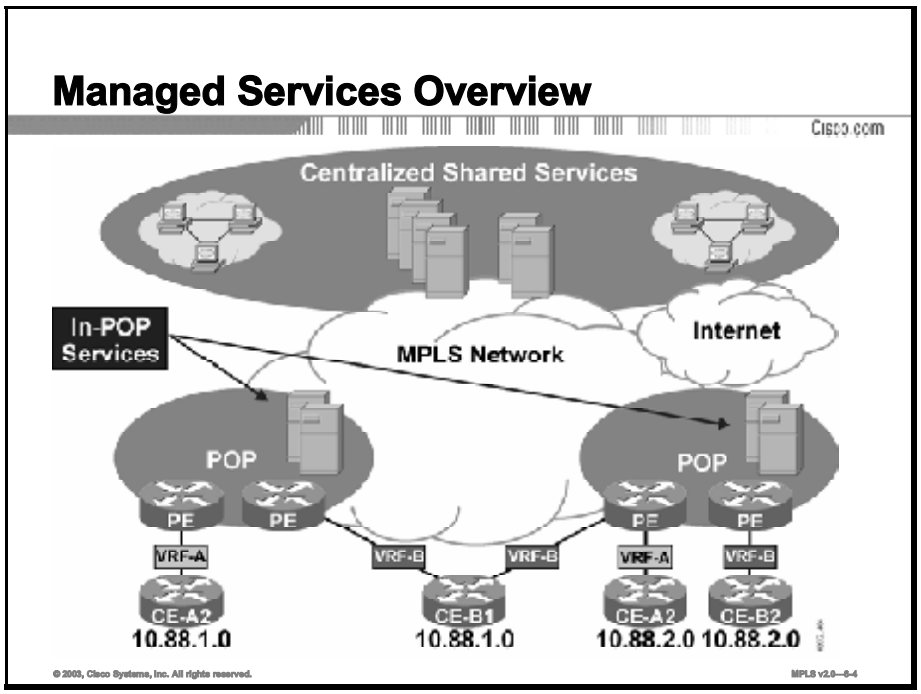

In modern networks, many end users have a need to connect to common services like e-mail, DHCP servers, and so on. Typically, these services have been provided by individual enterprises as part of their network.

Cisco MPLS for Managed Shared Services is a set of features delivered in Cisco IOS software for enabling managed shared services for MPLS VPNs. Building on leading Cisco MPLS capabilities, service providers now can provide their enterprise clients all the connectivity benefits associated with Cisco MPLS VPNs while creating additional revenue streams by also providing economically attractive, IP services.

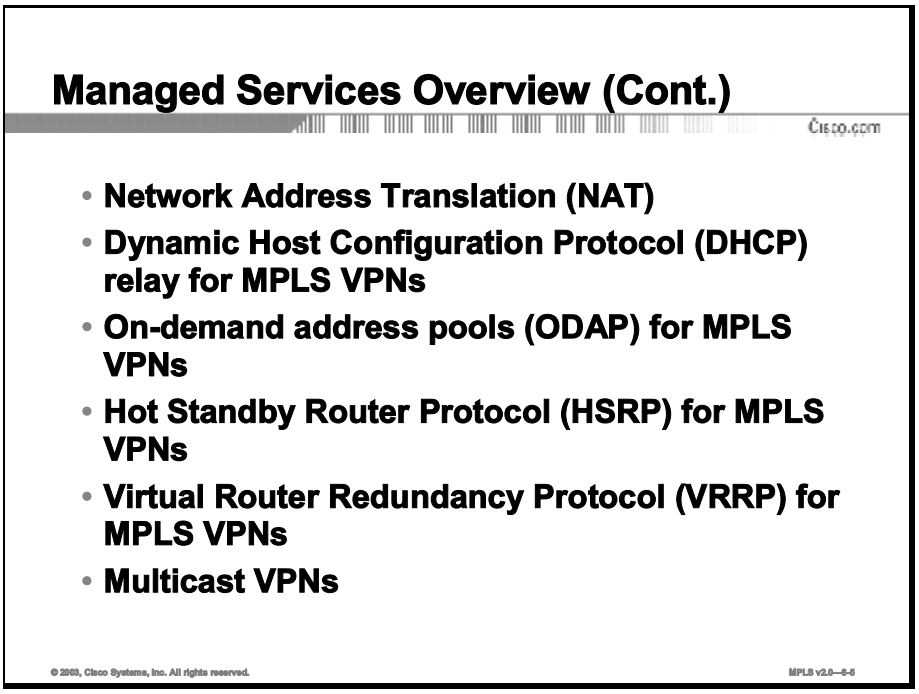

Cisco has expanded its widely deployed MPLS VPN solution to include the following technologies in Cisco IOS software:

- Network Address Translation (NAT) for MPLS VPNs
- Dynamic Host Configuration Protocol (DHCP) relay for MPLS VPNs
- Ondemand address pools (ODAP) for MPLS VPNs
- Hot Standby Router Protocol (HSRP) for MPLS VPNs
- Virtual Router Redundancy Protocol (VRRP) for MPLS VPNs
- Multicast VPNs

With these key new technologies, enterprise IP services can now be moved from the enterprise network into the MPLS VPN network of the service provider, and shared across multiple VPNs for greater operational leverage and economies of scale.

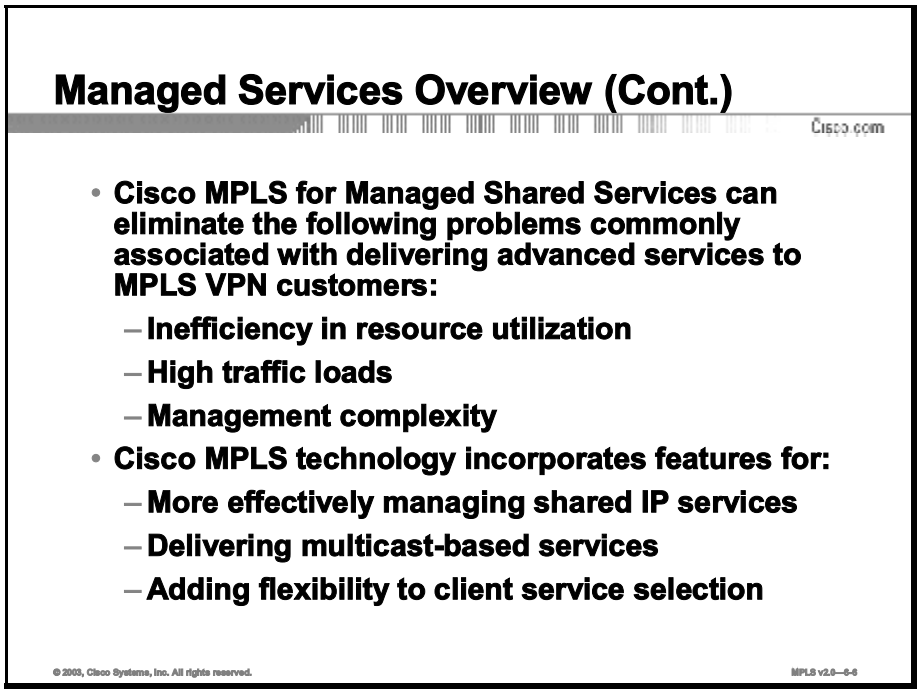

Cisco MPLS for Managed Shared Services eliminates many of the problems—such as inefficiency in resource utilization, high traffic loads, and management complexity—commonly associated with delivering advanced services to MPLS VPN customers. The Cisco MPLS technology incorporates features for more effectively managing shared IP services, delivering multicast-based services, and adding flexibility to client service selection.

The Cisco MPLS offering includes a number of VRF features that present opportunities for new IP services revenue streams, as well as for cost savings. NAT for MPLS VPNs, for instance, lets service providers more cost-effectively support services such as content hosting, enterprise resource planning (ERP) application hosting, and managed Internet access. Other features add support in the MPLS network for industry-standard protocols, as well as improve or automate routing control. The comprehensive collection of functions can help service providers eliminate many customer-expressed barriers to entry by ensuring that MPLS VPN business clients have access to the robust functionality that they expect in the enterprise environment.

Cisco MPLS for Managed Shared Services also incorporates multicast VPN functionality to help service providers meet enterprise market demands for IP services essential in applications such as telecommuting. By reducing packet replication in the MPLS network, multicast VPN technology allows for massively scalable distribution of data, voice, and video streams. Utilizing multicast VPN features, service providers can leverage existing infrastructure resources to offer competitive services in videoconferencing, e-learning, and other Internet-based streaming applications.

Taken together, the Cisco MPLS for Managed Shared Services features give service providers powerful new MPLS VPN functionality and versatility—without deployment or management complexity.

Network Address Translation

This topic describes the implementation of NAT services in an MPLS VPN environment.

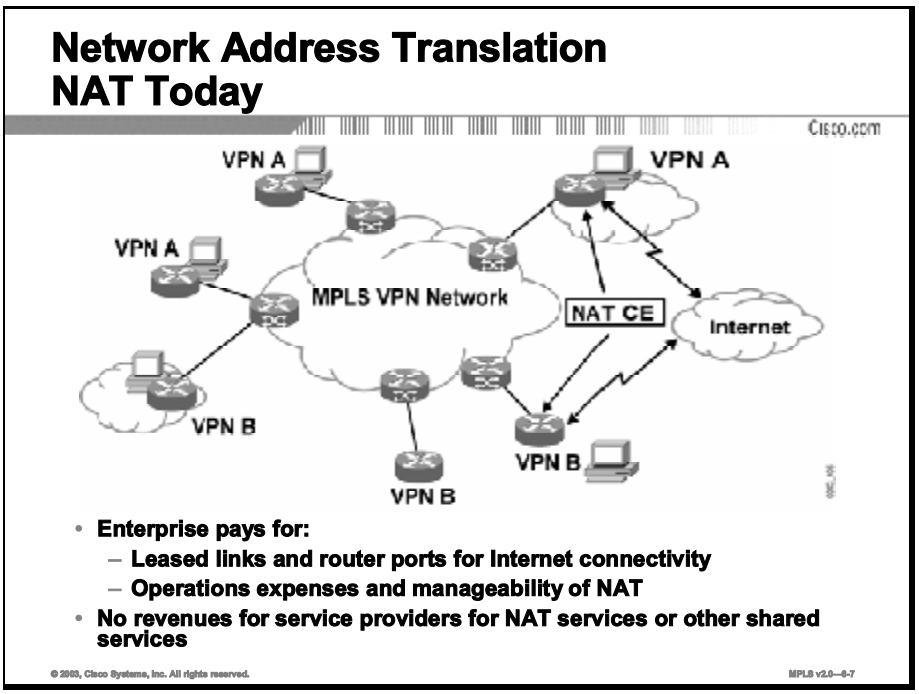

In modern MPLS networks, enterprises have to pay for leased links and router ports for Internet connectivity in addition to VPN connectivity, as well as the operational expenses associated with internally managing NAT. While service providers can currently provide NAT services to their enterprise clients with additional router and NAT devices, it is a highly complex design. NAT for MPLS VPNs is a simpler and more flexible way to integrate NAT services within MPLS VPNs with a single network connection that provides both MPLS VPN connectivity and access to shared services.

Because NAT for MPLS VPNs provides more economical NAT services, these services can be made more appealing to enterprise clients with a resulting revenue opportunity for service providers.

Network Address Translation (Cont.)
NAT for Shared Services

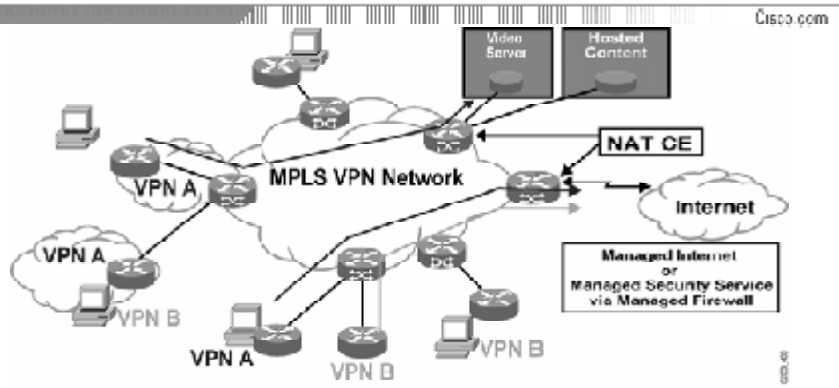

- **Reduced CapEx and OpEx for the enterprise:**
  - **Leased lines and router ports for Internet connectivity go away**
  - **Reduces network complexity, and operations and manageability cost**
- **Increased revenues for service providers:**
  - **Can provide outsourced NAT services to enterprises**

The integration of NAT with MPLS VPNs allows multiple MPLS VPNs to be configured on a single device to work together. NAT can differentiate which MPLS VPN that it receives IP traffic from even if the MPLS VPNs are all using the same IP addressing scheme. This enhancement enables multiple MPLS VPN customers to share services while ensuring that each MPLS VPN is completely separate from the other.

MPLS service providers would like to provide value-added services such as Internet connectivity, Domain Name System (DNS) servers, and Voice over IP (VoIP) to their customers. These additions require that customer IP addresses be different when reaching the services. Because MPLS VPNs allow customers to use overlapped IP addresses in their networks, NAT must be implemented to make the services possible.

Cisco NAT for MPLS VPNs provides the following:

- A simple and more flexible way of integrating NAT with MPLS VPNs
- Automatic management of the overlapping of VPN address spaces (allowable in MPLS VPNs) to ensure that addresses are mapped correctly in shared-services applications
- Centralized delivery of full-VPN NAT services
- NAT redundancy (NAT can be configured on one or more PE routers)

Cisco NAT for MPLS VPNs eliminates the requirement for physical connectivity between a shared service and the provider network that is performing NAT.

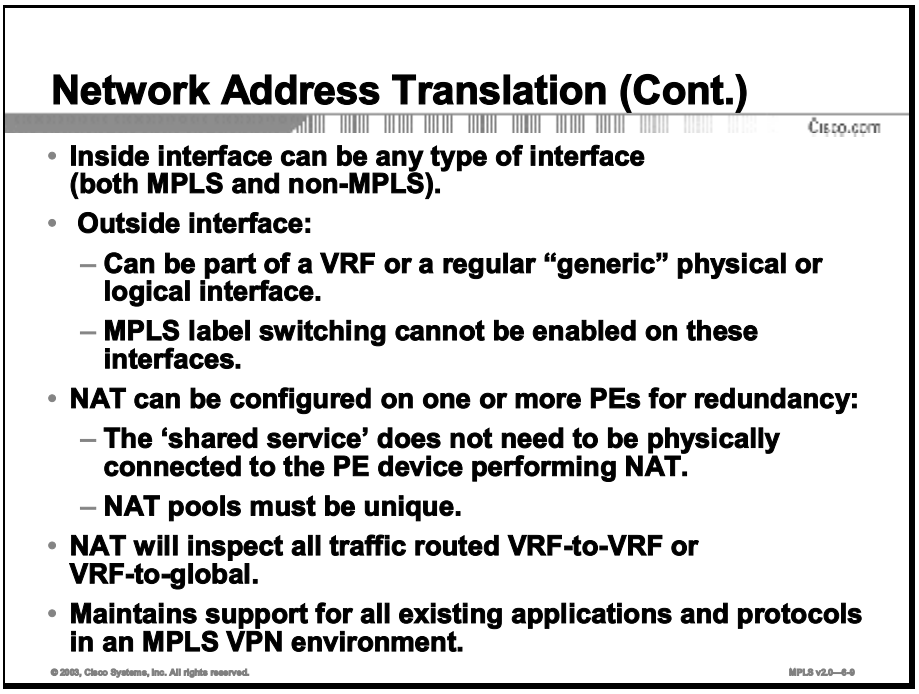

NAT can be implemented on the PE route in the following scenarios:

- **Service point:** Shared access can be from a generic interface or from a VPN interface.
- **NAT point:** NAT can be configured on the PE router that is directly connected to the shared access gateway, or on the PE router that is not directly connected to the shared access gateway.
- **NAT interface:** The shared access gateway interface is most often configured as the outside interface of NAT. The inside interface of NAT can be either the provider edge-customer edge (PE-CE) interface of a VPN, the interface to the MPLS backbone, or both. The shared access gateway interface can also be configured as the inside interface.
- **Routing type:** Common service can be Internet connectivity or a common server. For Internet connectivity, a default route should be propagated to all the VPN customers that use the service. For common server access, a static or dynamically learned route should be propagated to the VPN customers.
- **NAT configuration:** NAT can have different configurations: static, dynamic, pool/interface overloading, and route map.

The concept of VPN-aware NAT is very similar to classic NAT. Inside and outside interfaces serve the same function as in classic NAT; only the location of the NAT service is changed. An inside interface can be any type of interface. An outside interface can be part of a VRF or a regular "generic" physical or logical interface, but MPLS cannot be enabled on these interfaces.

The "shared service" does not need to be physically connected to the PE device performing NAT. In addition, NAT can be configured on one or more PEs for redundancy.

NAT will inspect all traffic routed VRF-to-VRF or VRF-to-global to determine when and where NAT should be applied.

VPN-aware NAT also maintains support for all existing applications and protocols in an MPLS VPN environment.

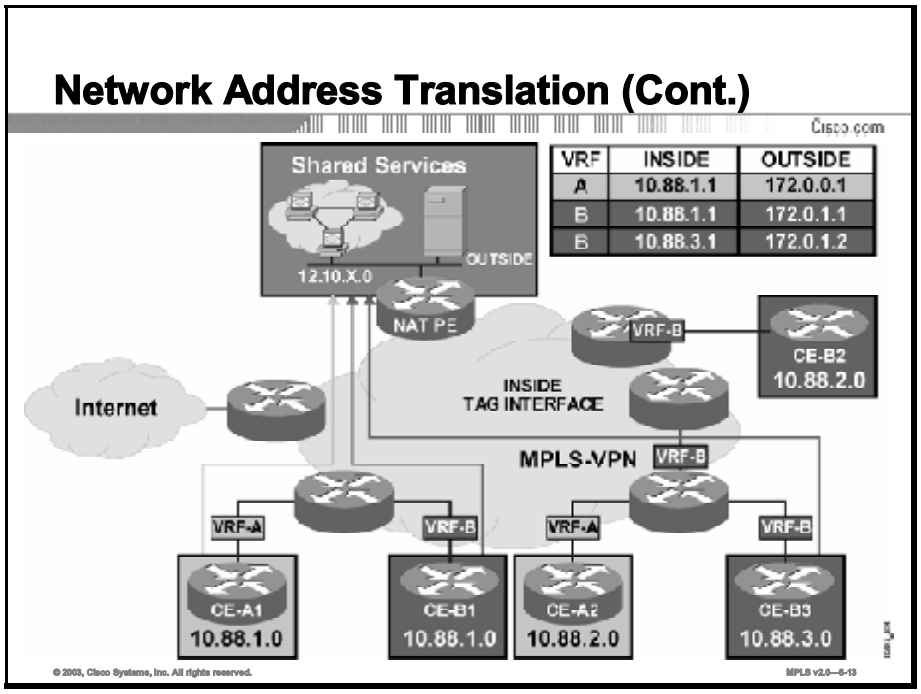

There are two approaches to implementing NAT in the MPLS VPN network. NAT can be implemented on the CE router, which is already supported by NAT, or it can be implemented on a PE router. Integration of NAT with MPLS VPNs enables the implementation of NAT on a PE router in an MPLS cloud.

The figure presents an example of VPN-aware NAT. CE-A1, CE-A2, CE-B1, and CE-B2 are clients in VRF-A and VRF-B. Packets from these clients destined for the shared service are routed to the inside interface of the NAT PE over their respective VPNs. At the NAT PE, the address translation process replaces the inside source address with the outside source address from the NAT table and forwards the packet to the shared service.

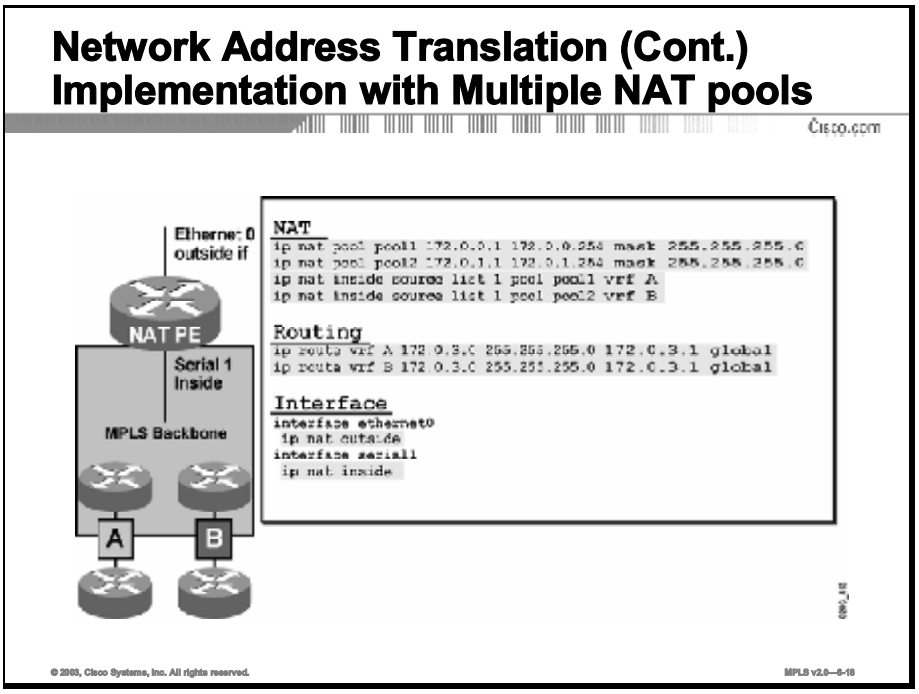

The figure presents an example of VPN-aware NAT configuration for two VPNs, A and B. NAT services are being configured on the PE router connected to the shared services. As indicated earlier, the NAT services can also be configured on any other router that is part of the VPN.

NAT pools are configured with a standard NAT configuration command: **ip nat pool**. Only one NAT pool is required; however, in this example, there are two pools, one for each VPN in order to allow for easy address administration. The NAT pools are assigned to their respective VPNs using the **ip nat inside pool** command.

NAT services are applied to the interfaces using the **ip nat** command.

Because the outside interface is not participating in the VPN, the VPN VRF would not normally know of its existence. A static default route is created, pointing to the next-hop address of the shared services for each VPN, by using the **ip route vrf** command with the **global** keyword.

For further information see the following:

- http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1839/products_feature_guide09186a00801145f5.html#22289

DHCP Relay

This topic describes the implementation of managed DHCP relay services in an MPLS VPN environment.

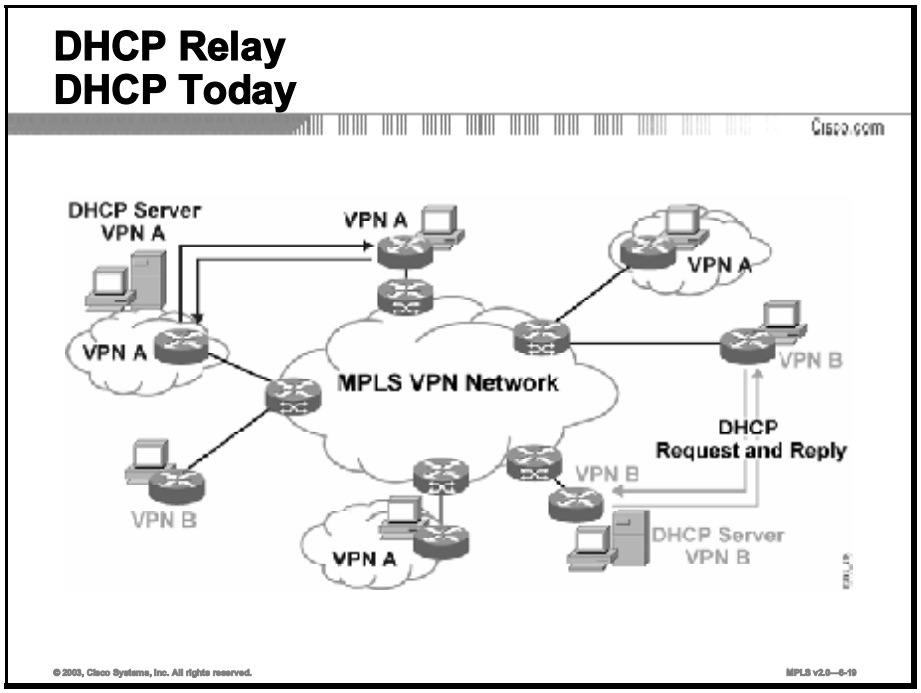

Current implementations of DHCP suffer from a couple issues:

- Even if they are collocated, there is a replication of DHCP servers per VPN.
- There is no added value from the service provider.

DHCP Relay (Cont.)
DHCP Support for Shared Services

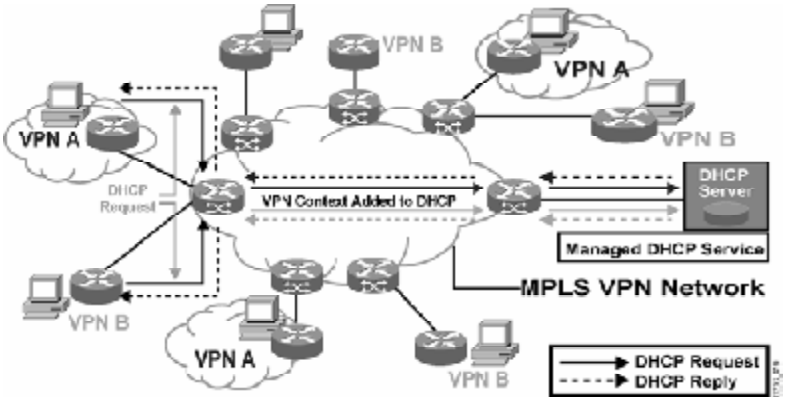

Service providers can take advantage of another centralized service to support DHCP clients. DHCP Relay for MPLS VPNs enables a DHCP relay agent to forward information about the DHCP request and the VPN association when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can then use that information to interpret IP addresses or other parameter-assignment policies. The DHCP relay agent information option is organized as a single DHCP option that contains one or more suboptions (VPN identifier, subnet selection, and server identifier override) to convey information known by the relay agent.

The DHCP relay agent information option (option 82) enables a DHCP relay agent to include information about itself when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can use this information to implement IP address or other parameter-assignment policies. The DHCP relay agent information option is organized as a single DHCP option that contains one or more suboptions that convey information known by the relay agent.

In some environments, a relay agent resides in a network element that also has access to one or more MPLS VPNs. A DHCP server that wants to offer service to DHCP clients on those different VPNs needs to know the VPN in which each client resides. The network element that contains the relay agent typically knows about the VPN association of the DHCP client and includes this information in the relay agent information option.

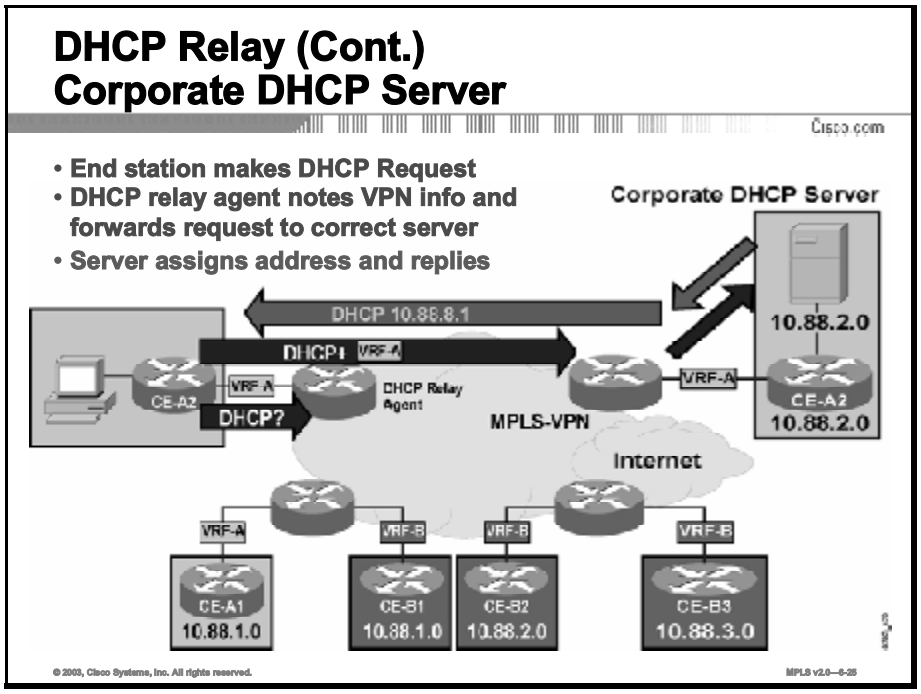

In this two-VPN example, a corporate DHCP server and a DHCP client have been added to VPN A. The client broadcasts a DHCP request to the local relay. The local relay converts the broadcast to a unicast request for the DHCP server and adds the VPN ID. This request is forwarded to the egress PE router based upon the DHCP server address. From the egress PE router, it is forwarded to the DHCP server.

The DHCP server assigns the client an address and replies to the DHCP relay, which in turn forwards the reply to the client.

The relay agent uses the VPN identifier suboption to tell the DHCP server the VPN for every DHCP request that it passes on to the DHCP server. This suboption is also used to properly forward any DHCP reply that the DHCP server sends back to the relay agent. The VPN identifier suboption contains the VPN ID configured on the incoming interface to which the client is connected. If you configure the VRF name but not the VPN ID, the VRF name is used as the VPN identifier suboption. If the interface is in global routing space, the VPN suboptions are not added.

The subnet selection suboption allows the separation of the subnet where the client resides from the IP address used to communicate with the relay agent. In typical DHCP processing, the gateway address specifies both the subnet on which a DHCP client resides and the IP address that the server can use to communicate with the relay agent. Situations exist where the relay agent needs to specify a subnet on which a DHCP client resides that is different from the IP address that the server can use to communicate with the relay agent. The subnet selection suboption is included in the relay agent information option and passed on to the DHCP server. The gateway address is changed to the outgoing interface of the relay agent toward the DHCP server. The DHCP server uses this gateway address to send reply packets back to the relay agent.

The server identifier override suboption value is copied in the reply packet from the DHCP server instead of the normal server ID address. The server identifier override suboption contains the incoming interface IP address, which is the IP address on the relay agent that is accessible from the client. Using this information, the DHCP client sends all renew and release packets to the relay agent. The relay agent adds all of the VPN suboptions and then forwards the renew and release packets to the original DHCP server.

After the relay agent has added these suboptions to the DHCP relay agent information option, the gateway address is changed to the outgoing interface of the relay agent toward the DHCP server. When the packets are returned from the DHCP server, the relay agent removes the relay agent information options and forwards the packets to the DHCP client on the correct VPN.

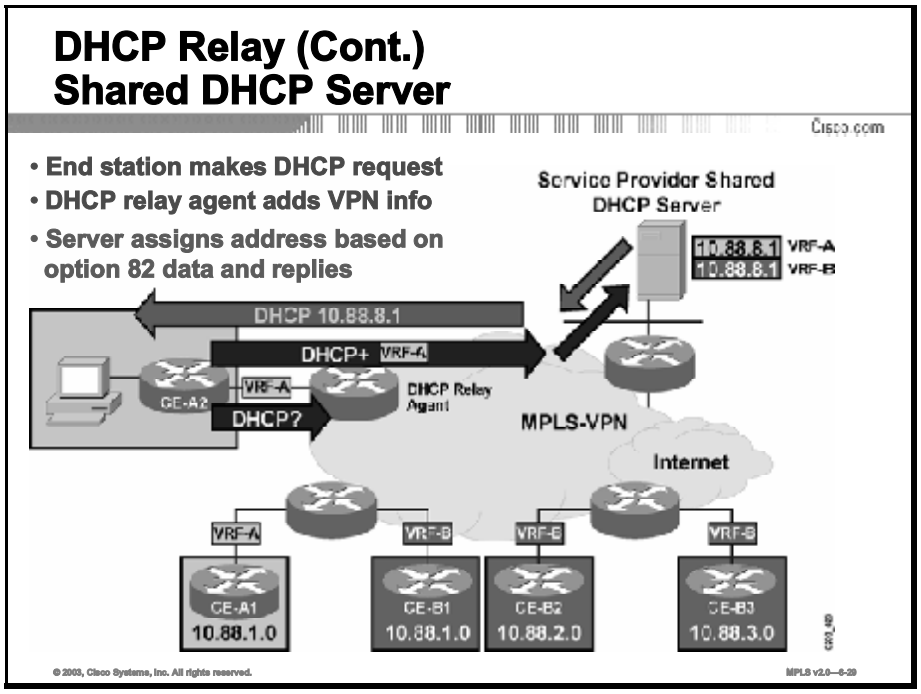

In this example, the corporate DHCP server has been replaced with a shared DHCP server provided by the service provider. Because the server is shared between VPNs, a NAT PE could also be included to provide address translation.

The client broadcasts a DHCP request to the local relay. The local relay converts the broadcast to a unicast request for the shared DHCP server and adds the VPN ID. This request is forwarded to the egress PE router via the NAT PE based upon the DHCP server address. At the NAT PE, an address translation is performed and the request is forwarded to the DHCP server.

The DHCP server assigns the client an address from the VPN pool and replies to the DHCP relay, which in turn forwards the reply to the client.

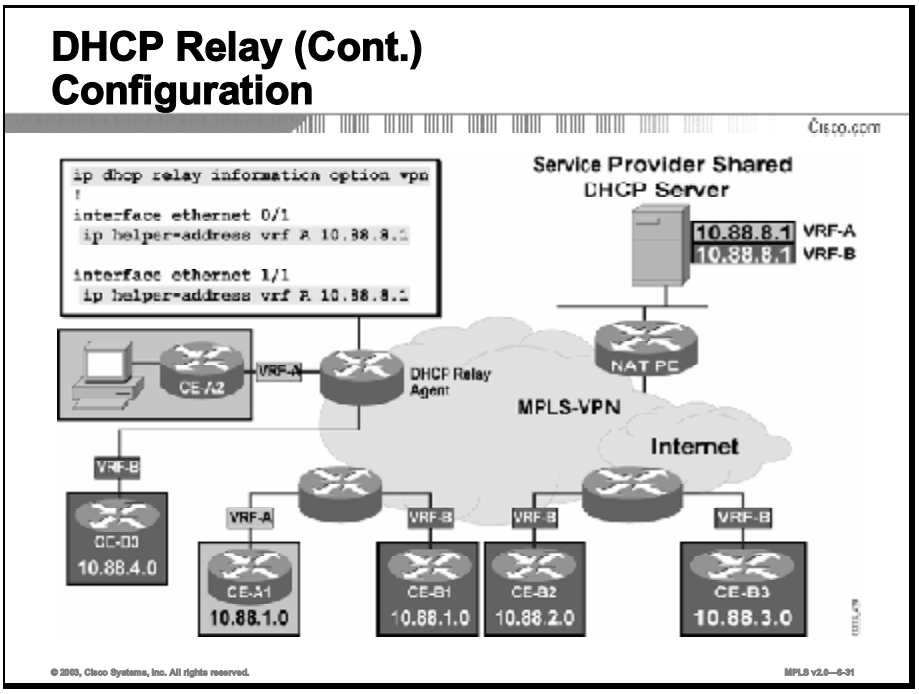

The figure presents an example of a typical DHCP relay configuration.

The **dhcp relay information option vpn** command enables the system to insert VPN suboptions into the DHCP relay agent information option in forwarded BOOTREQUEST messages to a DHCP server and sets the gateway address to the outgoing interface toward the DHCP server. The VPN suboptions are also added to the BOOTP broadcast packets when the command is configured.

The DHCP server address is configured using the **ip helper-address vrf** command.

Note	For simplicity, this example uses two separate interfaces for the VPNs. They could also be supported using Inter-Switch Link Protocol (ISL) and subinterfaces. Additional DHCP configuration options are available. See the Cisco IOS documentation for further information.

For further information see the following:

- http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087d3c.html

On-Demand Address Pools

This topic describes the implementation of managed ODAP services in an MPLS VPN environment.

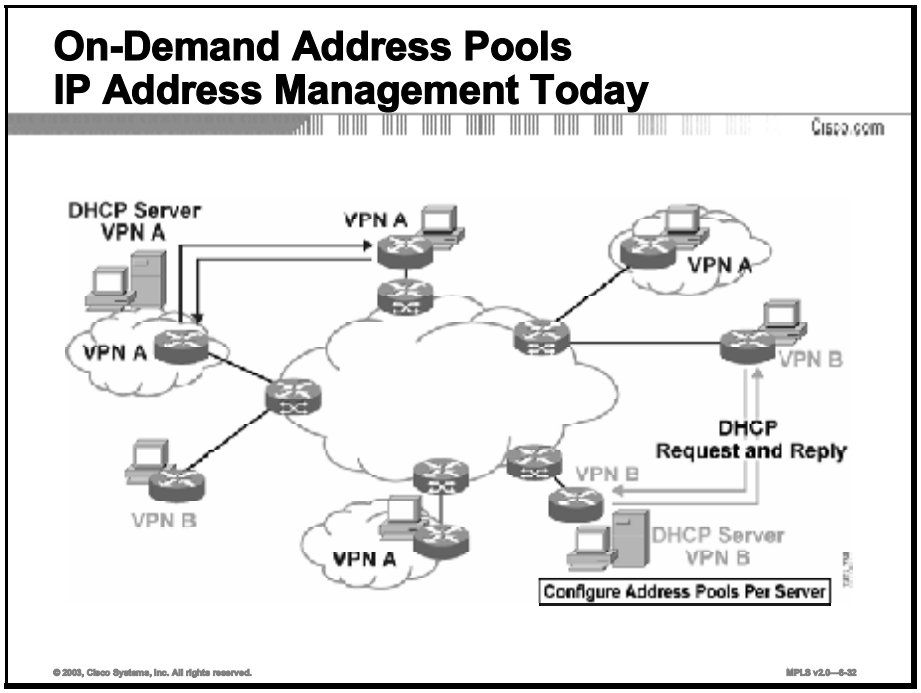

Today, service providers face the following challenges when it comes to efficient management of IP address space for customers:

- Address management is independent but inefficient.
- Providers need to manage addresses manually and allocate them to RADIUS or DHCP servers.
- Once site thresholds are reached, new addresses have to be manually allocated.

On-Demand Address Pools (Cont.)
Shared Service

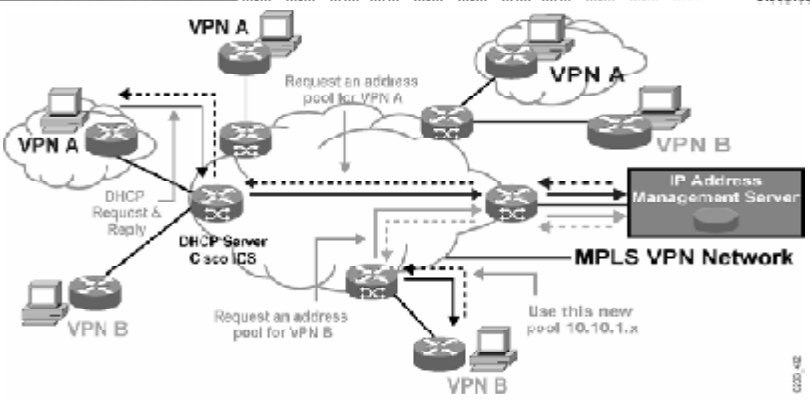

- **Use IOS DHCP server**
- **Request address pools on demand when a threshold is reached**
- **Much more efficient IP address management for MPLS VPNs**
- **Less network load and new revenue opportunity for service provider**

 MPLS v2.0—6-38

With MPLS VPNs, service providers have to allocate their IP address pools to independent RADIUS or DHCP servers for each VPN. Once the site threshold has been reached, new addresses have to be provided manually. With ODAP for MPLS VPNs, this process can now be fully automated and provided as a shared service on one or more servers. When the site threshold is exceeded, ODAP automates the process of expanding the overall address pool, reducing network loading, and performing configuration.

The Cisco ODAP for MPLS VPNs feature provides the following:

- Capabilities for automated control
- Support for MPLS VPNs, with addresses assigned per subnet, per interface
- Easy monitoring capabilities (pool manager can assess address utilization and expand the pool as needed)
- Simplified VPN setup (upon configuration, pool manager can request an initial subnet from the address pool server)

On-Demand Address Pools (Cont.)
Provisioning and Startup

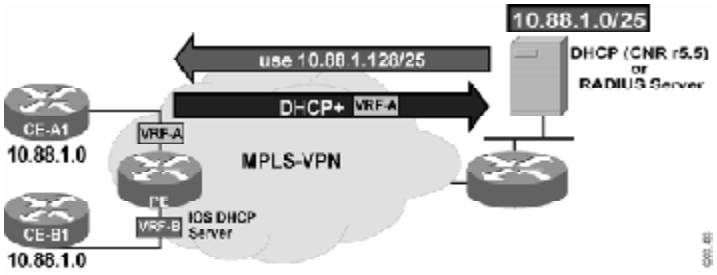

- **PE router is configured for ODAP.**
- **ODAP requests initial pool for VRF-A from server.**
- **Server replies with the initial address pool.**

As soon as the DHCP server is enabled on the PE router that has ODAP enabled, the PE router will request an initial address pool from its designated server.

On-Demand Address Pools (Cont.)
Address Pool Management

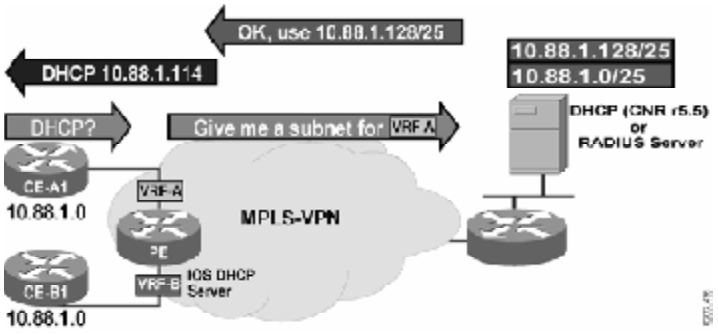

- End station makes DHCP request.
- DHCP server fulfills request from pool – reaches 90 percent.
- ODAP pool manager requests expansion.
- Server allocates another subnet and replies.
- PE adds subnet routing information to VRF.

The PE router will honor DHCP requests and assign addresses until its address pool is 90 percent depleted. At this point it will request an extension of the address pool from its designated server.

On-Demand Address Pools (Cont.)
VRF Pool Example

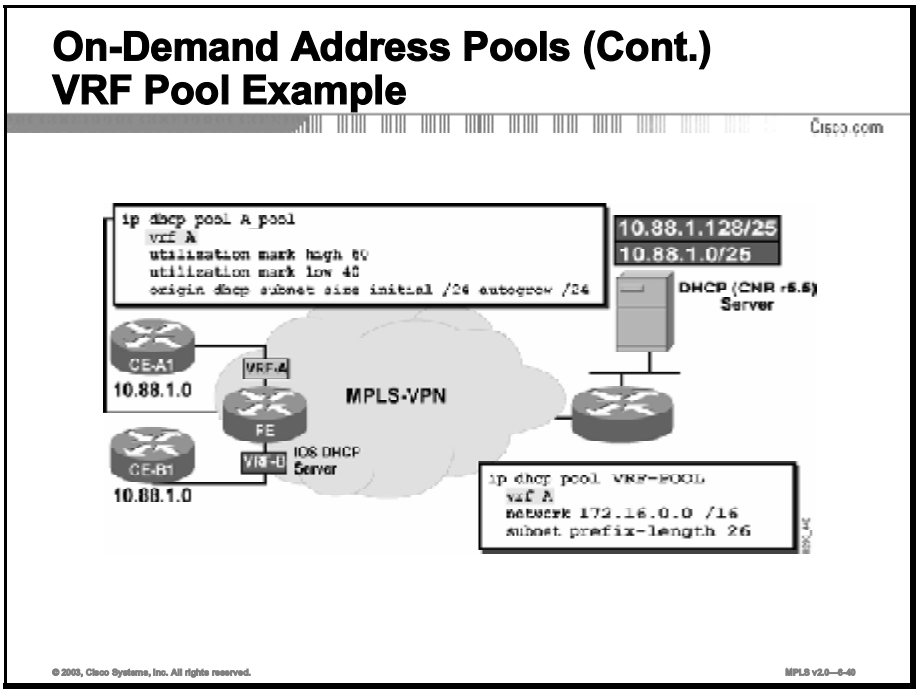

In this example, the PE router is configured as a DHCP server with ODAP, and a second router is configured to be a subnet allocation server for VNP A.

On the PE router, a DHCP pool named "A_pool" has been created. This pool is associated with VPN A. Three new commands have been introduced to the DHCP command presented in the previous topic. To configure an address pool as an ODAP, use the **origin** command. The **subnet size initial** *size* option is used to set the initial size of the first requested subnet. You can enter *size* as either the subnet mask (nnnn.nnnn.nnnn.nnnn) or prefix size (/nn). The **autogrow** *size* option is used to specify that the pool can grow incrementally. The *size* argument is the size of the requested subnets when the pool requests additional subnets (upon detection of high utilization). You can enter *size* as either the subnet mask (nnnn.nnnn.nnnn.nnnn) or prefix size (/nn).

To configure the high-utilization mark of the current address pool size, use the **utilization mark high** command in DHCP pool configuration mode. To configure the low utilization mark of the current address pool size, use the **utilization mark low** command in DHCP pool configuration mode.

On the ODAP server, a VRF subnet allocation pool named "VRF-POOL," which allocates subnets from the 172.16.0.0/16 network, has been configured to match the VRF named "A." The configuration of the **subnet prefix-length** command in this example configures the size of each subnet that is allocated from the subnet pool to support 62-host IP addresses.

Note	Additional DHCP and ODAP configuration options are available. Reference the Cisco IOS documentation for further information.

HSRP and VRRP

This topic describes the implementation of managed HSRP and VRRP services in an MPLS environment.

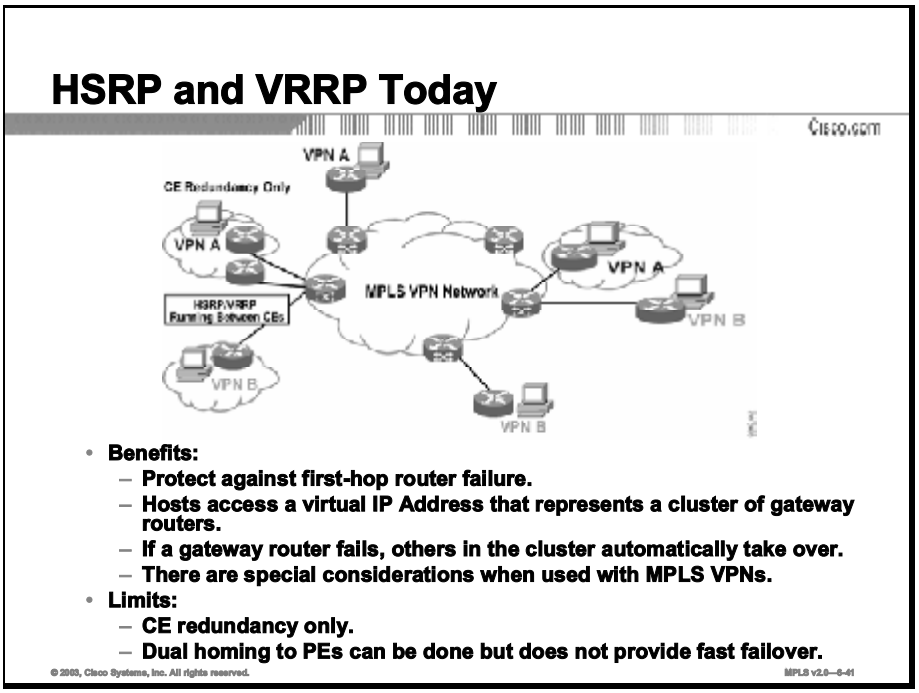

Today, service providers face challenges when it comes to implementing an efficient redundancy scheme. To address many of these issues, Cisco has offered HSRP and VRRP, which have provided several benefits:

- Both protocols protect against first-hop router failure.
- Hosts access a virtual IP address that represents a cluster of gateway routers.
- If a gateway router fails, others in the cluster automatically take over.

However, when implemented in a VPN environment, these protocols have limitations:

- They provide CE redundancy only.
- Dual homing to PEs can be done but does not provide fast failover.

HSRP and VRRP Today (Cont.)
Support for MPLS VPNs

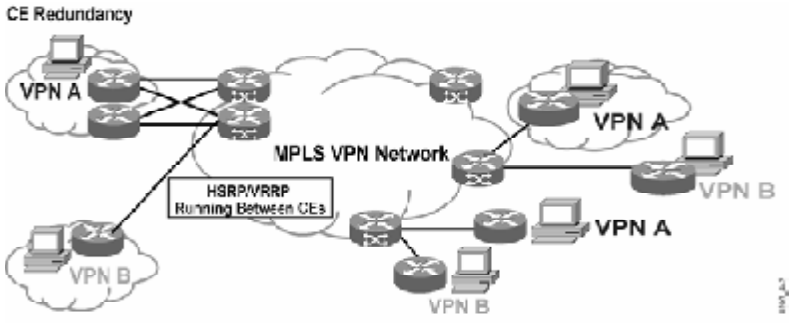

- **Improved network availability**
- **HSRP:**
  - **Transparent network topology modifications**
  - **Simple, centralized control of hot standby parameters**
- **VRRP:**
  - **Standards-based protocol**
  - **The flexibility to choose the protocol that best suits each environment**

Cisco MPLS for Managed Shared Services also provides HSRP support on MPLS VPN interfaces. This feature provides transparent "first-hop IP routing" redundancy for workstations or routers connected to interfaces within the MPLS VPN. HSRP creates a hot standby router group with a lead router that services all packets sent to the hot standby address. Other routers within the group monitor the lead router. If the lead fails, a standby router inherits the lead position as well as the hot standby address. The HSRP protocol allows specification of active routers, pre-emption delays, hold times, and interface status tracking.

The benefits of HSRP for MPLS VPNs include the following:

- Improved network availability
- Transparent network topology modifications
- Simple, centralized control of hot standby parameters

Similar to HSRP, VRRP allows a group of routers to function as one virtual router. Cisco MPLS for Managed Shared Services includes VRRP for MPLS VPNs by enabling the group of routers to share one virtual IP address and one virtual MAC address. One master router performs packet forwarding for the local hosts, and the rest of the routers within the group can act as backup routers to protect from failures of the master. With VRRP, the backup routers stay idle as far as packet forwarding is concerned.

The benefits of VRRP for MPLS VPNs include:

- Improved network availability
- Standards-based protocol
- The flexibility to choose the protocol that best suits each environment

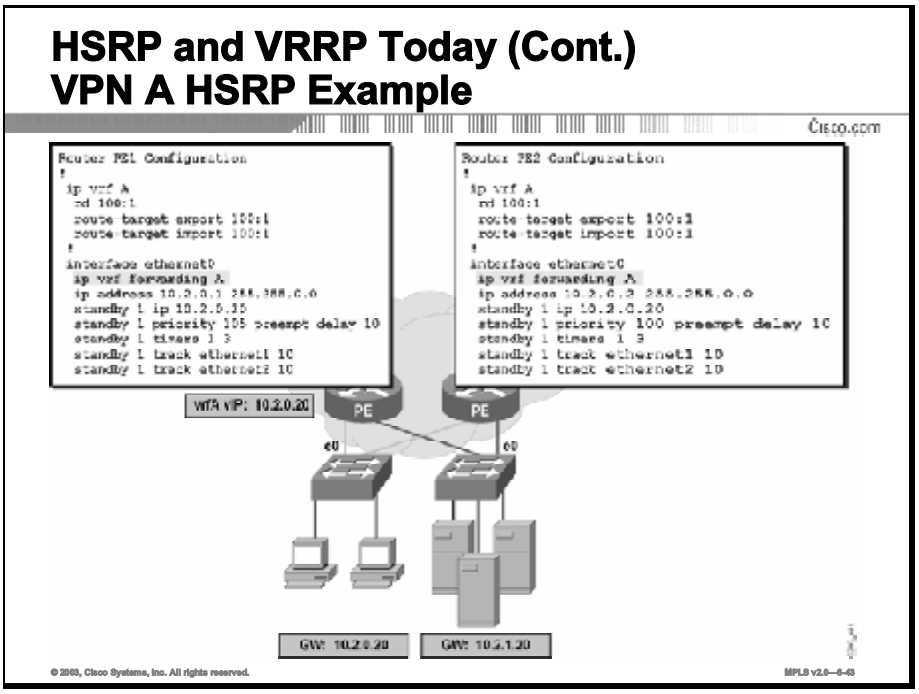

The creation of a VPN-aware HSRP is a combination of the standard MPLS VPN and HSRP commands. In this example, a VPN-aware HSRP is created for VPN A. MPLS forwarding has been enabled on Ethernet0, and the virtual IP address has been configured as 10.2.0.20.

Multicast VPNs

This topic describes the implementation of managed multicast VPN services in an MPLS VPN environment.

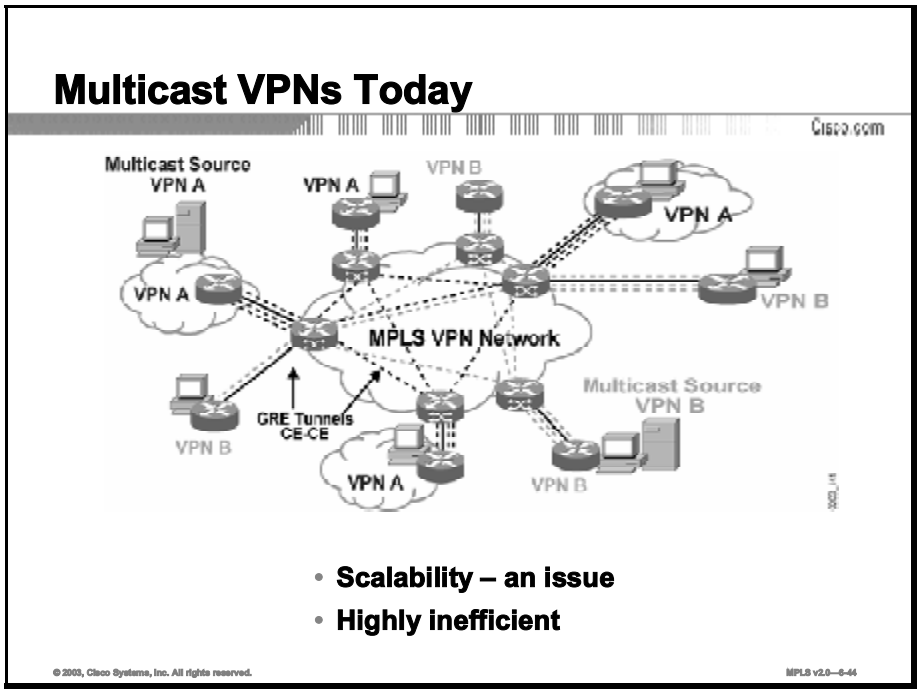

Historically, IPinIP generic routing encapsulation (GRE) tunnels were the only way to connect through a service provider network. Although such tunneled networks tend to have scalability issues, they represented the only means of passing IP multicast traffic through a VPN.

MPLS was derived from tag switching, and various other vendor methods of IP switching support enhancements in the scalability and performance of IP-routed networks by combining the intelligence of routing with the high performance of switching. MPLS is now used for VPNs, which is an appropriate combination because MPLS decouples information used for forwarding of the IP packet (the label) from the information carried in the IP header.

A multicast VPN allows an enterprise to transparently interconnect its private network across the network backbone of a service provider. The use of a multicast VPN to interconnect an enterprise network in this way does not change the way that the enterprise network is administered nor does it change general enterprise connectivity.

The multicast VPN feature in Cisco IOS software provides the ability to support the multicast feature over a Layer 3 VPN. As enterprises extend the reach of their multicast applications, service providers can accommodate these enterprises over their MPLS core network. IP multicast is used to stream video, voice, and data to an MPLS VPN network core.

A VPN represents network connectivity across a shared infrastructure, such as an Internet service provider (ISP). Its function is to provide the same policies and performance as a private network, at a reduced cost of ownership, thus creating many opportunities for cost savings through operations and infrastructure.

Multicast VPNs Today (Cont.)

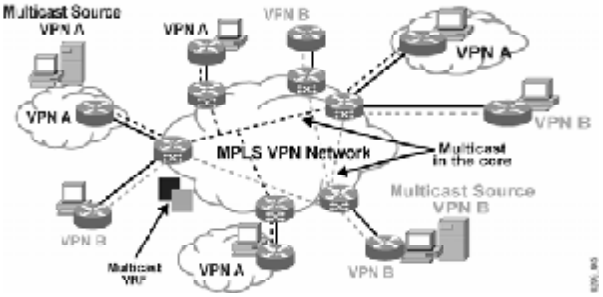

- **Enabling service providers with MPLS VPN networks to offer multicast services to their enterprise clients**
- **Minimizing configuration time and complexity—configuration is required only at edge routers**
- **Ensuring transparency of the service provider network**
- **Providing the ability to easily build advanced enterprise-friendly services such as virtual multicast networks**
- **Increasing network scalability**

By implementing native multicast functionality inside their MPLS VPN networks, service providers can now monetize multicast services. Service providers can utilize current resources to support bandwidth-hungry streaming services such as telecommuting, videoconferencing, e-learning, and a host of other business applications. Cisco multicast VPN technology helps improve the efficiency of the bandwidth-hungry applications of enterprise networks by eliminating the packet replication and performance issues associated with distribution of multicast traffic.

Multicast VPNs benefit service providers by accomplishing the following:

- Enabling service provider with MPLS VPN networks to offer multicast services to their enterprise clients
- Minimizing configuration time and complexity—configuration is required only at edge routers
- Ensuring transparency of the service provider network
- Providing the ability to easily build advanced enterprise-friendly services such as virtual multicast networks
- Increasing network scalability

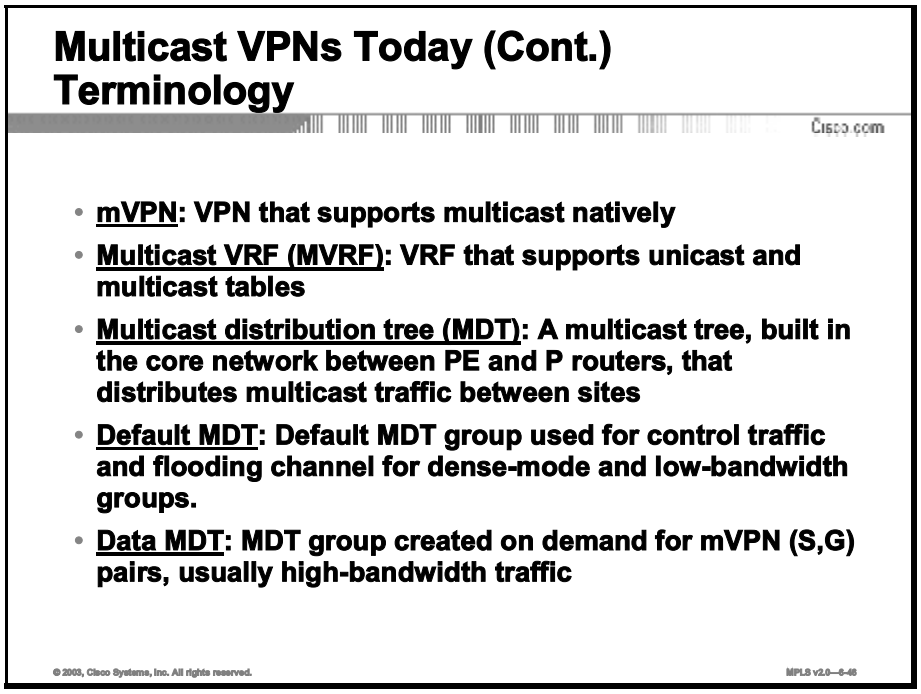

VPN-aware multicast technology has introduced a new set of terminology.

Multicast VPNs introduce multicast routing information to the VRF table. When a PE router receives multicast data or control packets from a CE router, forwarding is performed according to information in the multicast virtual routing and forwarding instance (MVRF).

A set of MVRFs that can send multicast traffic to each other constitutes a multicast domain. For example, the multicast domain for a customer that wanted to send certain types of multicast traffic to all global employees would consist of all CE routers associated with that enterprise.

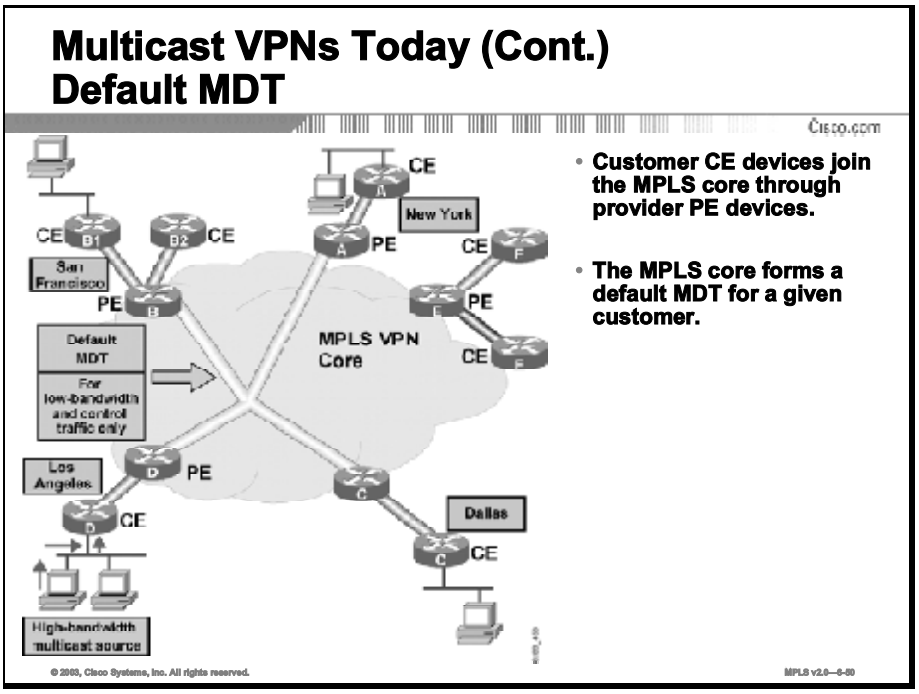

Multicast VPNs establish a static default MDT for each multicast domain. The default MDT defines the path used by PE routers to send multicast data and control messages to every other PE router in the multicast domain.

In the example here, a service provider has a multicast customer with offices in San Francisco, Los Angeles, New York, and Dallas. A one-way multicast presentation is occurring in Los Angeles. The service provider network supports all three sites associated with this customer, in addition to the site of a different enterprise customer.

The default MDT for the enterprise customer consists of provider routers and their associated PE routers. The other PE router is not part of the default MDT, because it is associated with a different customer.

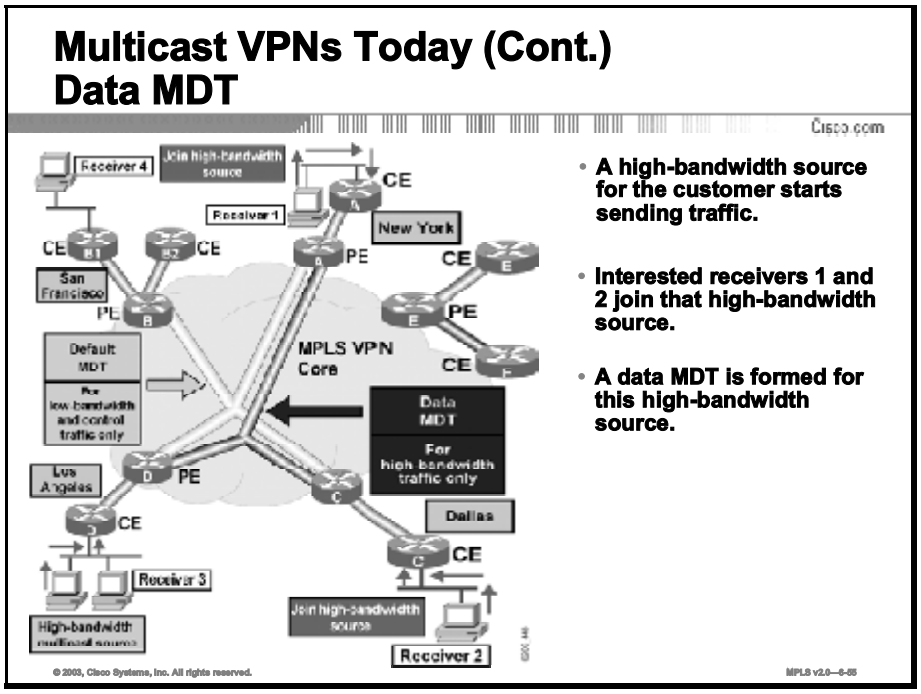

Multicast VPNs also support the dynamic creation of MDTs for high-bandwidth transmission. Data MDTs are a feature unique to Cisco IOS software. Data MDTs are intended for high-bandwidth sources such as full-motion video inside the VPN to ensure optimal traffic forwarding in the MPLS VPN core. The threshold at which the data MDT is created can be configured on a per-router or a per-VRF basis.

When the multicast transmission exceeds the defined threshold, the sending PE router creates the data MDT and sends a User Datagram Protocol (UDP) message that contains information about the data MDT to all routers in the default MDT. The statistics to determine whether a multicast stream has exceeded the data MDT threshold are examined once every 10 seconds. If multicast distributed switching is configured, the time period can be up to twice as long.

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. They are not created for (*, G) entries regardless of the value of the individual source data rate.

In the example here, an employee joins the multicast session. The PE router associated with the employee site sends a join request that flows across the default MDT for the multicast domain of the customer. The PE router associated with the multicast session source receives the request.

The source CE router begins to send the multicast data to the associated PE router, which sends the multicast data along the default MDT. Immediately after sending the multicast data, the source PE router recognizes that the multicast data exceeds the bandwidth threshold at which a data MDT should be created. Therefore, the PE router creates a data MDT, sends a message to all routers using the default MDT that contains information about the data MDT, and, 3 seconds later, begins sending the multicast data for that particular stream using the data MDT. Only the PEs routers that have interested receivers for this source will join the data MDT and receive traffic on it.

PE routers maintain a Protocol Independent Multicast (PIM) relationship with other PE routers over the default MDT, and a PIM relationship with their directly attached PE routers.

The figure depicts the final flow of multicast data sourced from the multicast sender in Los Angeles to the multicast clients in New York and Dallas. Multicast data sent from the multicast sender is delivered in its original format to its associated PE router using sparse mode, bidir, or Source Specific Multicast (SSM). This PE router then encapsulates the multicast data and sends it across the data MDT using the configured MDT data groups. The mode used to deliver the multicast data across the data MDT is determined by the service provider and has no direct correlation with the mode used by the customer. The PE router in New York receives the data along the data MDT. That PE router de-encapsulates the packet and forwards it in its original format toward the multicast client using the mode configured by the customer.

Multicast VPNs Today (Cont.)
Solution Concept

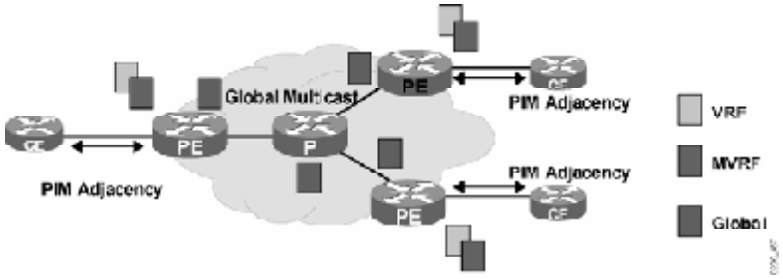

- P and PE routers multicast-enabled.
- Global multicast routing tables in the provider network.
- Globally, PEs configured to run PIM (global instance) with adjacent P routers.
- Multicast-enabled VPNs have a VPN multicast routing table (MVRF).
- No requirement to run same multicast protocols in the customer and provider network.
- If PIM is configured, PEs maintain PIM adjacencies with CE devices:
  - No PIM adjacency between CEs devices non directly connected
- Normal PIM configuration in customer network:
  - RPs, and so on

For every multicast domain of which an MVRF is a part, the PE router creates a multicast tunnel interface. A multicast tunnel interface is an interface that the MVRF uses to access the multicast domain. It can be thought of as a conduit that connects an MVRF and the global MVRF. One tunnel interface is created per MVRF.

To support VPN-aware multicast systems, PIM multicast (PIM, SSM, and so on) capability must be enabled on all affected P and PE routers. This addition results in a global multicast routing table being created in the provider network routers. The PE routes that have been configured to run PIM (global instance) will establish a PIM adjacency with neighboring P routers.

Multicast-enabled VPNs will create a VPN multicast routing table (MVRF).

There is no requirement to run the same multicast protocols in the customer and provider network. If PIM is configured as the CE-to-PE multicast protocol, the PE devices maintain PIM adjacencies with CE devices. No PIM adjacency will be established between CE devices that are not directly connected.

Normal PIM configuration (for example, rendezvous point [RP], RR, confederations) is accomplished in the customer network.

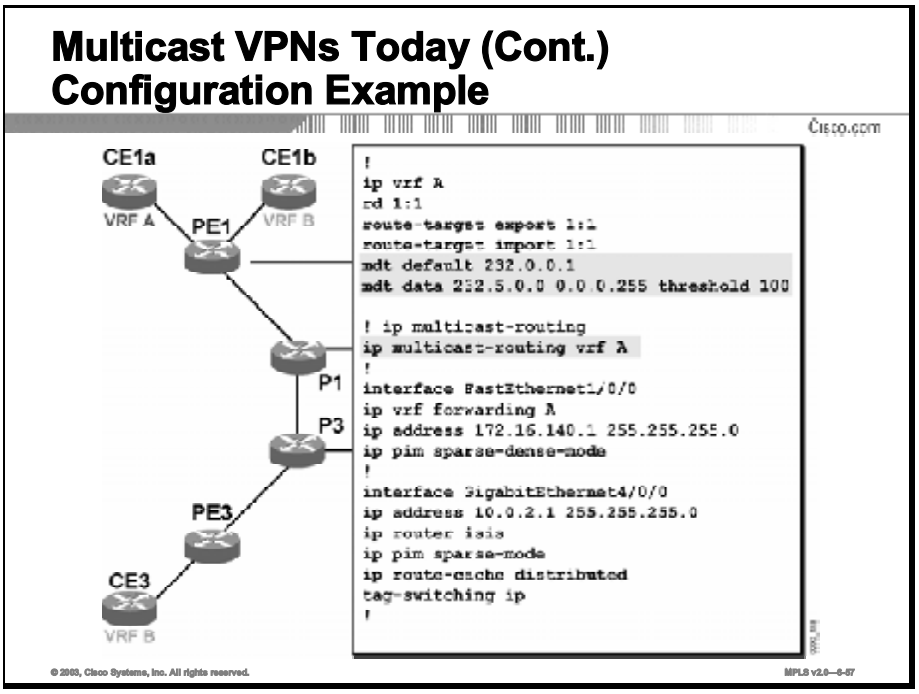

Configuring VPN-aware multicast capability is a combination of standard VPN, standard multicast, and new VPN-aware multicast commands.

Enabling a VPN for Multicast Routing

This task enables a VPN for multicast routing.

PIM

PIM can operate in dense mode or sparse mode. It is possible for the router to handle both sparse groups and dense groups at the same time.

In dense mode, a router assumes that all other routers want to forward multicast packets for a group. If a router receives a multicast packet and has no directly connected members or PIM neighbors present, a prune message is sent back to the source. Subsequent multicast packets are not flooded to this router on the pruned branch. PIM builds source-based multicast distribution trees.

In sparse mode, a router assumes that other routers do not want to forward multicast packets for a group, unless there is an explicit request for the traffic. When hosts join a multicast group, the directly connected routers send PIM join messages toward the RP. The RP keeps track of multicast groups. Hosts that send multicast packets are registered with the RP by the first-hop router of that host. The RP then sends join messages toward the source. At this point, packets are forwarded on a shared distribution tree. If the multicast traffic from a specific source is sufficient, the first-hop router of the host may send join messages toward the source to build a source-based distribution tree.

Fast Switching and IP Multicast

Fast switching of IP multicast packets is enabled by default on all interfaces (including GRE and Distance Vector Multicast Routing Protocol [DVMRP] tunnels), with one exception: It is disabled and not supported over X.25 encapsulated interfaces. Note the following properties of fast switching:

- If fast switching is disabled on an incoming interface for a multicast routing table entry, the packet is sent at process level for all interfaces in the outgoing interface list.
- If fast switching is disabled on an outgoing interface for a multicast routing table entry, the packet is process-level switched for that interface, but may be fast-switched for other interfaces in the outgoing interface list.

Disable fast switching if you want to log debug messages, because when fast switching is enabled, debug messages are not logged.

Prerequisites

You must enable PIM sparse mode on the interface that is used for BGP peering. Configure PIM on all interfaces used for IP multicast. Cisco recommends configuring PIM sparse mode on all physical interfaces of PE routers connecting to the backbone. Cisco also recommends configuring PIM sparse mode on all loopback interfaces if they are used for BGP peering or if their IP address is used as an RP address for PIM.

In order to be able to use Auto-RP within a VRF, the interface facing the CE must be configured for PIM sparse-dense mode.

In this example, multicast and MPLS capability have been enabled at the global and interface levels using the standard commands. A VPN named "A" has been created and multicast-enabled using the new **ip multicast-routing vrf** command. The default MDT group for a VRF is configured using the **mdt default** command. The multicast group address range for data MDT groups is configured using the **mdt data** *group-address-range wildcard-bits* [**threshold** *threshold-value*] [**list** *access-list*] command.

For further information see the following:

- http://www.cisco.com/en/US/partner/tech/tk828/tk363/tech_digest09186a00801a64a3.html
- http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/products_feature_guide09186a00801039b0.html

Summary

This topic summarizes the key points discussed in this lesson.

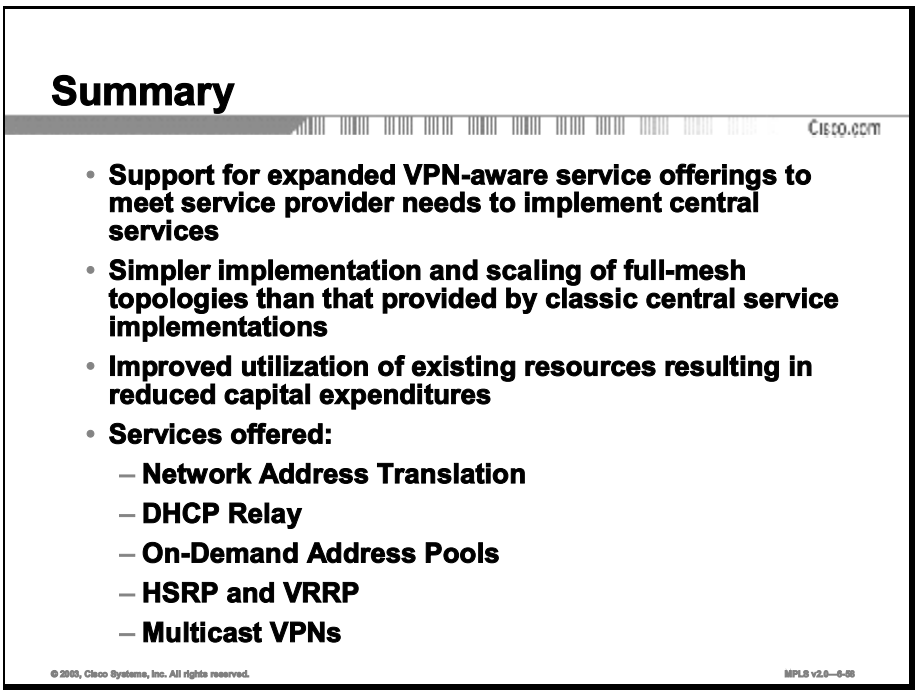

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about Cisco managed services.

Next Steps

For the associated lab exercise, refer to these sections of the course Lab Guide:

- Lab Exercise 6-1: Overlapping VPNs
- Lab Exercise 6-2: Merging Service Providers
- Lab Exercise 6-3: Common Services VPN

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are the steps that you need to take to enable a VPN-aware NAT service on an existing MPLS network?

Q2) When you are implementing a VPN-aware DHCP relay service, which command do you use to configure the DHCP server address?

Q3) When you are implementing a VPN-aware DHCP relay service, how does the DHCP server know which VPN that a request comes from?

Q4) When you are implementing a VPN-aware ODAP service, where does the DHCP server get its initial address pool from?

A) It requests it from the ODAP server.

B) It requests it from the DHCP relay agent.

C) It is configured with the **ip nat pool** command on the DHCP server.

D) It is configured with the **ip dhcp pool** command on the ODAP server.

Q5) How is a VPN-aware HSRP or VRRP service implemented?

Q6) Which type of traffic flows over the default MDT?

Q7) What is a data MDT?

Quiz Answer Key

Q1) Create the NAT address pool (**ip nat pool**)
Assign the address pool to the VRF (**ip nat inside source**)
Enable NAT on the interfaces (**ip nat outside** and **ip nat inside**)
Create a static address to the next hop (**ip route vrf**)

Relates to: Network Address Translation

Q2) **ip helper-address vrf**

Relates to: DHCP Relay

Q3) The DHCP relay agent inserts the VPN ID into the DHCP request.

Relates to: DHCP Relay

Q4) A

Relates to: On-Demand Address Pools

Q5) Enable a VRF on the interface and then configure the HSRP or VRRP service using the standard commands.

Relates to: HSRP and VRRP

Q6) Control traffic and flooding channel for dense-mode and low-bandwidth groups.

Relates to: Multicast VPNs

Q7) an MDT group created on demand for mVPN (S,G) pairs, usually high-bandwidth traffic

Relates to: Multicast VPNs

Module 7

Internet Access from an MPLS VPN

Overview

Integrating Internet access with a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) solution is one of the most common service provider business requirements. This module provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Module Objectives

Upon completing this module, you will be able to describe the various Internet access implementations available and the benefits and drawbacks of each model. This ability includes being able to do the following:

- Describe MPLS VPN Internet topologies
- Describe VPN Internet access implementation methods
- Describe methods to separate Internet access from VPN services
- Describe the characteristics of Internet access solutions in which Internet access is provided through a separate VPN

Module Outline

The module contains these lessons:

- VPN Internet Access Topologies
- VPN Internet Access Implementation Methods
- Separating Internet Access from VPN Services
- Internet Access as a Separate VPN

VPN Internet Access Topologies

Overview

This lesson identifies the characteristics of the different modules that are used to combine Internet access with VPN services.

Relevance

This lesson is a crucial one for learners planning to enhance their usage of network resources using MPLS VPNs.

Objectives

This lesson identifies the characteristics of the different modules that are used to combine Internet access with VPN services.

Upon completing this lesson, you will be able to:

- Describe the characteristics of classical Internet access for a VPN customer
- Describe the characteristics of Internet access from every customer site
- Describe the characteristics of Internet access through a central firewall service
- Describe the characteristics of a wholesale Internet access service

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Complex MPLS VPNs” module of this course

Outline

This lesson includes these topics:

- Overview
- Classical Internet Access for a VPN Customer
- Internet Access from Every Customer Site
- Internet Access Through a Central Firewall Service
- Wholesale Internet Access
- Summary
- Quiz

Classical Internet Access for a VPN Customer

This topic describes the typical way that VPN customers who have not implemented MPLS access the Internet.

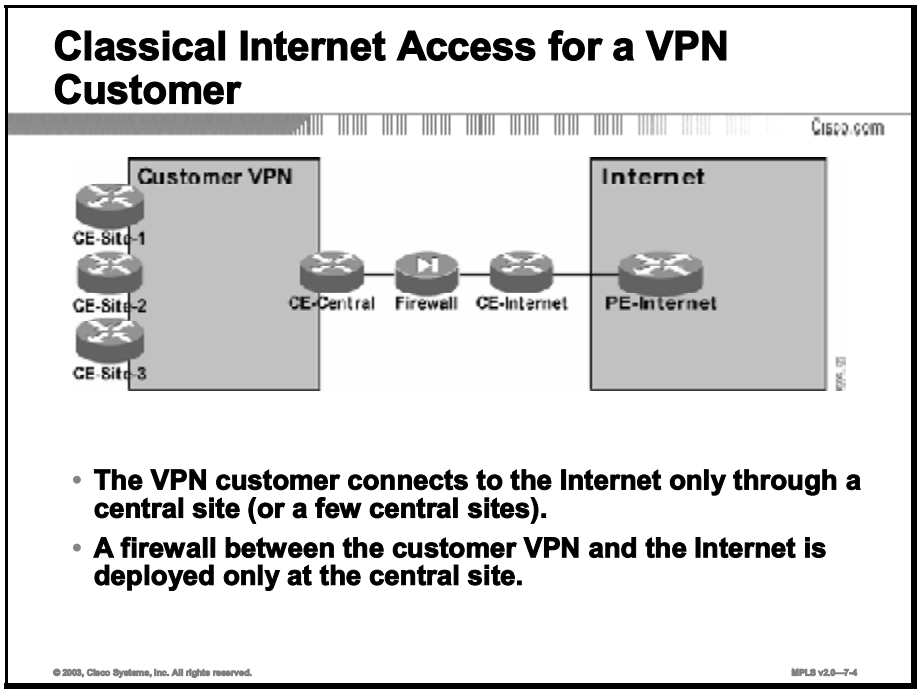

Classical Internet access is implemented through a (usually central) firewall that connects the customer network to the Internet in a secure fashion. The private network of the customer (or VPN if the customer is using a VPN service) and the Internet are connected only through the firewall.

Classical Internet Access Addressing

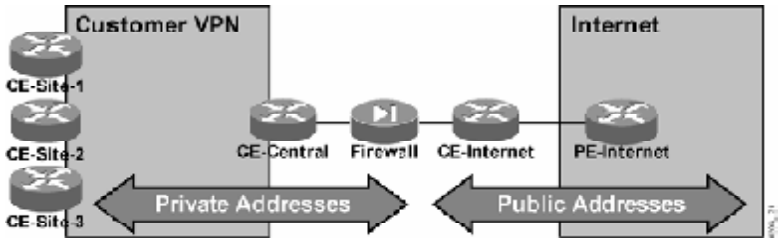

- **Customer can use private address space.**
- **The firewall provides Network Address Translation (NAT) between the private address space and the small portion of public address space assigned to the customer.**

Addressing requirements for this type of connection are very simple:

- The customer is assigned a small block of public address space used by the firewall.
- The customer typically uses private addresses inside the customer network.
- The firewall performs Network Address Translation (NAT) between the private addresses of the customer and the public addresses assigned to the customer by the Internet service provider (ISP). Alternatively, the firewall might perform an application-level proxy function that isolates private and public IP addresses.

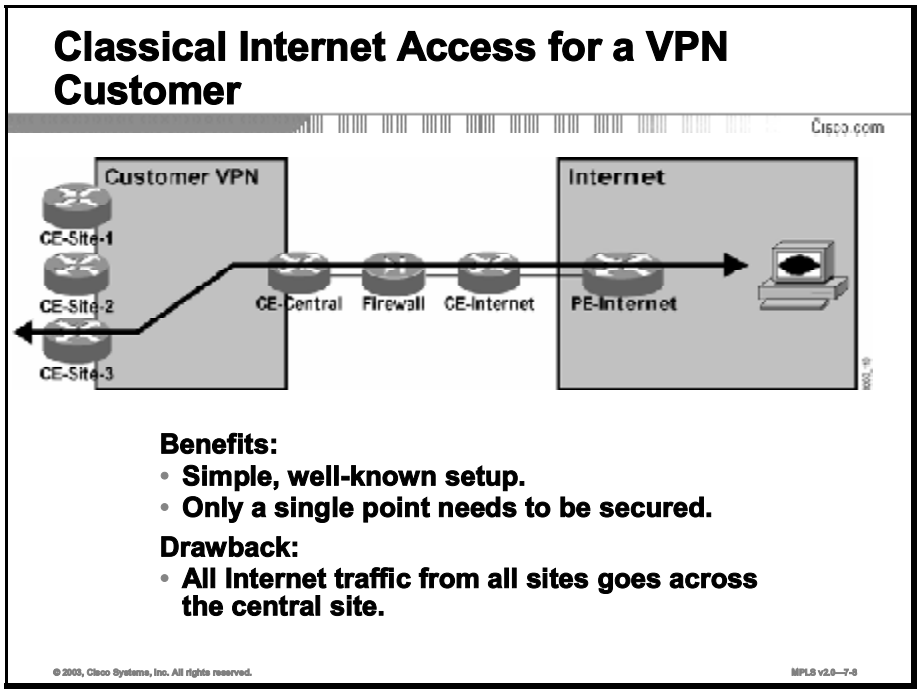

There are a number of benefits associated with the classical design:

- It is a well-known setup used worldwide for Internet connectivity from a corporate network. Access to expertise needed to implement such a setup is thus simple and straightforward.
- There is only one interconnection point between the secure customer network and the Internet. Security of the Internet access needs to be managed only at this central point.

The major drawback of this design is the traffic flow—all traffic from the customer network to the Internet has to pass through the central firewall. While this might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when the users are geographically separated.

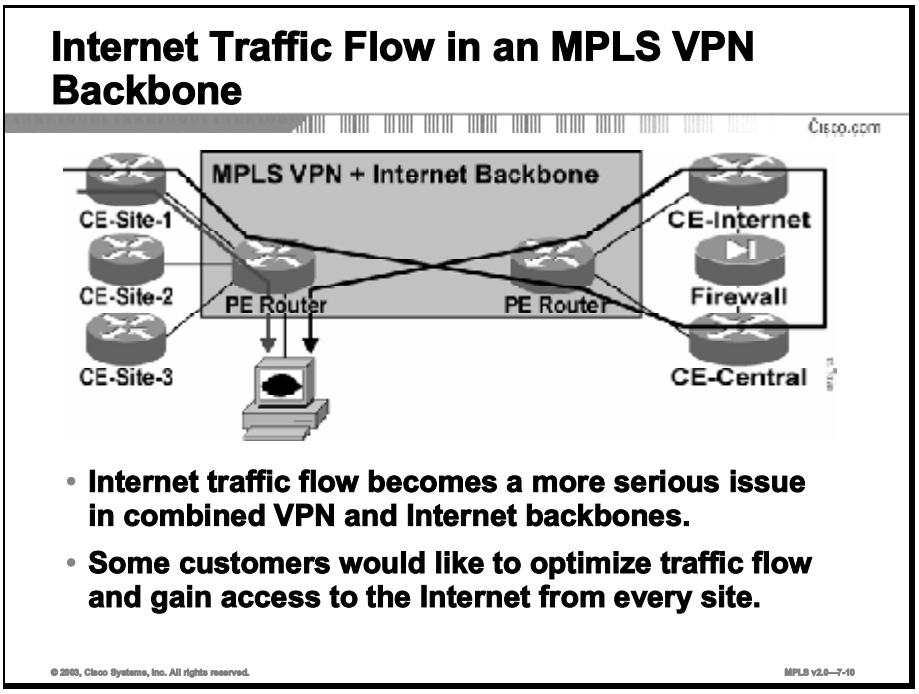

The traffic flow issue becomes even more pronounced when the customer VPN (based on, for example, MPLS VPN services) and the Internet traffic share the same service provider backbone. In this case, the traffic from a customer site may have to traverse the service provider backbone as VPN traffic, and then return into the same backbone by the corporate firewall, ending up at a server very close to the original site.

Based on this analysis, the drawbacks of the central firewall design can be summarized as follows:

- The link between the central site and the provider backbone has to be over-dimensioned, because it has to transport all of the customer Internet traffic.
- The provider backbone is overutilized, because the same traffic crosses the backbone twice, first as VPN traffic and then as Internet traffic (or vice versa).
- Response times and quality of service (QoS) may suffer because the traffic between the customer site and an Internet destination always has to cross the central firewall, even when the Internet destination is very close to the customer site.

These drawbacks have prompted some large users and service providers to consider alternate designs in which every customer site can originate and receive Internet traffic directly.

Internet Access from Every Customer Site

This topic describes ways in which customers may access the Internet from their own sites.

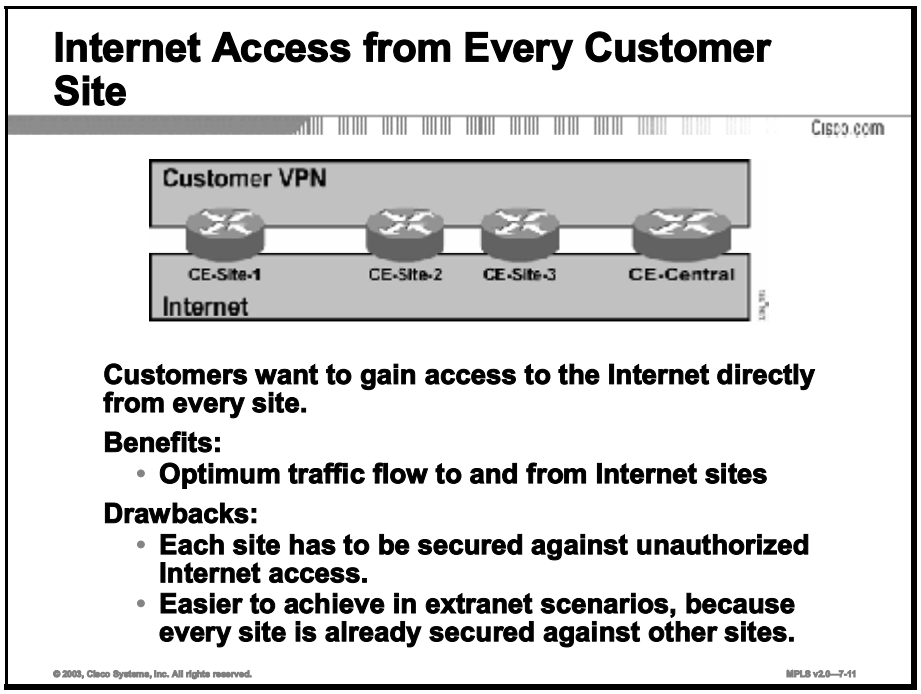

To bypass the limitations of Internet access through a central firewall, some customers are turning toward designs in which each customer site has its own independent Internet access. While this design clearly solves all traffic flow issues, the associated drawback is higher exposure—each site has to be individually secured against unauthorized Internet access. This design is applicable primarily for larger sites (concentrating traffic from nearby smaller sites) or for extranet VPNs, in which each site is already secured against the other sites participating in the extranet VPN.

Internet Access from Every Site—Addressing

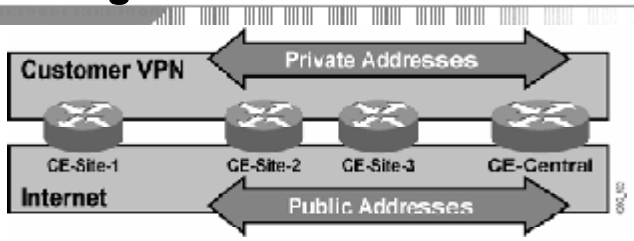

Two addressing options:

- **Every CE router performs NAT functionality—a small part of public address space has to be assigned to each CE router.**
- **Customer uses only public IP addresses in the private network—not realistic for many customers.**

In order to gain Internet access from every site, each site requires at least some public IP addresses. Two methods can be used to achieve this goal:

- A small part of public address space can be assigned to each customer site. NAT between the private IP addresses and the public IP addresses needs to be performed at each site.
- If the customer is already using public IP addresses in the VPN, NAT functionality is not needed. Unfortunately, this option is open only to those customers that own large address blocks of public IP addresses.

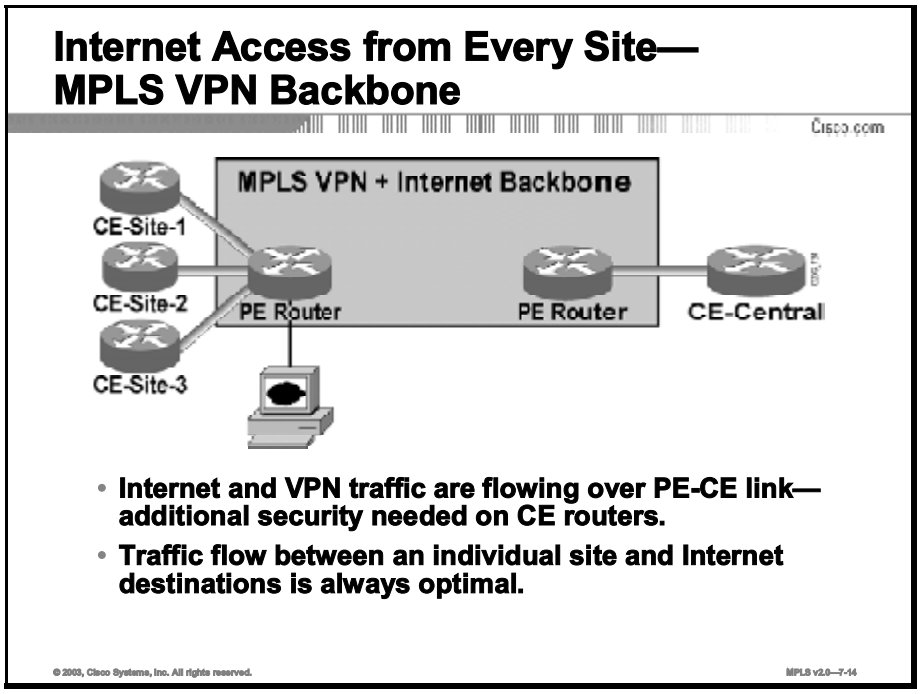

To achieve Internet access from every customer site, each customer edge (CE) router must forward VPN traffic toward other customer sites, and forward Internet traffic toward Internet destinations. The two traffic types are usually sent over the same physical link to minimize costs. Switched WAN encapsulation (Frame Relay or ATM) can be used to separate the VPN and Internet traffic onto different virtual circuits, or the traffic can share the same logical link as well, resulting in reduced security. On the other hand, the weaker (and less complex) security of this design is offset by optimal traffic flow between every site and Internet destinations.

Internet Access Through a Central Firewall Service

This topic describes how Internet access is obtained when a central firewall service has been implemented.

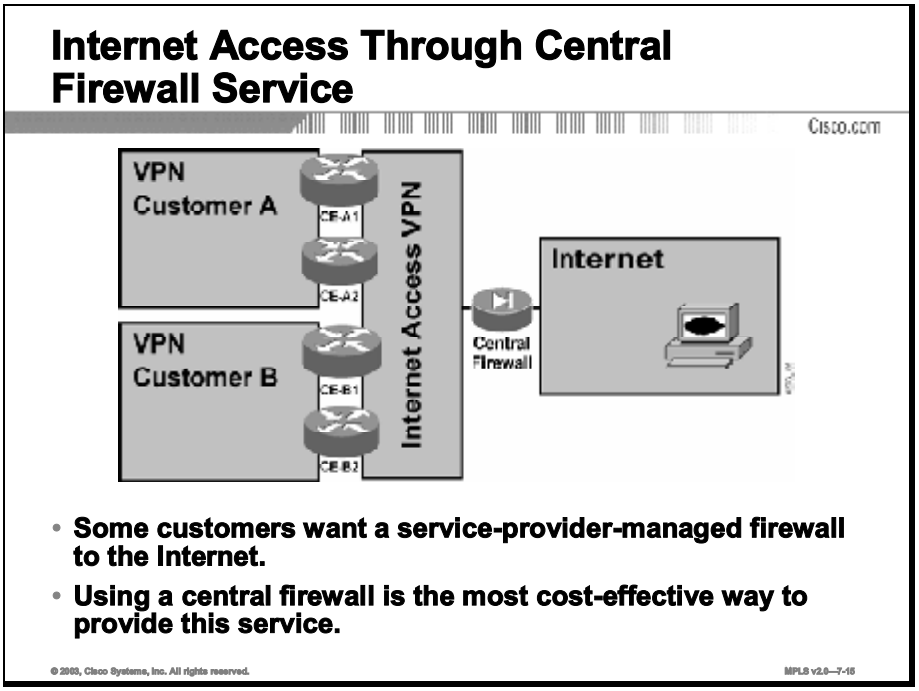

For customers who do not want the complexity of managing their own firewall, a managed firewall service offered by the service provider is a welcome relief. These customers typically want the service provider to take care of the security issues of their connection to the Internet.

The service provider can implement the managed firewall service by deploying a dedicated firewall at each customer site or (for a more cost-effective approach) by using a central firewall that provides secure Internet access to all customers.

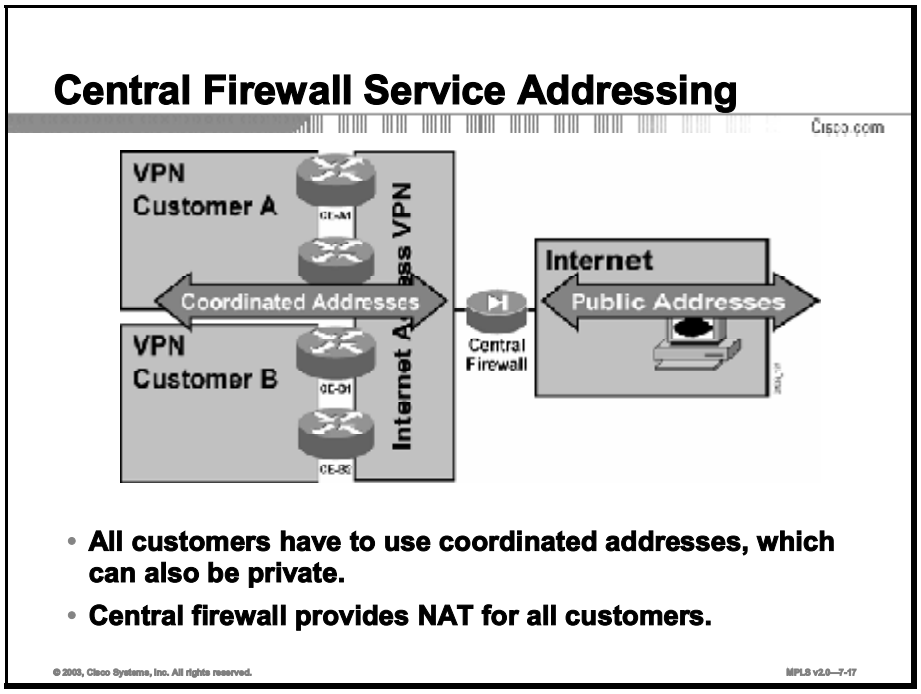

The central firewall, hosted by the service provider, has to use public addresses toward the Internet. Private addresses can be used between the central firewall and the individual customers. However, these addresses need to be coordinated between the service provider and the customers to prevent routing conflicts and overlapping addresses visible to the central firewall. Customers using a central firewall service are thus limited to IP addresses assigned to them by the service provider, much in the same way that Internet customers are limited to the public IP addresses assigned by their ISP.

Central Firewall Service Addressing (Cont.)

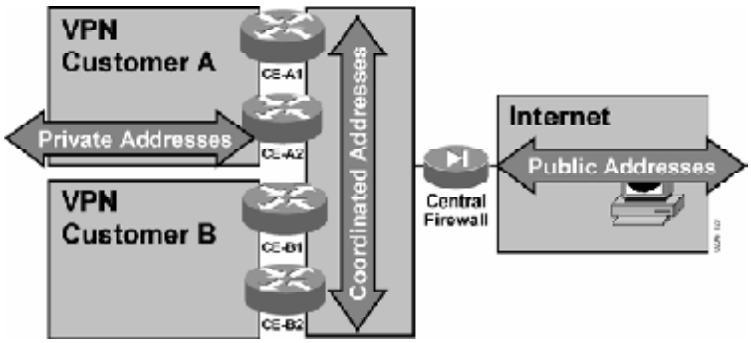

- **Each customer can use private address space if the CE routers provide address translation between private and coordinated address space.**

Customers of a central firewall service that still want to retain their own private addresses inside their network can use NAT on the CE routers, connecting their private network to the transit network that links customer sites to the central firewall.

Note	Service providers usually use private IP addresses as the address space between the central firewall and the customers. There is always a potential for overlapping addresses between the coordinated address space and the address space of an individual customer. The CE device providing NAT functionality therefore has to support address translation between overlapping sets of IP addresses.

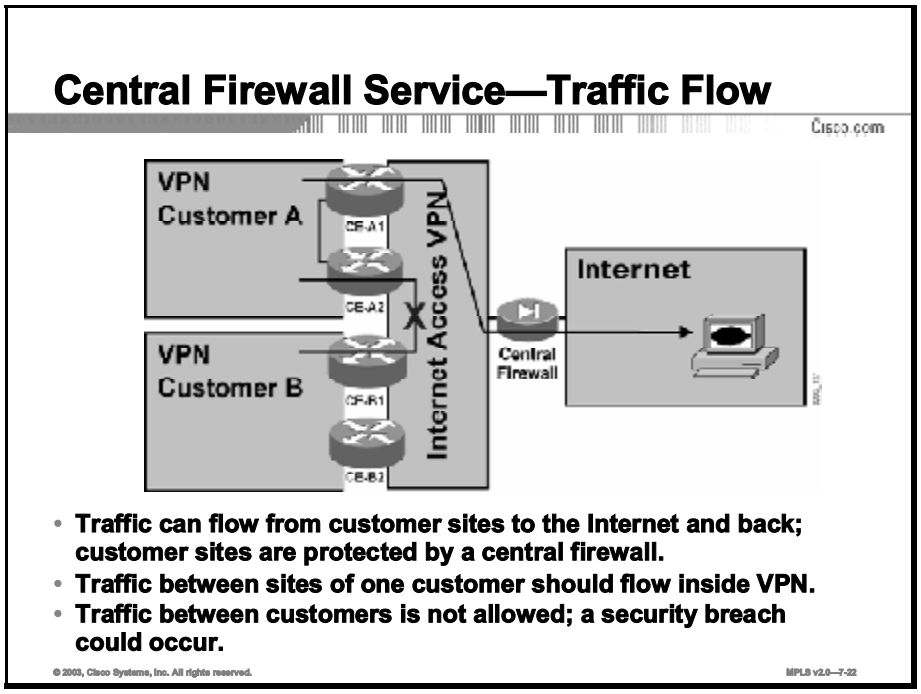

The traffic flow between sites participating in a central firewall service is limited by the security requirements of the service:

- Traffic between the customer sites and the Internet must flow freely, restricted only by the security functions of the central firewall.
- Traffic between sites of an individual customer should never flow across the VPN that links the customer sites with the central firewall. This traffic must flow inside the customer VPN.
- Traffic between customers using the central firewall is not allowed, because the individual customers are not protected from outside access (this is the task of the service provider, handled by the central firewall). Intercustomer traffic could lead to potential security problems.

Note	The restrictions on intercustomer traffic prevents customers from deploying publicly accessible servers in their networks, because these servers would not be available to other customers of the same service.

Wholesale Internet Access

This topic describes wholesale Internet access.

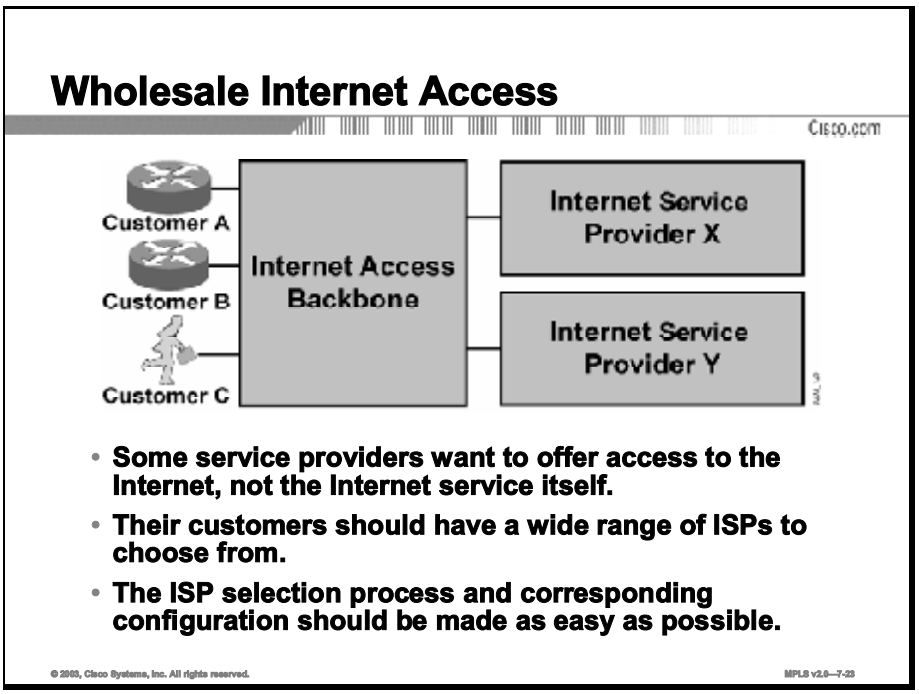

Similar to wholesale dial service (where an ISP uses the modem pools of other service providers) is the wholesale Internet access service, where an ISP uses the IP transport infrastructure of another service provider to reach the end users. The business model of this service varies—the end users might be customers of the service provider that owns the transport backbone (for example, a cable operator), who offers Internet access through a large set of ISPs as a value-added service. Alternatively, the service provider owning the Internet access backbone might act as a true wholesaler, selling transport infrastructure to Internet service providers, who then charge end users for the whole package.

When a service provider owns the backbone and provides Internet access to customers, it usually wants to offer a wide range of upstream ISPs to choose from, in order to satisfy various customer connectivity and reliability requirements. The selection of upstream ISPs and the corresponding configuration process should therefore be as easy as possible.

Wholesale Internet Access Addressing

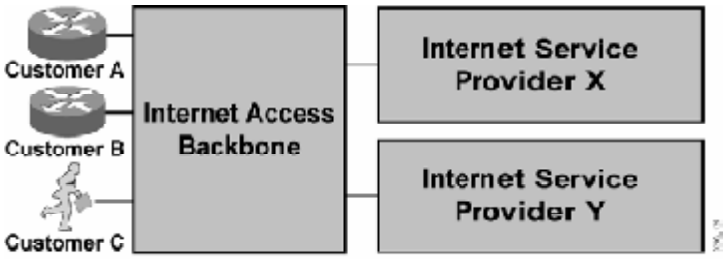

- **Customers get address space from the ISP that they connect to.**
- **When using dynamic addresses, the wholesale Internet access provider has to use a different address pool for every upstream service provider.**

Regardless of the business model used in the wholesale Internet access service, the addressing requirements are always the same: the upstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone. The wholesale Internet access provider consequently has to use a different address pool for every upstream ISP.

Summary

This topic summarizes the key points discussed in this lesson.

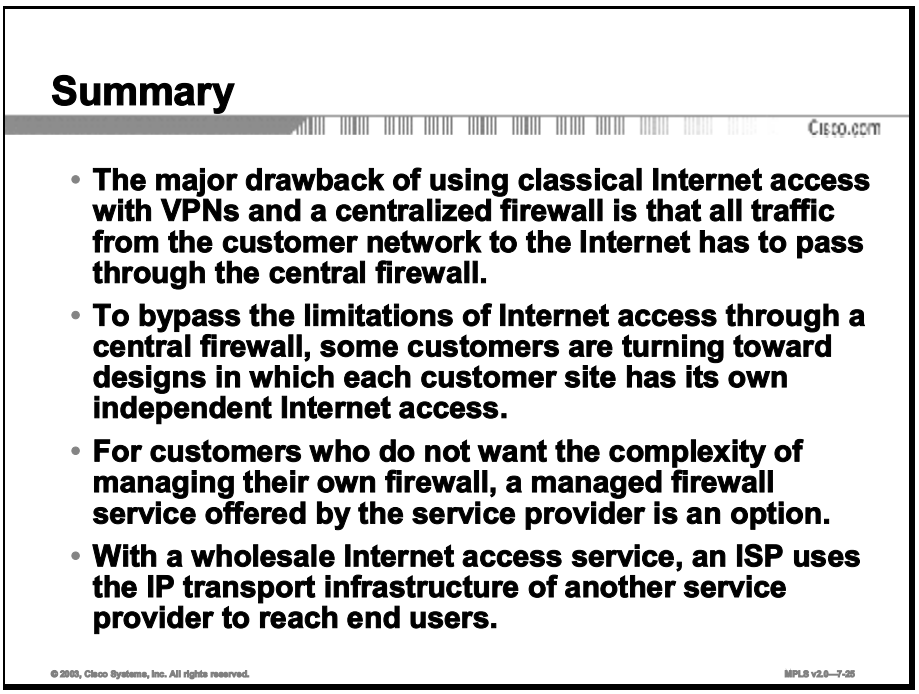

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about this topic.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is the major drawback of the classical Internet access model?

A) All of the customer traffic passes through the central firewall service.

B) None of the customer traffic passes through the central firewall service.

C) Only some of the customer traffic passes through the central firewall service.

D) There is no drawback.

Q2) Which two of the following statements are NOT correct regarding Internet access from every customer site? (Choose two.)

A) It is easier to achieve in intranet scenarios.

B) It provides optimum traffic flow to and from Internet sites.

C) It is more difficult to achieve in extranet scenarios.

D) Each site has to be secured against unauthorized Internet access.

Q3) Customers of a central firewall service who still want to retain their own private addresses inside their network can use NAT in which of the following ways?

A) NAT on the CE routers

B) NAT on the PE routers

C) NAT on both the CE and PE routers

D) NAT is not required if customers use their own private IP addresses.

Q4) When the wholesale Internet access solution is implemented, which of the following statements is correct?

A) The downstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone.

B) The upstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone.

C) Both the upstream and downstream ISPs must allocate a portion of their address space to the end users connected to the Internet access backbone.

D) None of the above is correct.

Quiz Answer Key

Q1) A

Relates to: Classical Internet Access for a VPN Customer

Q2) A, C

Relates to: Internet Access from Every Customer Site

Q3) A

Relates to: Internet Access Through a Central Firewall Service

Q4) B

Relates to: Wholesale Internet Access

VPN Internet Access Implementation Methods

Overview

This lesson identifies different design models for combining Internet access with VPN services. It lists the benefits and drawbacks of these models, and then explains the implications of their use.

Relevance

This lesson is a crucial one for learners planning to enhance their usage of network resources using MPLS VPNs.

Objectives

This lesson identifies the different design models used with VPN services.

Upon completing this lesson, you will be able to:

- Describe the major design models for combining Internet access with VPN services
- Describe the characteristics of Internet access in VPNs
- Describe the characteristics of Internet access through global routing
- Describe the characteristics of Internet access through a separate (sub)interface

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Complex MPLS VPNs” module of this course

Outline

This lesson includes these topics:

- Overview
- Major Design Models
- Internet Access in VPNs
- Internet Access Through Global Routing
- Internet Access Through Separate (Sub)interfaces
- Summary
- Quiz

Major Design Models

This topic describes two major design models for combining Internet access with VPN services.

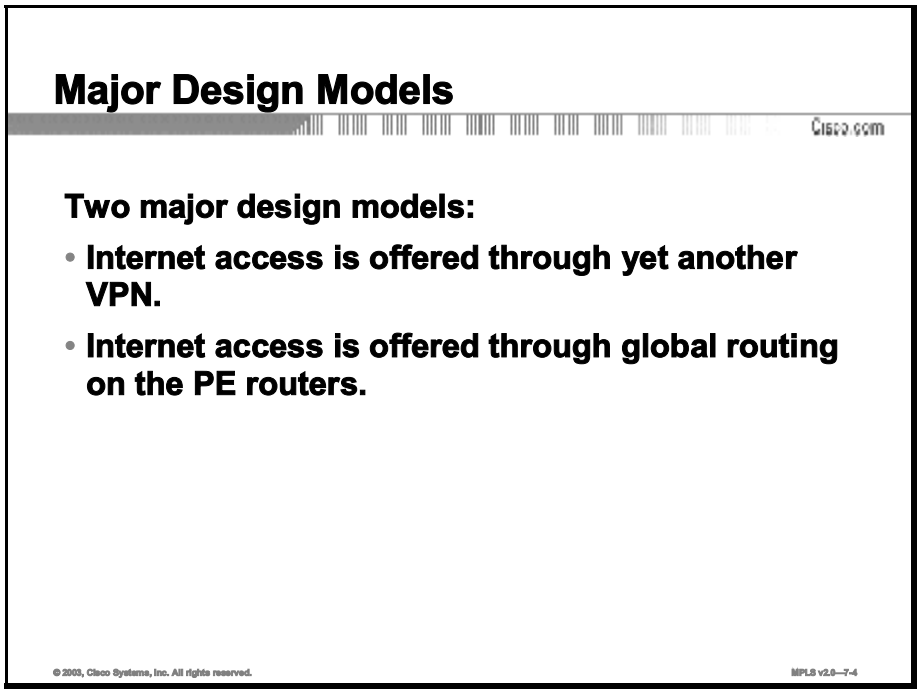

Network designers that want to offer Internet access and MPLS VPN services in the same MPLS backbone can choose between two major design models:

- Internet routing that is implemented as yet another VPN
- Internet routing that is implemented through global routing on the provider edge (PE) routers

Internet Access in VPNs

This topic describes the benefits and drawbacks of having Internet access in VPNs.

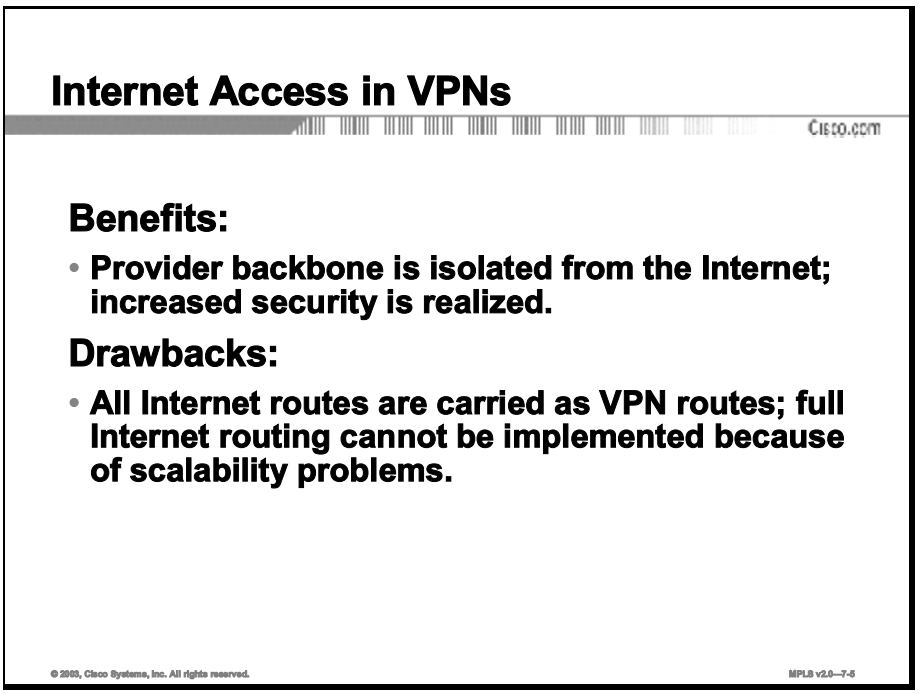

The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet, which results in increased security. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow the service providers to offer services that were simply not possible to implement with pure IP routing.

The obvious drawback of running the Internet as a VPN in the MPLS VPN architecture is the scalability of such a solution. An Internet VPN simply cannot carry full Internet routing because of scalability problems associated with carrying close to a hundred thousand routes inside a single VPN.

Internet Access Through Global Routing

This topic describes the implementation options for providing Internet access using global routing.

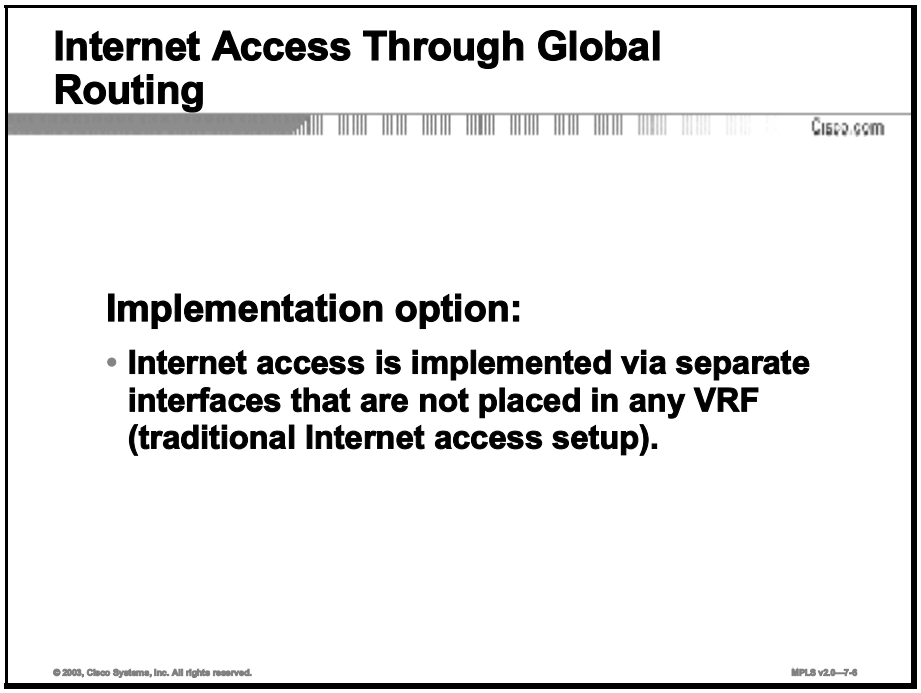

Implementing Internet access through global routing is identical to building an IP backbone offering Internet services. IP version 4 (IPv4) Border Gateway Protocol (BGP) is deployed between the PE routers to exchange Internet routes, and the global routing table on the PE routers is used to forward the traffic toward Internet destinations.

VPN customers can reach the global routing table by this method:

- The customers can use a separate logical link for Internet access. This method is equivalent to traditional VPN and Internet access.

Internet Access Through Separate (Sub)interfaces

This topic describes the use of separate interfaces and subinterfaces to provide Internet access.

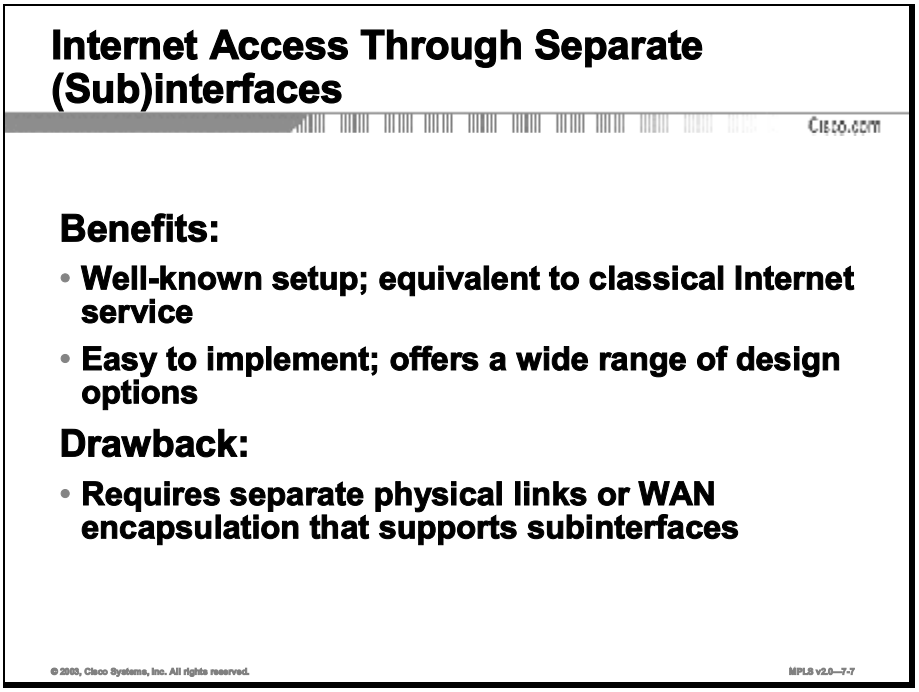

Internet access through separate logical links is easy to set up, because it is equivalent of the classical combination of Internet and VPN services that many customers are using today. This setup is also compatible with all the Internet services required by some customers (for example, the requirement to receive full Internet routing from a service provider).

The drawback of this design is the increased complexity, or cost, of the provider edge-customer edge (PE-CE) connectivity. Separation of Internet and VPN connectivity requires either two separate physical links or a single physical link with WAN encapsulation that supports subinterfaces (for example, Frame Relay).

Note	Some customers might be reluctant to change their encapsulation type to Frame Relay, because the IP QoS mechanisms on Frame Relay differ from those provided on PPP links.

Summary

This topic summarizes the key points discussed in this lesson.

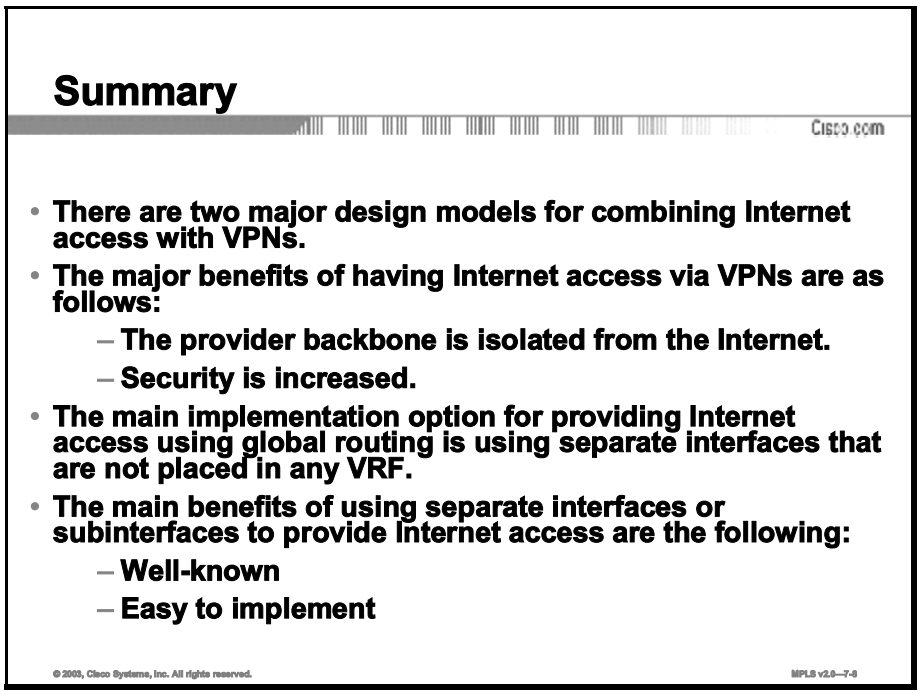

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about this topic.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The two major design models for implementing Internet access via VPNs are using another _____ and using _____ _____ on PE routers.

Q2) The two major benefits of using VPNs to provide Internet access are that the _____ _____ is isolated from the Internet and that _____ is increased.

Q3) When you are using global routing to provide Internet access, the main implementation option is that separate interfaces not be placed in a _____.

Q4) The main benefits of using separate interfaces or subinterfaces to provide Internet access are that they are _____ and are _____ to implement.

Quiz Answer Key

Q1) VPN, global routing

Relates to: Major Design Models

Q2) provider backbone, security

Relates to: Internet Access in VPNs

Q3) VRF

Relates to: Internet Access Through Global Routing

Q4) well-known, easy

Relates to: Internet Access Through Separate (Sub)interfaces

Separating Internet Access from VPN Services

Overview

This lesson describes the characteristics of an Internet access service in which the Internet access is totally separate from MPLS VPN services. It identifies the PE-CE requirements for this solution and how to implement the solution in an MPLS VPN network.

Relevance

This lesson is a crucial one for learners planning to enhance their usage of network resources using MPLS VPNs.

Objectives

Describe methods to separate Internet access from VPN services.

Upon completing this lesson, you will be able to:

- Describe the characteristics of Internet access separated from VPNs
- Describe the characteristics of implementing separate subinterfaces
- Describe the characteristics of classical Internet access for a VPN customer
- Describe the characteristics of Internet access from every customer site
- Identify the limitations of separate Internet access

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Complex MPLS VPNs” module of this course

Outline

This lesson includes these topics:

- Overview
- Designing Internet Access Separated from VPNs
- Implementing Separate Subinterfaces
- Classical Internet Access for a VPN Customer
- Internet Access from Every Customer Site
- Limitations of Separate Internet Access
- Summary
- Quiz

Designing Internet Access Separated from VPNs

This topic describes how Internet access is separated from traditional VPN services.

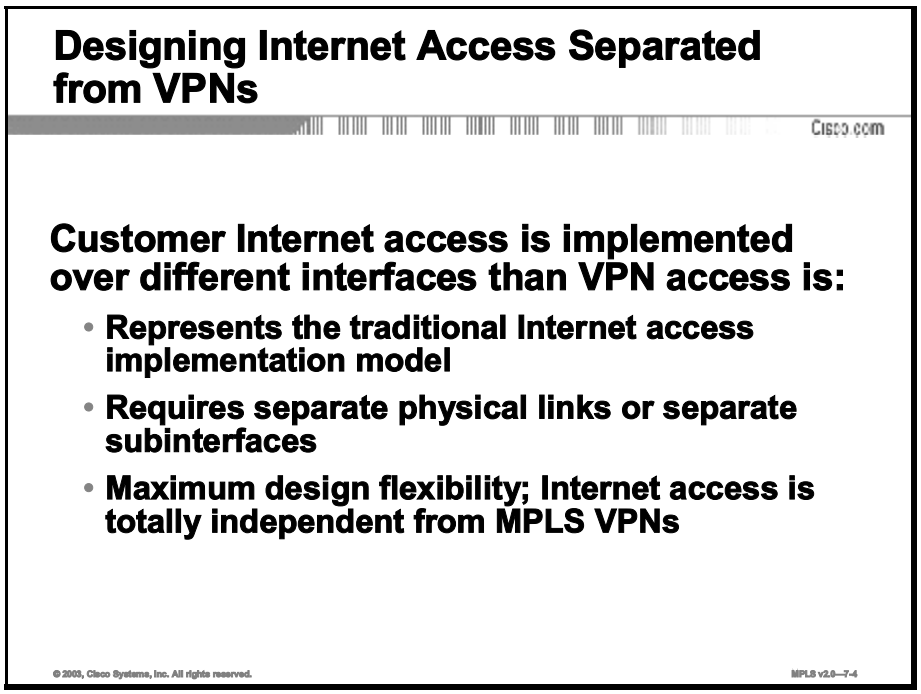

Internet access can always be implemented with the traditional implementation model, with two links between the customer site(s) and the service provider network: a VPN link and an Internet link. The two links can be implemented with one physical link if you use a Layer 2 encapsulation that supports subinterfaces (Frame Relay, ATM, or a VLAN).

The traditional Internet access implementation model provides maximum design flexibility, because the Internet access is completely separated from the MPLS VPN services. Nevertheless, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet access solutions such as wholesale Internet access.

Implementing Separate Subinterfaces

This topic describes some implementation methods for using subinterfaces to provide Internet access.

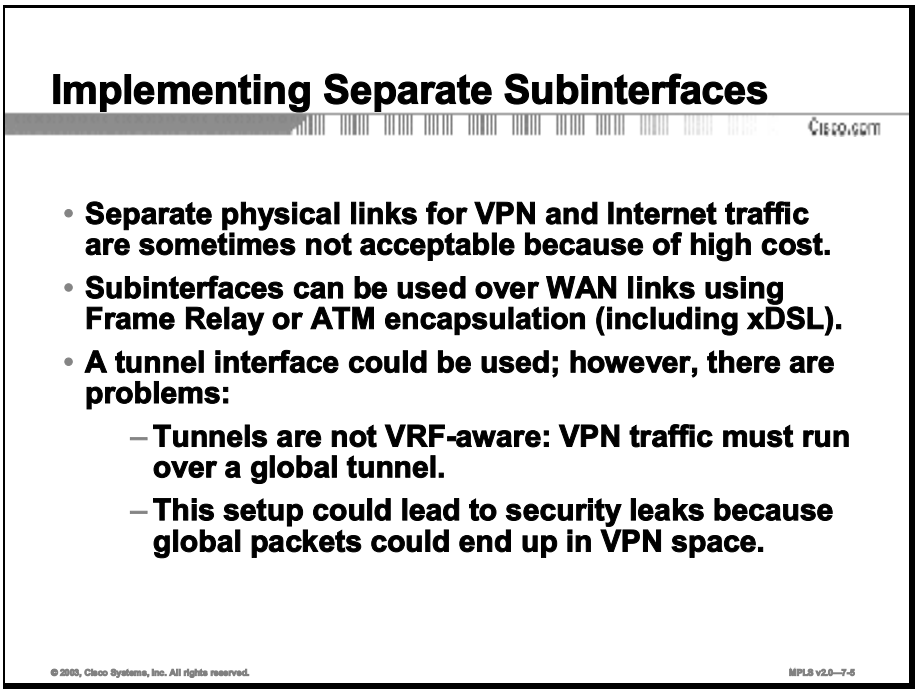

In situations where the cost factor prohibits separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can be configured only on WAN links using Frame Relay or ATM encapsulation (including xDSL) and on LAN links using any VLAN encapsulation (Inter-Switch Link Protocol [ISL] or 802.1q). For all other encapsulation types, a tunnel interface can be used between the CE and the PE router.

However, the use of tunnel interfaces is strongly discouraged for security reasons:

- A tunnel interface on a PE router is not VRF-aware. The endpoints of the tunnel have to be in a global routing table—the VPN traffic must be tunneled across an Internet interface.
- It is also very easy to spoof generic routing encapsulation (GRE) tunnels (if the tunnel key is configured, and the key is known). An intruder from the Internet could easily generate traffic that could appear as if it were coming over the GRE tunnel from the CE router and would therefore be forwarded into the customer VPN.

Example of Internet Access Through a Dedicated Subinterface

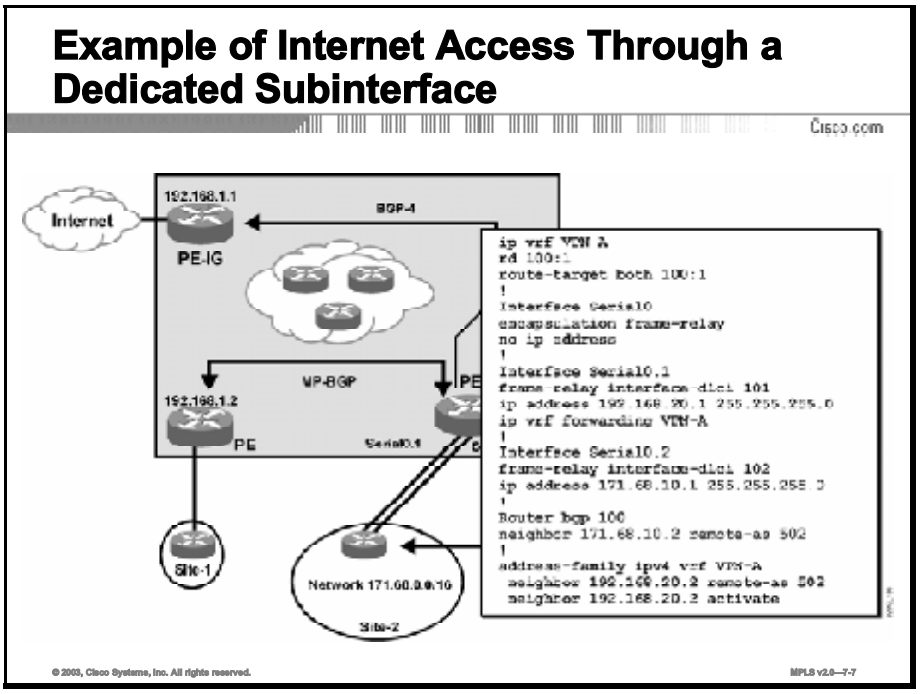

The example here illustrates the configuration needed to implement Internet access through a dedicated Frame Relay interface. The following configuration steps are performed:

- The customer virtual routing and forwarding instance (VRF) (VPN-A) is created.
- Frame Relay encapsulation is configured on the PE-CE link (Serial0).
- The VPN subinterface (Serial0.1) is created and associated with data-link connection identifier (DLCI) 101.
- The Internet subinterface (Serial0.2) is created and associated with DLCI 102.
- The CE router is configured as a BGP neighbor in both the global BGP process and inside the VPN in the VRF VPN-A.

Note The allowas-in feature needs to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

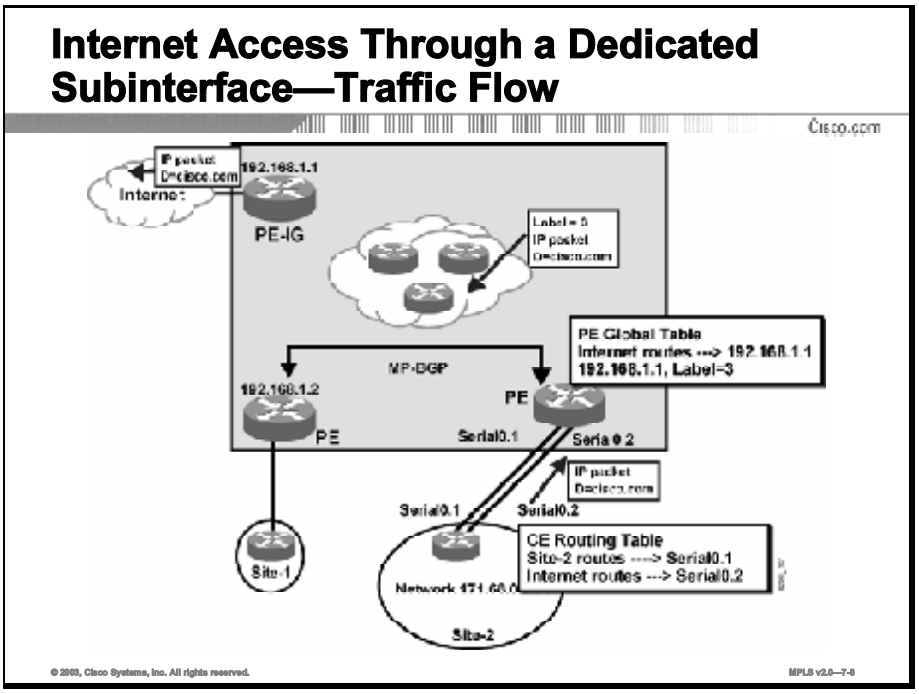

The Internet traffic flow in this setup is identical to the traditional Internet traffic flow—when a packet is received from the CE router through the Internet subinterface, a lookup is performed in the global Forwarding Information Base (FIB) on the PE router and the packet is forwarded toward the BGP next hop.

Classical Internet Access for a VPN Customer

This topic describes how classical Internet access is achieved by a VPN customer.

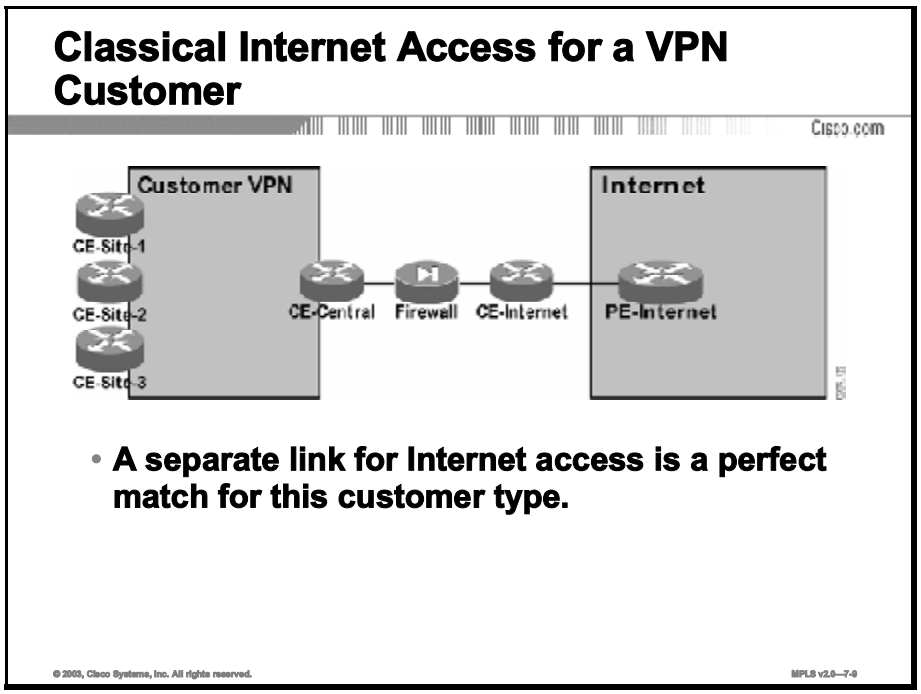

The classical Internet access setup for a VPN customer is based on a separated Internet access design model. This design model is thus a perfect match for customers looking for classical Internet access service.

Internet Access from Every Customer Site

This topic describes how Internet access is obtained from every customer site.

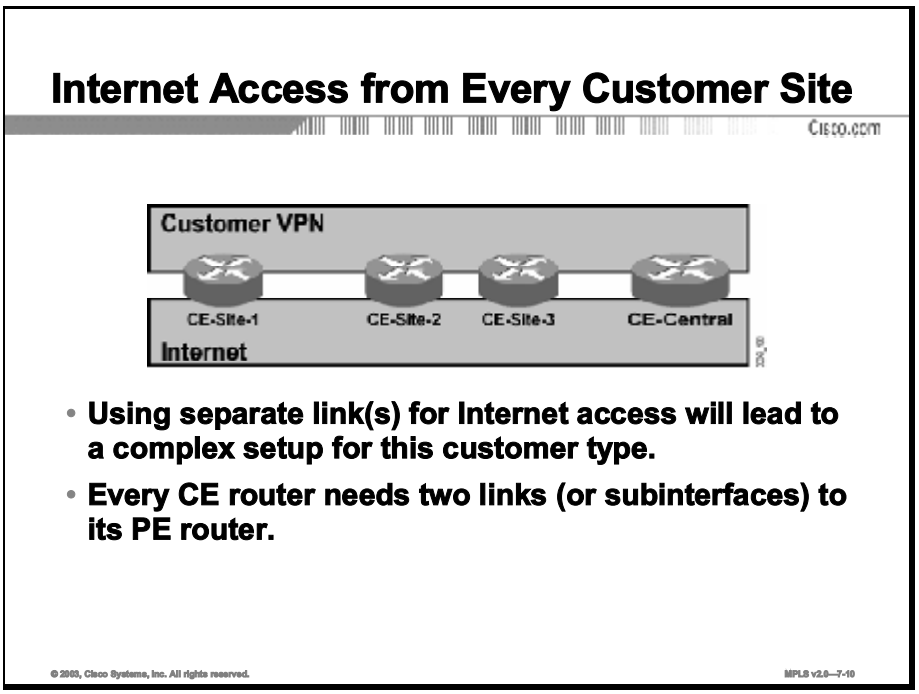

For customers that need Internet access from every site, two physical (or logical) links between every CE router and its PE router might prove to be too complex or too expensive.

Limitations of Separate Internet Access

This topic describes some of the limitations of having separate Internet access.

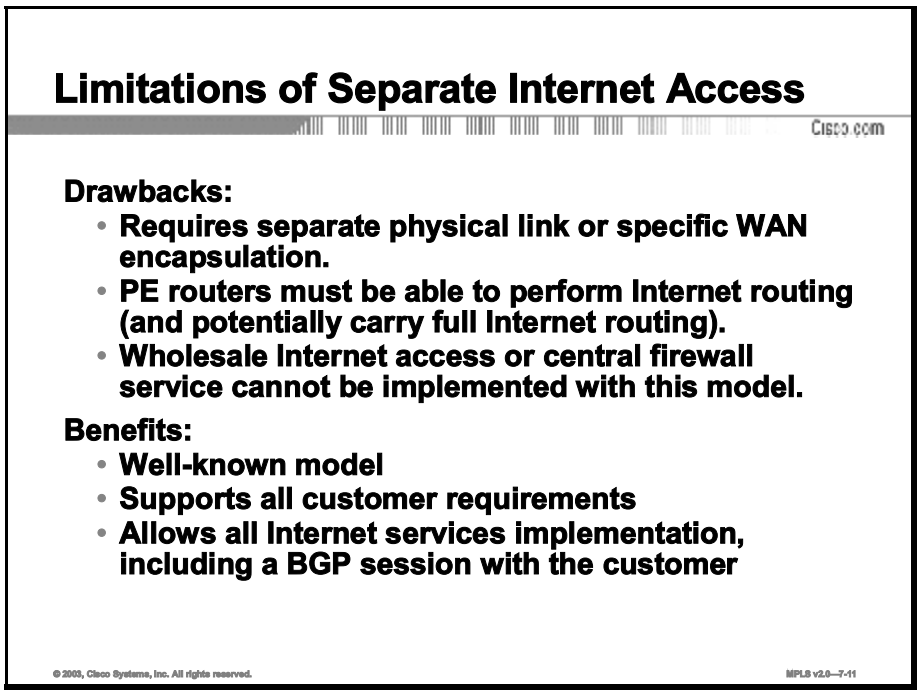

The benefits of a separate Internet access design model are obvious:

- It is a well-known and widely understood model.
- It supports all customer requirements, including multihomed customer connectivity with full Internet routing.
- In addition it allows all Internet services implementation, including BGP sessions with customers.

The drawbacks of this model are as follows:

- It requires two dedicated physical links between the PE and the CE router or specific WAN or LAN encapsulations that might not be suitable for all customers.
- The PE routers must be able to perform hop-by-hop Internet routing and use either the default route to reach the Internet or carry the full Internet routing table.
- Advanced Internet access services (central managed firewall service or wholesale Internet access service) cannot be realized with this model at all.

Summary

This topic summarizes the key points discussed in this lesson.

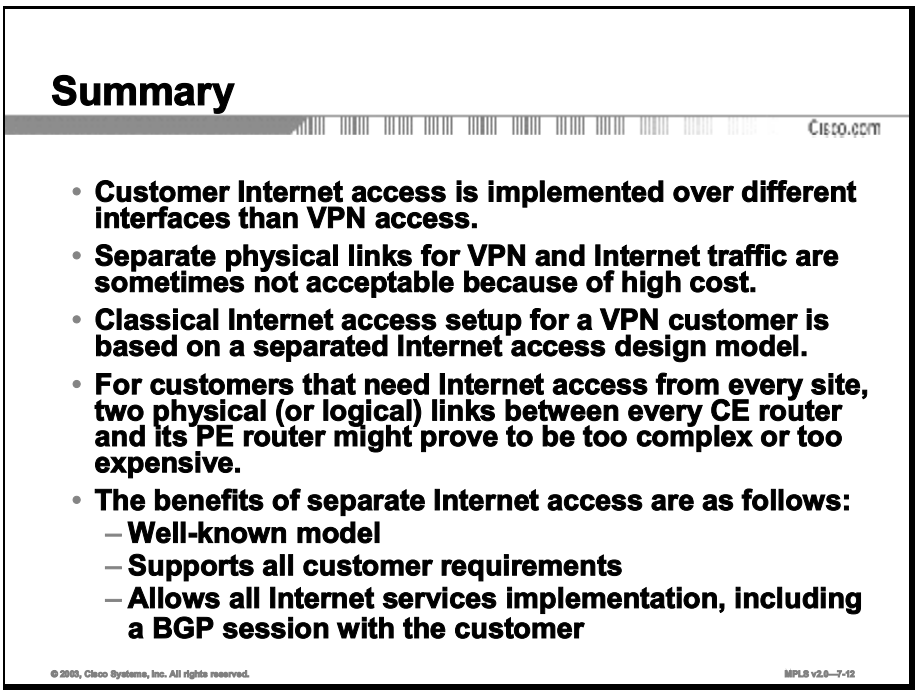

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about this topic.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The traditional Internet access implementation model provides _____ _____ flexibility because the Internet access is completely separated from the MPLS VPN services.

Q2) In situations where the cost factor prohibits separate physical links for VPN and Internet traffic, _____ can be used to create two logical links over a single physical link.

Q3) The _____ _____ _____ setup for a VPN customer is based on a separated Internet access design model.

Q4) For customers that need Internet access from every site, two physical (or logical) links between every CE router and its PE router might prove to be too _____ or too _____ to implement.

Q5) One of the drawbacks of a separate Internet access design model is that PE routers must be able to perform hop-by-hop Internet routing and use _____ _____ to reach the Internet or carry the full Internet routing table.

Quiz Answer Key

Q1) maximum design

Relates to: Designing Internet Access Separated from VPNs

Q2) subinterfaces

Relates to: Implementing Separate Subinterfaces

Q3) classical Internet access

Relates to: Classical Internet Access for a VPN Customer

Q4) complex, expensive

Relates to: Internet Access from Every Customer Site

Q5) default route

Relates to: Limitations of Separate Internet Access

Internet Access as a Separate VPN

Overview

This lesson describes the characteristics of Internet access solutions in which the Internet access is provided as a separate VPN. It identifies the scaling issues of this design and discusses how to implement the design in an MPLS VPN network.

Relevance

This lesson is a crucial one for learners planning to improve their usage of network resources using MPLS VPNs.

Objectives

This lesson describes the characteristics of Internet access solutions in which the Internet access is provided as a separate VPN.

Upon completing this lesson, you will be able to:

- Describe the characteristics of Internet access as a separate VPN
- Describe the characteristics of a redundant Internet access implementation
- Describe the characteristics of classical Internet access for a VPN customer
- Describe the characteristics of Internet access from every customer site
- Describe the characteristics of Internet access through a central firewall service
- Describe the characteristics of wholesale Internet access
- Identify the limitations of running an Internet backbone in a VPN

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Complex MPLS VPNs” module of this course

Outline

This lesson includes these topics:

- Overview
- Internet Access as a Separate VPN
- Redundant Internet Access
- Classical Internet Access for a VPN Customer
- Internet Access from Every Customer Site
- Internet Access Through a Central Firewall Service
- Wholesale Internet Access
- Limitations of Running an Internet Backbone in a VPN
- Summary
- Quiz

Internet Access as a Separate VPN

This topic considers how to use the Internet as a separate VPN.

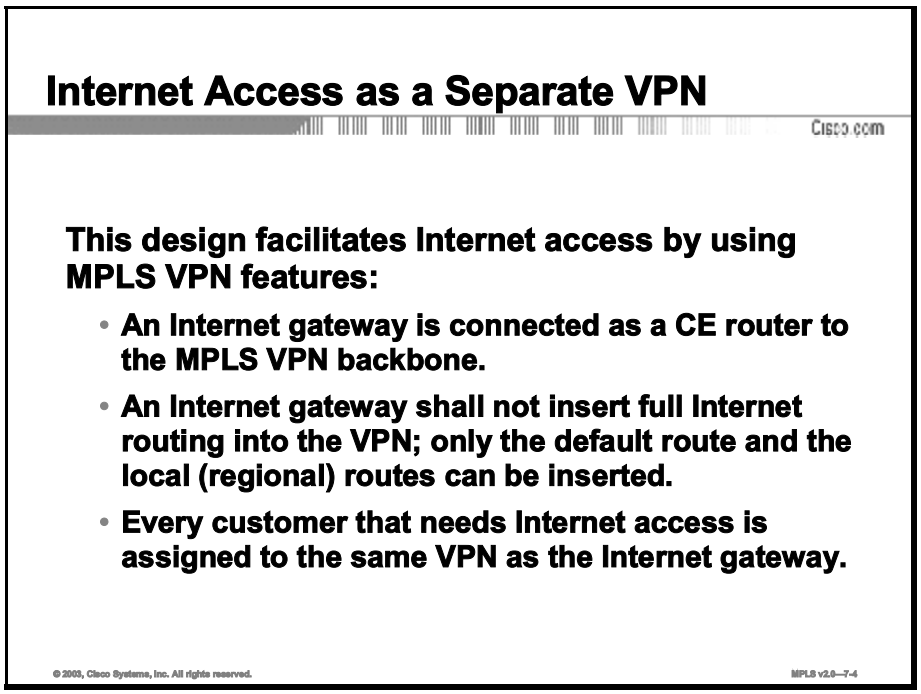

MPLS VPN architecture suggests an obvious solution to Internet access for VPN customers: define the Internet as another VPN and use various MPLS VPN topologies to implement various types of Internet access. Under this design model, the Internet gateways appear as CE routers to the MPLS VPN backbone and customer Internet access is enabled by combining an Internet VPN with a customer VPN in the VRFs of the customer (overlapping VPN topology).

The Internet VPN should not contain the full set of Internet routes, because that would make the solution completely nonscalable. The Internet gateway routers (CE routers) should announce a default route toward the PE routers. To optimize local routing, the local (or regional) Internet routes should also be inserted in the Internet VPN.

Internet Access as a Separate VPN (Cont.)

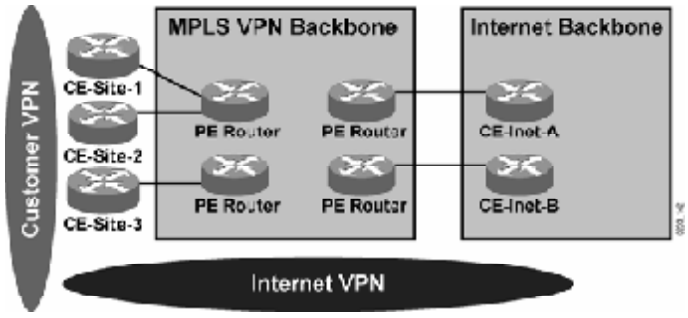

- **The Internet backbone is separate from the VPN backbone.**
- **VPN customers are connected to the Internet through a proper VPN VRF setup.**

When you are implementing Internet access as a separate VPN, the Internet backbone is kept separate from the MPLS VPN backbone, resulting in increased security for the MPLS VPN backbone (for example, Internet hosts can reach only PE routers, but not the provider routers [P routers]). The VPN customers are connected to the Internet simply through proper VRF setup.

Redundant Internet Access

This topic describes how to implement redundant Internet access.

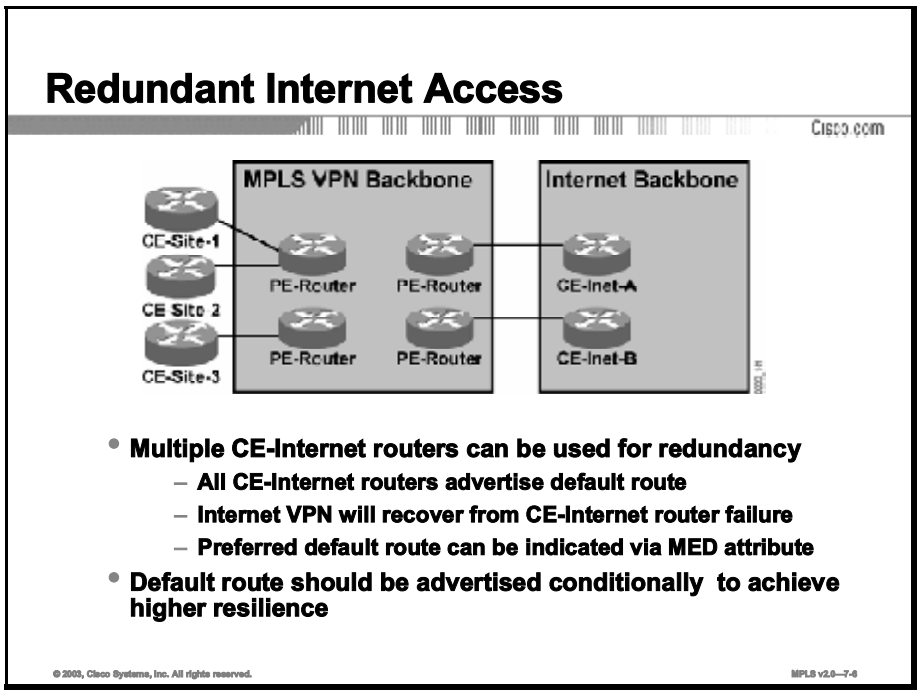

Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone:

- Multiple Internet gateways (acting as CE routers) have to be connected to the MPLS VPN backbone to ensure router and link redundancy.
- All Internet gateways advertise the default route to the PE routers, resulting in routing redundancy.
- The Internet gateways also announce local Internet routes. Because these routes are announced with different BGP attributes (most notably multi-exit discriminator [MED]), the PE routers select the proper customer edge-Internet (CE-Internet) router as the exit point toward those destinations.
- The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one CE-Internet router acts as a primary Internet gateway and the other CE-Internet routers act as a backup.
- The redundancy established so far covers the path between customer sites and the CE-Internet routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the CE-Internet routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the CE-Internet routers—they announce the default route to the PE routers only if they can reach an upstream destination.

Redundant Internet Access (Cont.)

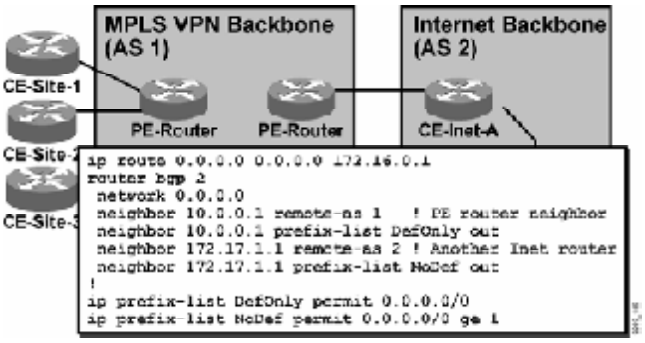

Example: CE-Inet-A should advertise default route only if it can reach network 172.16.0.0/16 (upstream ISP core).

The figure shows a sample configuration of a CE-Internet router with conditional default route advertisement. Router CE-Inet-A will advertise the default route to the PE router only if it can reach the network 172.16.0.0/16.

The following steps are used to configure this functionality:

Step 1 A static default route is configured toward a next hop in network 172.16.0.0. If the network 172.16.0.0 is not reachable, this static route will not enter the IP routing table.

Step 2 The default route origination is configured in the BGP routing process with the network command. The default route will be originated in BGP only if it is present in the IP routing table (which, based on the previous step, means that the network 172.16.0.0/16 is reachable).

Step 3 Prefix lists are used to filter BGP routing updates—the default route is sent only to the PE routers, not to the other Internet routers.

Classical Internet Access for a VPN Customer

This topic describes the characteristics of classical Internet access for a VPN customer.

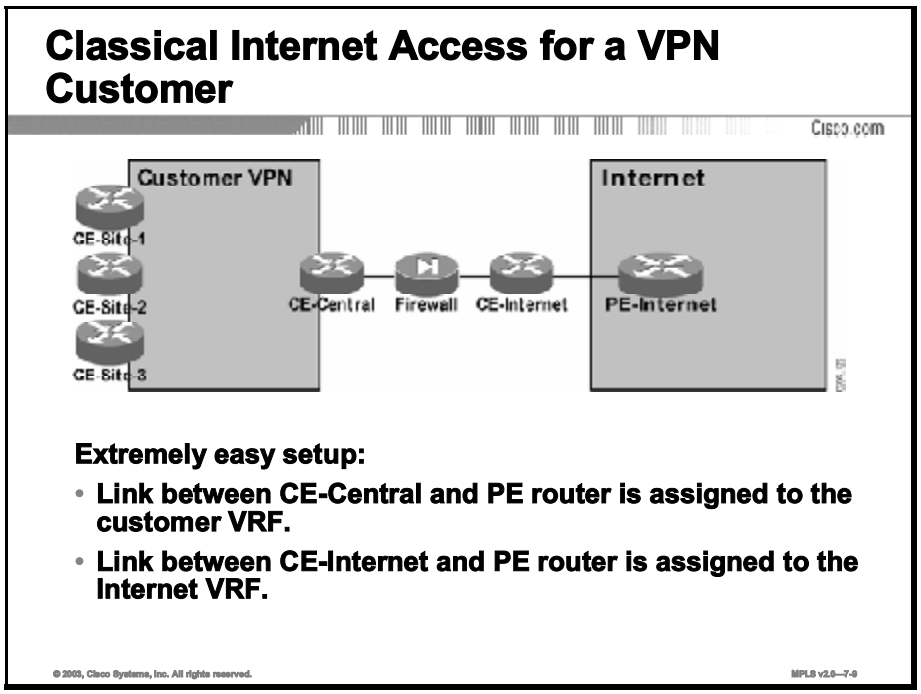

The classical Internet access model can be easily implemented with the Internet configured as a VPN over the MPLS VPN backbone. The link between a PE router and the CE-Internet router is assigned to the Internet VRF, and the link between a PE router and the CE-Central router is assigned to the customer VRF. The external Border Gateway Protocol (EBGP) multihop session can be configured between the Internet gateway (CE-Internet router in the previous figure) and the CE-Internet router in this figure to give full Internet routing to the customer.

Internet Access from Every Customer Site

This topic describes the characteristics of having Internet access from every customer site.

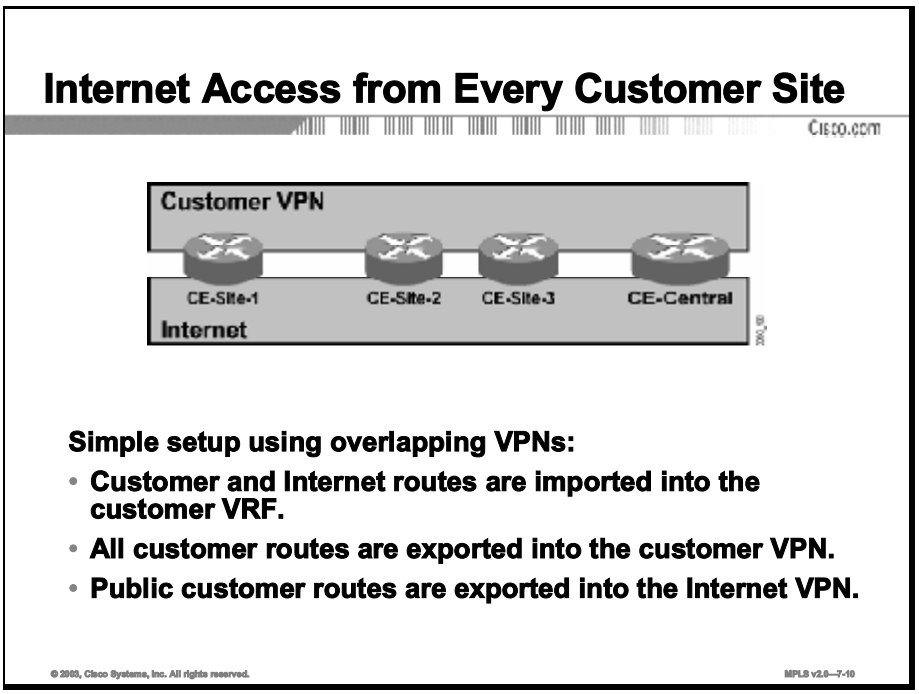

Internet access from every customer site is best implemented with an overlapping VPN solution:

- Customer routes are marked with a customer-specific (customer) route target (RT).
- Internet routes are marked with a special (Internet) RT.
- Customer sites that need to reach the Internet are placed in a separate VRF. Customer and Internet routes are imported into this VRF, and the routes exported from this VRF are marked with customer and Internet RTs.

Internet Access Through a Central Firewall Service

This topic describes the characteristics of Internet access through a central firewall service.

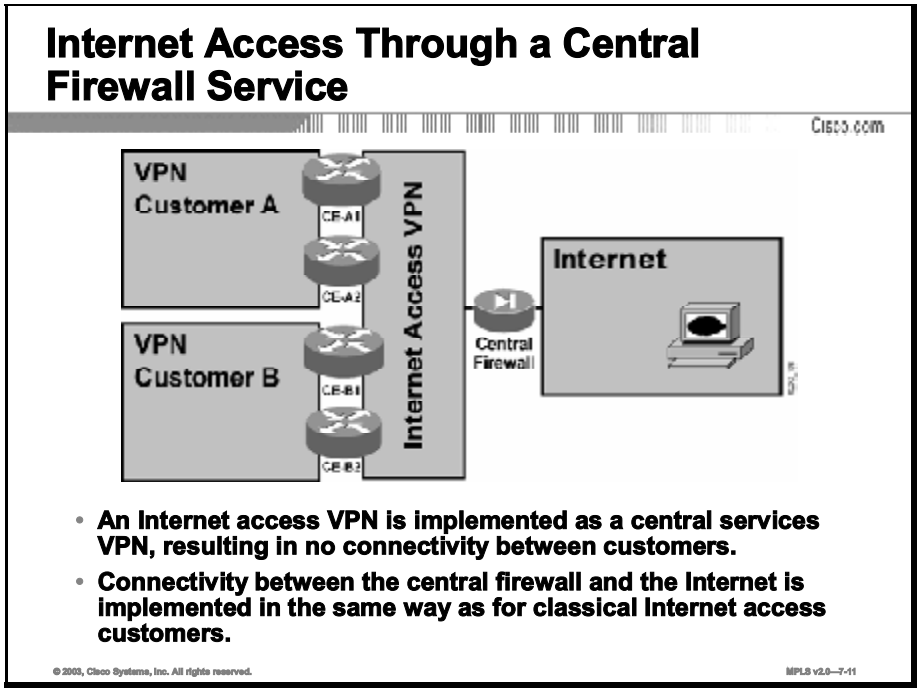

The central managed firewall service should be implemented with the central services VPN topology, with the central firewall being the server site and all customer CE routers residing in client sites. For customers with their own VPNs implemented over the same MPLS VPN backbone, the topology that overlaps customer VPN and central services VPN should be used.

The central services VPN prevents direct exchange of traffic between client sites, resulting in effective security for the customers of this service.

Connectivity between the central firewall and the Internet is implemented in the same way as Internet access for classical Internet customers. If the Internet is configured in a VPN, the public interface of the firewall is connected to an interface on a PE router, which is placed in the Internet VRF.

Wholesale Internet Access

This topic describes the characteristics of implementing the Wholesale Internet Access model.

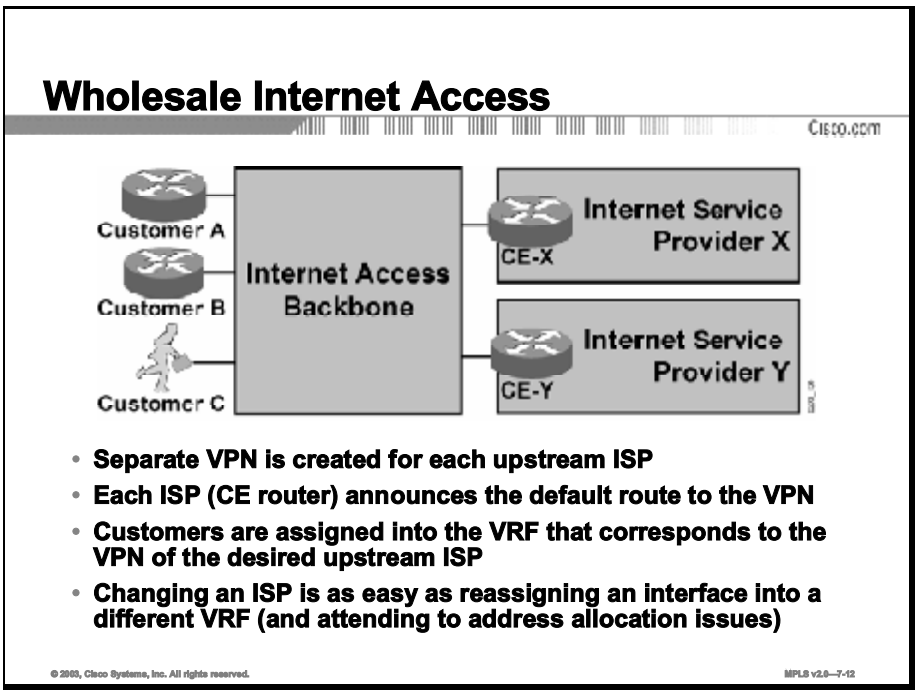

Wholesale Internet access is implemented by creating a separate VPN for every upstream ISP. The Internet gateway of the upstream ISP (acting as a CE router toward the MPLS VPN-based Internet access backbone) announces a default route, which is used for routing inside the VPN.

Customers are tied to upstream service providers simply by placing the PE-CE link into the VRF associated with the upstream service provider. Changing an ISP becomes as easy as reassigning the interface into a different VRF and attending to address allocation issues. For customers using access methods supporting dynamic address allocation (for example, dialup or cable), the new customer IP address is assigned automatically from the address space of the new ISP.

Limitations of Running an Internet Backbone in a VPN

This topic describes the benefits and limitations of running an Internet backbone in a VPN.

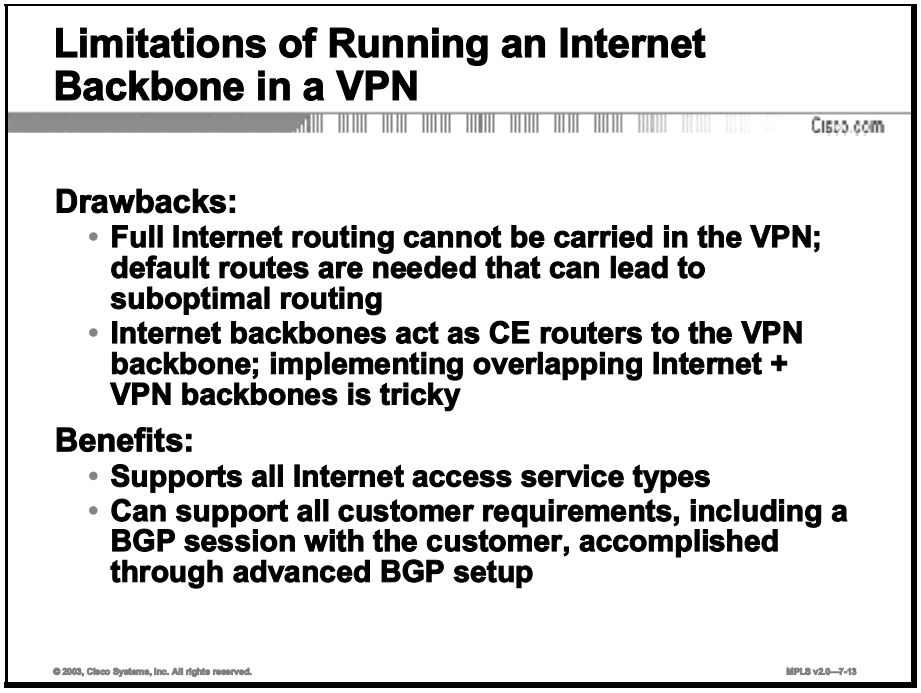

Internet access implemented as a separate VPN has a few drawbacks:

- Full Internet routing cannot be carried inside a VPN, and therefore default routing toward the Internet gateways has to be used, potentially resulting in suboptimal routing.

Note	With future MPLS VPN extensions, called "recursive VPN," or "Carrier's Carrier model," even full Internet routing will be able to be propagated across a VPN.

- The Internet backbone is positioned as a customer toward the MPLS VPN backbone. If the service provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

The benefits of this design far outweigh the limitations:

- This design model supports all Internet access services, ranging from traditional Internet access to innovative services like wholesale Internet access.
- It also supports all customer requirements, including full Internet routing on the customer routes.

Summary

This topic summarizes the key points discussed in this lesson.

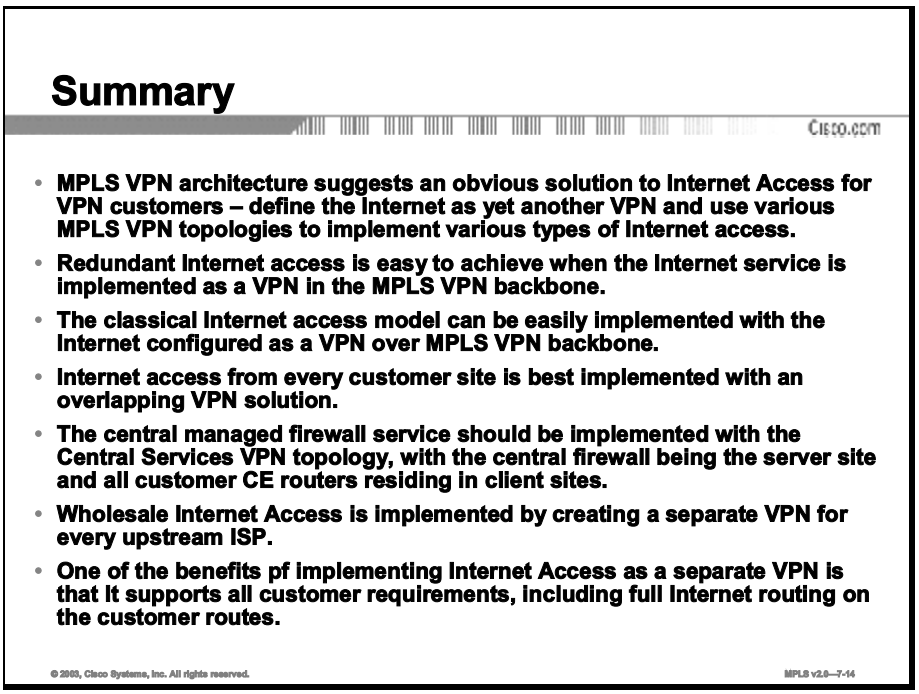

References

For additional information, refer to this resource:

- Access www.cisco.com for additional information about this topic.

Next Steps

For the associated lab exercise, refer to these sections of the course Lab Guide:

- Lab Exercise 7-1: Separate Interface for Internet Connectivity
- Lab Exercise 7-2: Multisite Internet Access
- Lab Exercise 7-3: Internet Connectivity in an MPLS VPN

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The Internet VPN should not contain the full set of _____ routes because that would make the solution completely nonscalable.

Q2) All Internet gateways (CE routers) advertise the _____ route to the PE routers, which results in routing redundancy.

Q3) In a classical Internet access for a Customer VPN model, the link between a PE router and the CE-Internet router is assigned to the _____ VRF, and the link between a PE router and the CE-Central router is assigned to the _____ VRF.

Q4) The main benefits of having Internet access from every customer site is best implemented with an _____ VPN solution.

Q5) The central managed firewall service should be implemented with the _____ _____ VPN topology.

Q6) Wholesale Internet access is implemented by creating a separate VPN for every _____ _____.

Q7) One of the drawbacks of Internet access that is implemented as a separate VPN is that _____ _____ routing cannot be carried inside a VPN. Therefore, default routing toward the Internet gateways has to be used, which can potentially result in suboptimal routing.

Quiz Answer Key

Q1) Internet

Relates to: Internet Access as a Separate VPN

Q2) default

Relates to: Redundant Internet Access

Q3) Internet, customer

Relates to: Classical Internet Access for a VPN Customer

Q4) overlapping

Relates to: Internet Access from Every Customer Site

Q5) central services

Relates to: Internet Access Through a Central Firewall Service

Q6) upstream ISP

Relates to: Wholesale Internet Access

Q7) Full Internet

Relates to: Limitations of Running an Internet Backbone in a VPN

Course Glossary

The Course Glossary for *Implementing Cisco MPLS* (MPLS) v2.0 highlights and defines key terms and acronyms used throughout this course. Many of these terms are also described in the Cisco Internetworking Terms and Acronyms resource, available via http://www.cisco.com.

Acronym or Term	Definition
access control list	*See* ACL.
ACL	access control list. A filter list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router).
AF	Assured Forwarding. There are four AF classes, AF1x through AF4x. Within each class, there are three drop probabilities. Depending on the policy of a given network, packets can be selected for a PHB based on required throughput, delay, jitter, or loss, or according to priority of access to network services.
Any Transport over MPLS	*See* AtoM.
Assured Forwarding	*See* AF.
Asynchronous Transfer Mode	*See* ATM.
ATM	Asynchronous Transfer Mode. The international standard for cell relay, in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3.
ATM edge LSR	A label switch router that is connected to the ATM-LSR cloud through LSC-ATM interfaces. The ATM edge LSR adds labels to unlabeled packets and strips labels from labeled packets. *See also* LSR.
ATM-LSR	A label switch router with a number of LSC-ATM interfaces. The router forwards the cells among these interfaces using labels carried in the VPI/VCI field. *See also* LSR.
AToM	Any Transport Over MPLS. Allows service providers who offer Layer 2 connectivity to expand service offerings by connecting Ethernet, ATM, Frame Relay, and serial and PPP networks through an MPLS backbone. AToM is a scalable architecture based on label switching that allows multiplexing of connections. It is also a standards-based open architecture and can be extended to other transport types.
BA	Behavior Aggregate. A collection of packets with the same codepoint crossing a link in a particular direction. The terms "aggregate" and "behavior aggregate" are used interchangeably.
Behavior Aggregate	*See* BA.
BGP	Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.
Border Gateway Protocol	*See* BGP.
CAR	committed access rate. A traffic policing and marking mechanism. The CAR and DCAR (distributed CAR) services limit the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.
CBWFQ	class-based weighted fair queuing. CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. Allows the user to define traffic classes based on customer-defined match criteria such as access control lists (ACLs), input interfaces, protocol, and quality of service (QoS) label. When traffic classes have been defined, they can be assigned a bandwidth, queue limit, or drop policy such as weighted random early detection (WRED).
CE	customer edge. A CE is part of a customer network and connects to a provider edge router (PE router). A CE can join any set of virtual private networks (VPNs). Each CE connects a customer site to a PE, obtaining the VPN service for that customer site, and belongs to exactly one customer. CE routers are not aware of associated VPNs. Each CE may have many configlets and may be configured by multiple SRVC service requests. *See also* PE.
cell-loss priority	*See* CLP.

Acronym or Term	Definition
CE router	customer edge router. A router that is part of a customer network and that interfaces to a provider edge router (PE router). *See also* PE router.
CEF	Cisco Express Forwarding. An advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns. Virtual routing and forwarding tables (VRFs) use CEF technology; therefore, MPLS VPNs must be enabled with CEF.
CIR	committed information rate. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. CIR, measured in bits per second, is one of the key negotiated tariff metrics.
Cisco Express Forwarding	*See* CEF.
Cisco IOS software	Cisco system software that provides common functionality, scalability, and security for all products under the CiscoFusion architecture. Cisco IOS software allows centralized, integrated, and automated installation and management of internetworks while ensuring support for a wide variety of protocols, media, services, and platforms.
class-based weighted fair queuing	*See* CBWFQ.
classification	Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, by using the three precedence bits in the type of service (ToS) field of the IP packet header—two of the values are reserved for other purposes—packets can be categorized into a limited set of up to six traffic classes. After packet classification, other QoS features can be used to assign the appropriate traffic-handling policies, including congestion management, bandwidth allocation, and delay bounds for each traffic class.
class of service	*See* CoS.
CLI	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.
CLP	cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions to free up resources to handle insured traffic.
command-line interface	*See* CLI.
committed access rate	*See* CAR.
committed information rate	*See* CIR.
congestion	Traffic in excess of network capacity.
congestion avoidance	Mechanism by which an ATM network controls the traffic entering the network to minimize delays. To use resources most efficiently, lower-priority traffic is discarded at the edge of the network if conditions indicate that it cannot be delivered.
Constraint Route-LDP	*See* CR-LDP.
CoS	class of service. CoS refers to the methods that provide "differentiated service," in which the network delivers a particular kind of service based on the class of service specified for each packet. CoS provides specific categories of service such as Gold, Silver, and Best-Effort service classes. CoS is a set of concrete device features in which a single network router treats traffic in different classes differently. CoS techniques provide a means of specifying policies to control network resource allocation in support of customer and applications requirements. The implementation of CoS techniques delivers measurable quality of service (QoS). *See also* QoS.

Acronym or Term	Definition
CR-LDP	Constraint Route-LDP. *See also* LDP.
customer edge	*See* CE.
customer edge router	*See* CE router.
data-link connection identifier	*See* DLCI.
differentiated services	*See* DiffServ.
differentiated services code point	*See* DSCP.
DiffServ	differentiated services. A paradigm for providing QoS on the Internet by employing a small, well-defined set of building blocks from which a variety of services can be built.
DLCI	data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices).
DSCP	differentiated services code point . In the IP header, this octet classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a per-hop behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.
EBGP	Exterior Border Gateway Protocol. EBGP communicates among different network domains or autonomous systems. The primary function of EBGP is to exchange network reachability information between autonomous systems, including information about the list of autonomous system routes. The autonomous systems use EGBP border edge routers to distribute the routes, which include label-switching information. Each border edge router rewrites the next-hop and MPLS labels.
EF	Expedited Forwarding. The EF PHB can be used to build a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service through DiffServ (DS) domains. Such a service appears to the endpoints like a point-to-point connection or a "virtual leased line." This service has also been described as a premium service. Codepoint 101110 is recommended for the EF PHB.
Expedited Forwarding	*See* EF.
Exterior Border Gateway Protocol	*See* EBGP.
EXP	MPLS EXP bit. MPLS is a switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on pre-established IP routing information.
FEC	forwarding equivalence class. An FEC is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. An FEC might correspond to a destination IP subnet, but it also might correspond to any traffic class that the edge LSR considers significant. For example, all traffic with a certain value of IP precedence might constitute an FEC.
FIB	Forwarding Information Base. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table. Because there is a one-to-one correlation between FIB entries and routing table entries, the FIB contains all known routes and eliminates the need for route cache maintenance that is associated with switching paths such as fast switching and optimum switching.

Acronym or Term	Definition
forwarding equivalence class	*See* FEC.
Forwarding Information Base	*See* FIB.
Frame Relay	Industry-standard, switched data link layer protocol that handles multiple virtual circuits using HDLC encapsulation between connected devices. Frame Relay is more efficient than X.25, the protocol for which it is generally considered a replacement.
Frame Relay traffic shaping	*See* FRTS.
FRF.11	Frame Relay Forum implementation agreement for Voice over Frame Relay (v1.0 May 1997). This specification defines multiplexed data, voice, fax, DTMF digit-relay, and CAS and robbed-bit signaling frame formats, but does not include call setup, routing, or administration facilities. See www.frforum.com.
FRF.12	Frame Relay Forum implementation agreement for Frame Relay Fragmentation. The FRF.12 Implementation Agreement (also known as FRF.11 Annex C) was developed to allow long data frames to be fragmented into smaller pieces and interleaved with real-time frames. In this way, real-time voice and nonreal-time data frames can be carried together on lower speed links without causing excessive delay to the real-time traffic. See www.frforum.com.
FRTS	Frame Relay traffic shaping. Queuing method that uses queues on a Frame Relay network to limit surges that can cause congestion. Data is buffered and sent into the network in regulated amounts to ensure that the traffic can fit within the promised traffic envelope for the particular connection.
generic routing encapsulation	*See* GRE.
Generic Traffic Shaping	*See* GTS.
GRE	generic routing encapsulation. Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment.
GTS	Generic Traffic Shaping. Provides a mechanism to control the traffic flow on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate (also known as the token bucket approach), while queuing bursts of the specified traffic. Thus, traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data-rate mismatches.
IETF	Internet Engineering Task Force. Task force consisting of more than 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.
integrated services	*See* IntServ.
Internet Engineering Task Force	*See* IETF.
Internet service provider	*See* ISP.
IntServ	integrated services.
IP	Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type of service specification, fragmentation and reassembly, and security. Defined in RFC 791.
IP precedence	A 3-bit value in the type of service (TOS) byte used for assigning precedence to IP packets.

Acronym or Term	Definition
IPSec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPSec provides these security services at the IP layer. IPSec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPSec. IPSec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.
IP Security	*See* IPSec.
ISP	Internet service provider. Provider of Internet access and services through a single BGP autonomous system.
L2F Protocol	Layer 2 Forwarding Protocol. Protocol that supports the creation of secure virtual private dialup networks over the Internet.
L2TP	Layer 2 Tunneling Protocol. Protocol used for implementing VPDNs and VPNs by tunneling PPP with multivendor interoperability and acceptance. This protocol was proposed as an alternative to IPSec but is often used with IPSec for authentication. This protocol merges Microsoft PPTP and Cisco Layer 2 Forwarding (L2F) technologies. *See also* IPSec.
label	A label is a header used by an LSR to forward packets. The header format depends upon network characteristics. In router networks, the label is a separate, 32-bit header. In ATM networks, the label is placed into the virtual path identifier/virtual channel identifier (VPI/VCI) cell header. In the core, LSRs read only the label, not the packet header. One key to the scalability of MPLS is that labels have only local significance between two devices that are communicating.
label-controlled ATM	*See* LC-ATM.
Label Distribution Protocol	*See* LDP.
label forwarding information base	*See* LFIB.
label imposition	The act of putting the first label on a packet.
label information base	*See* LIB.
label switched path	*See* LSP.
label switch router	*See* LSR.
LAN	local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
Layer 2 Forwarding Protocol	*See* L2F Protocol.
Layer 2 Tunneling Protocol	*See* L2TP.
LC-ATM	label-controlled ATM. An interface on a router or switch that uses label distribution procedures to negotiate label virtual channels.
LDP	Label Distribution Protocol. Provides communication between edge and core devices. It assigns labels in edge and core devices to establish label switched paths (LSPs) in conjunction with routing protocols such as OSPF, IS-IS, Enhanced Interior Gateway Routing Protocol (EIGRP), or BGP.
LFIB	label forwarding information. base. A data structure and way of managing forwarding in which destinations and incoming labels are associated with outgoing interfaces and labels.
LIB	label information base. A database used by an LSR to store labels learned from other LSRs, as well as labels assigned by the local LSR.

Acronym or Term	Definition
local-area network	*See* LAN.
LSP	label switched path. A sequence of hops (R0...Rn) in which a packet travels from R0 to Rn through label-switching mechanisms. A label switched path can be established dynamically, based on normal routing mechanisms, or it can be established through configuration.
LSP tunnel	A configured connection between two routers in which MPLS is used to carry the packet.
LSR	label switch router. The core device that switches labeled packets according to precomputed switching tables. It can also be a switch or a router.
Management Information Base	*See* MIB.
MIB	Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as SNMP or CMIP. The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.
Modular QoS CLI	*See* MQC.
MP-BGP	Multiprotocol Border Gateway Protocol. An extension to the BGP protocol that allows VPN information to remain unique within the MPLS VPN backbone. MP-BGP also allows BGP speakers to identify routing updates that do not carry standard IPv4 prefix information. Internal Border Gateway Protocol (IBGP) refers to BGP running within a single autonomous system. Exterior Border Gateway Protocol (EBGP) refers to BGP running between autonomous systems. *See* also EBGP.
MPLS	Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on pre-established IP routing information.
MPLS EXP bit	*See* EXP.
MPLS-TE	MPLS Traffic Engineering. The techniques and processes used to cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods had been used.
MPLS Traffic Engineering	*See* MPLS-TE.
MPLS VPN	Multiprotocol Label Switching Virtual Private Network. For the MPLS VPN solution, it is a set of PEs that are connected via a common “backbone” network to supply private IP interconnectivity between two or more customer sites for a given customer. Each VPN has a set of provisioning templates and policies and can span multiple provider administrative domains (PADs). CE routing communities (CERCs) in a VPN break down complex topologies into manageable subgroups.
MQC	Modular QoS CLI. The Modular QoS CLI is a CLI structure that allows users to create traffic polices and attach these polices to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the traffic policy determine how to treat the classified traffic. *See also* CLI *and* QoS.
Multiprotocol Border Gateway Protocol	*See* MP-BGP.
Multiprotocol Label Switching	*See* MPLS.
Multiprotocol Label Switching Virtual Private Network	*See* MPLS VPN.

Acronym or Term	Definition
NBAR	network-based application recognition. A new classification engine that can recognize a wide variety of application-level protocols, including HTTP via universal resource locator/Multipurpose Internet Mail Extensions (URL/MIME) type and protocols that use dynamic port assignments. When the traffic is classified by NBAR, appropriate quality of service (QoS) policies can be applied to the traffic classes using existing Cisco IOS QoS features.
network-based application recognition	*See* NBAR.
packet	Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms “datagram,” “frame,” “message,” and “segment” also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.
payload	Portion of a cell, frame, or packet that contains upper-layer information (data).
PBR	policy-based routing. Routing scheme that forwards packets to specific interfaces based on user-configured policies. Such policies might specify that traffic sent from a particular network should be forwarded out one interface, while all other traffic should be forwarded out another interface.
PE	provider edge. A router at the edge of a provider network that interfaces to the CE routers of a customer. All VPN processing occurs in the PE router. Each PE belongs to exactly one region of a provider administrative domain and connects to one or more customer sites. Each PE can have many VRF definitions and configlets, and each can be configured by many SRVC service requests. *See also* CE.
per-hop behavior	*See* PHB.
permanent virtual circuit	*See* PVC.
permanent virtual connection	*See* PVC.
PE router	provider edge router. A router that is part of a service provider network and that is connected to a customer edge (CE) router. The PE router function is a combination of an MLS edge label switch router (LSR) function and some additional functions to support VPNs. *See also* CE router.
PHB	per-hop behavior. RFC 2475 defines PHB as the externally observable forwarding behavior applied at a DiffServ-compliant node to a DiffServ Behavior Aggregate (BA). With the ability of the system to mark packets according to DSCP setting, collections of packets with the same DSCP setting, and sent in a particular direction, can be grouped into a BA. Packets from multiple sources or applications can belong to the same BA. In other words, a PHB refers to the packet scheduling, queuing, policing, or shaping behavior of a node on any given packet belonging to a BA, as configured by a service-level agreement (SLA) or a policy map. *See also* BA.
PIM	Protocol Independent Multicast. Multicast routing architecture that allows the addition of IP multicast routing on existing IP networks. PIM is unicast routing protocol independent and can be operated in two modes: dense and sparse.
point of presence	*See* POP.
Point-to-Point Protocol	*See* PPP.
Point-to-Point Tunneling Protocol	*See* PPTP.
policy-based routing	*See* PBR.
POP	point of presence. The access point, or device, into which the user dials to a service.

Acronym or Term	Definition
PPP	Point-to-Point Protocol. Successor to SLIP that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: LCP and NCP.
PPTP	Point-to-Point Tunneling Protocol. RFC 2637 describes the PPTP protocol.
Protocol Independent Multicast	*See* PIM.
provider edge	*See* PE.
provider edge router	*See* PE router.
PVC	permanent virtual circuit. Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and teardown in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.
quality of service	*See* QoS.
QoS	quality of service. The goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network.
random early detection	*See* RED.
RD	route distinguisher. An RD creates routing and forwarding tables and specifies the default route distinguisher for a VPN. The RD is an 8-byte value added to the beginning of the IPv4 prefixes of the customer to change them into globally unique VPN IPv4 prefixes. An RD is either ASN-relative, in which case it is composed of an autonomous system number and an arbitrary number, or it is IP-address-relative, in which case it is composed of an IP address and an arbitrary number. Each virtual routing and forwarding table (VRF) has an RD. Prefixes should use the same RD if they are associated with the same set of route targets (RTs). The community of interest association is based on the route target (RT) extended community attributes distributed with the Network Layer Reachability Information (NLRI). The RD value must be a globally unique value to avoid conflict with other prefixes.
Real-Time Transport Protocol	*See* RTP.
RED	random early detection. This class of algorithms is designed to avoid congestion in internetworks before it becomes a problem. RED works by monitoring traffic load at points in the network and stochastically discarding packets if the congestion begins to increase. The result of the drop is that the source detects the dropped traffic and slows its transmission. RED is designed to work primarily with TCP in IP internetwork environments.
Resource Reservation Protocol	*See* RSVP.
route distinguisher	*See* RD.
route target	*See* RT.
RSVP	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams that they want to receive. RSVP depends on IPv6. Also known as Resource Reservation Setup Protocol.

Acronym or Term	Definition
RT	route target. A 64-bit value by which Cisco IOS software discriminates routes for route updates in VRFs. An RT defines the destination of the route by specifying a target VPN community. RTs identify a group of sites that exchange routing information. BGP uses the RT value when determining the eligibility of routes for installation to a local routing table.
RTP	Real-Time Transport Protocol. Commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, time-stamping, and delivery monitoring to real-time applications.
service class	Collection of service types required for a specific service offered. Each service class includes the attributes and values that define the type or quality of service associated with a given class. For example, data connectivity is a service class you might define that includes the service type data-bandwidth.
service-level agreement	*See* SLA.
Site of Origin	*See* SOO.
SLA	service-level agreement. SLAs are negotiated contracts between VPN providers and their subscribers. An SLA defines the criteria for the specific services that the subscriber expects the provider to deliver. The SLA is the only binding mechanism at the disposal of the subscriber to ensure that the VPN provider delivers the services as agreed.
SOO	Site of Origin. SOO is a concept in MPLS VPN architecture that prevents routing loops in a site that is multihomed to the MPLS VPN backbone and in a site that uses AS-override. SOO is a BGP extended community attribute used to identify an IP address that originated from a site to prevent that IP address from being advertised back to the site. This attribute uniquely identifies the site from which the PE router learned the route. SOO is tagged at a PE in peering with BGP neighbors using an inbound route map and works in conjunction with the BGP CE-PE routing protocol. SOO must be unique for each customer site for each VPN. Therefore, the same value of SOO must be used on PE routers connected to the same CE router or to the same customer site.
Tag Distribution Protocol	*See* TDP.
TCB	transaction control block.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.
TDP	Tag Distribution Protocol. The protocol used to distribute label bindings to LSRs. (Cisco proprietary version of LDP)
ToS	type of Service. A byte in the IPv4 header.
traffic engineering tunnel	A label switched path tunnel that is used for engineering traffic. It is set up through means other than normal Layer 3 routing and is used to direct traffic over a path different from the one that Layer 3 routing would cause it to take.
transaction control block	*See* TCB.
Transmission Control Protocol	*See* TCP.
tunnel	Secure communication path between two peers, such as two routers.

Acronym or Term	Definition
tunneling	Architecture providing the services necessary to implement any standard point-to-point data encapsulation scheme.
type of service	*See* ToS.
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
User Datagram Protocol	*See* UDP.
VC merge	virtual circuit merge. VC merge allows the switch to transmit cells coming from different VCIs over the same outgoing VCI toward the same destination. In other words, it allows multipoint-to-point connections. VC merge is accomplished by queuing complete AAL5 frames in input buffers until the end of the frame has been received. The cells from the same AAL5 frame are all transmitted before sending cells from any other frame. This approach requires sufficient buffering capabilities inside the switch, but no more buffering than is required in IP networks.
Versatile Interface Processor	*See* VIP.
VIP	Versatile Interface Processor. Interface card used in Cisco 7000 and Cisco 7500 series routers. The VIP provides multilayer switching and runs Cisco IOS software. The most recent version of the VIP is VIP2.
virtual channel	*See* virtual circuit.
virtual circuit	Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either permanent (PVC) or switched (SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a virtual channel.
virtual circuit merge	*See* VC merge.
virtual LAN	*See* VLAN.
virtual path identifier/virtual channel identifier	*See* VPI/VCI.
Virtual Private Network	*See* VPN.
virtual routing and forwarding instance	*See* VRF.
VLAN	virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
Voice over IP	*See* VoIP.
VoIP	Voice over IP. The capability to carry normal telephony-style voice over an IP-based Internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323.

Acronym or Term	Definition
VPI/VCI	virtual path identifier/virtual channel identifier . VPI - virtual path identifier. 8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. VCI - virtual channel identifier. 16-bit field in the header of an ATM cell. The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination.
VPN	Virtual Private Network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses "tunneling" to encrypt all information at the IP level.
VPNv4	Used as a keyword in commands to indicate VPN IPv4 prefixes. These prefixes are customer VPN addresses, each of which has been made unique by the addition of an 8-byte route distinguisher.
VRF	virtual routing and forwarding instance. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.
WAN	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.
weighted fair queuing	*See* WFQ.
WFQ	weighted fair queuing. Congestion management algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission.
wide-area network	*See* WAN.

Lab Guide

Overview

Use the tasks here to complete the lab exercises for this course. The solutions information is found in the Lab Exercise Answer Key.

Outline

This Lab Guide includes these exercises:

- Lab Exercise 3-1: Establishing the Service Provider IGP Routing Environment
- Lab Exercise 3-2: Establishing the Core MPLS Environment
- Lab Exercise 5-1: Initial MPLS VPN Setup
- Lab Exercise 5-2: Running EIGRP Between PE and CE Routers
- Lab Exercise 5-3: Running OSPF Between PE and CE Routers
- Lab Exercise 5-4: Running BGP Between PE and CE Routers
- Lab Exercise 6-1: Overlapping VPNs
- Lab Exercise 6-2: Merging Service Providers
- Lab Exercise 6-3: Common Services VPN
- Lab Exercise 7-1: Separate Interface for Internet Connectivity
- Lab Exercise 7-2: Multisite Internet Access
- Lab Exercise 7-3: Internet Connectivity in an MPLS VPN

Lab Exercise 3-1: Establishing the Service Provider IGP Routing Environment

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

In this exercise, you will use the tasks and commands necessary to implement the service provider Interior Gateway Protocol (IGP) and routing environment. After completing this exercise, you will be able to meet these objectives:

- Verify the service provider IP addressing scheme, data-link connection identifier (DLCI) assignment, and interface status
- Enable the service provider IGP and configure appropriate IP addressing

Visual Objective

The figure illustrates what you will accomplish in this exercise. This exercise contains information about your laboratory setup, and details of the physical and logical connectivity in the laboratory, as well as information about the addressing scheme and IGP routing. The class will be divided into workgroups; each workgroup has its own customer edge (CE), provider edge (PE), and P routers that will provide service to two customers, customer A and customer B.

The names of all routers in your workgroup follow the naming convention detailed in this table.

Router Naming Convention

Router Role	Description
`Routers of customer A`	CE*wg*1A and CE*wg*2A, where *wg* equals your assigned workgroup number.
`Routers of customer B`	CE*wg*1B and CE*wg*2B, where *wg* equals your assigned workgroup number.
`Provider edge router`	PE*wg*1 and PE*wg*2, where *wg* equals your assigned workgroup number.
`Provider router`	P*wg*1 and P*wg*2, where *wg* equals your assigned workgroup number.

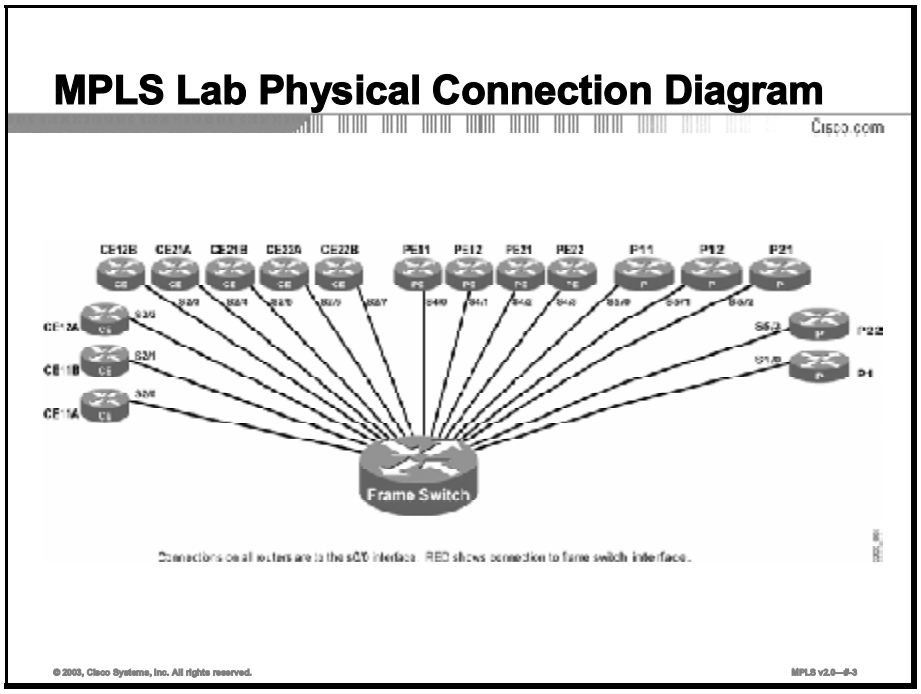

The first serial interface of each routers is connected to a Frame Relay switch. Logical connectivity has been provided by preconfigured permanent virtual circuits (PVCs), or DLCIs.

The DLCI values for all Frame Relay virtual circuits are shown in Table 2.

DLCI Identification

Source Router Type	Destination Router Type	DLCI
CE*wg*1A	PE*wg*1	101
CE*wg*1B	PE*wg*1	102
CE*wg*2A	PE*wg*2	101
CE*wg*2B	PE*wg*2	102
PE*wg*1	CE*wg*1A	101
PE*wg*1	CE*wg*1B	102
PE*wg*1	P*wg*1	111
PE*wg*2	CE*wg*2A	101
PE*wg*2	CE*wg*2B	102
PE*wg*2	P*wg*2	111
P*wg*1	PE*wg*1	111
P*wg*1	P*gw*2	112
P*wg*2	PE*wg*2	111
P*wg*2	P*gw*2	112

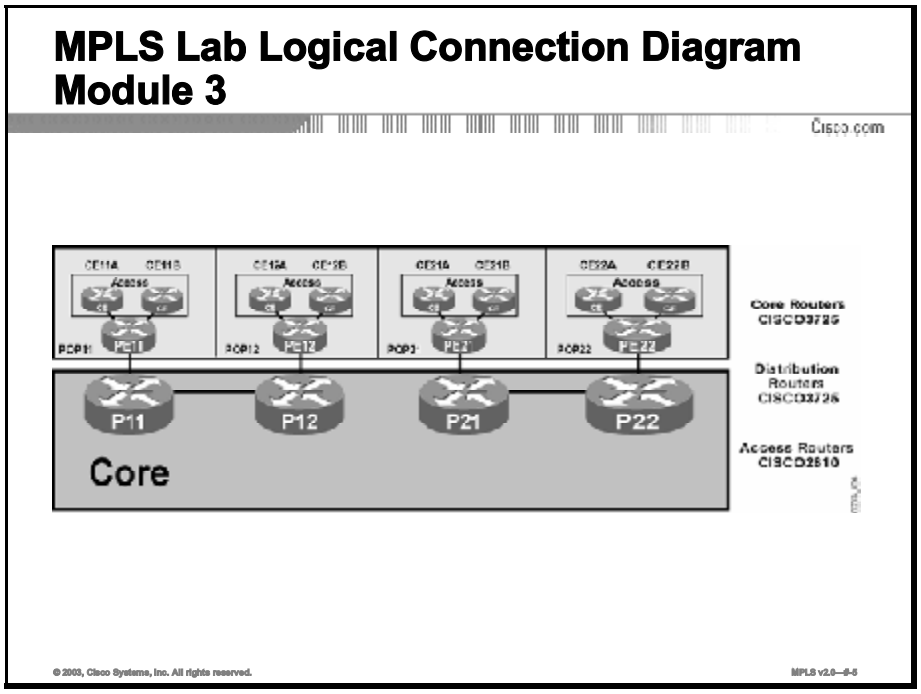

Each workgroup has two points of presence (POPs), POP*wg*1and POP*wg*2, where *wg* equals your workgroup number. For example:

- Workgroup 1 has POP11 and POP12
- Workgroup 2 has POP21 and POP22
- and so on

Each POP contains three routers. The "A" CE router provides support for customer A, and the "B" CE router provides support for customer B. Routing Information Protocol version 2 (RIPv2) is running on all CE routers.

The PE router provides connectivity to the core for both customers.

In addition to the three routers located in each POP, each workgroup also has two provider routers (P routers), one supporting each POP, in the core.

Each workgroup is further divided into two teams: team A and team B. Team A will configure the PE router on POP*wg*1 and its associated P router. Team B will configure the PE router on POP*wg*2 and its associated P router.

The routers in your workgroup have different roles as detailed in the following table.

Workgroup Router Roles

Router Name	Router Role in the Laboratory
CE*wg*1A, CE*wg*2A	Customer routers of customer A. The customer has two sites connected to different PE routers.
CE*wg*1B, CE*wg*2B	Customer routers of customer B. The customer has two sites connected to different PE routers.
PE*wg*1, PE*wg*2	PE routers—routers in your service provider backbone that connect to customer routers and to other service providers or customers.
P*wg*1	Provider core routers—routers in your service provider backbone that connect to PE routers and to other provider core routers.

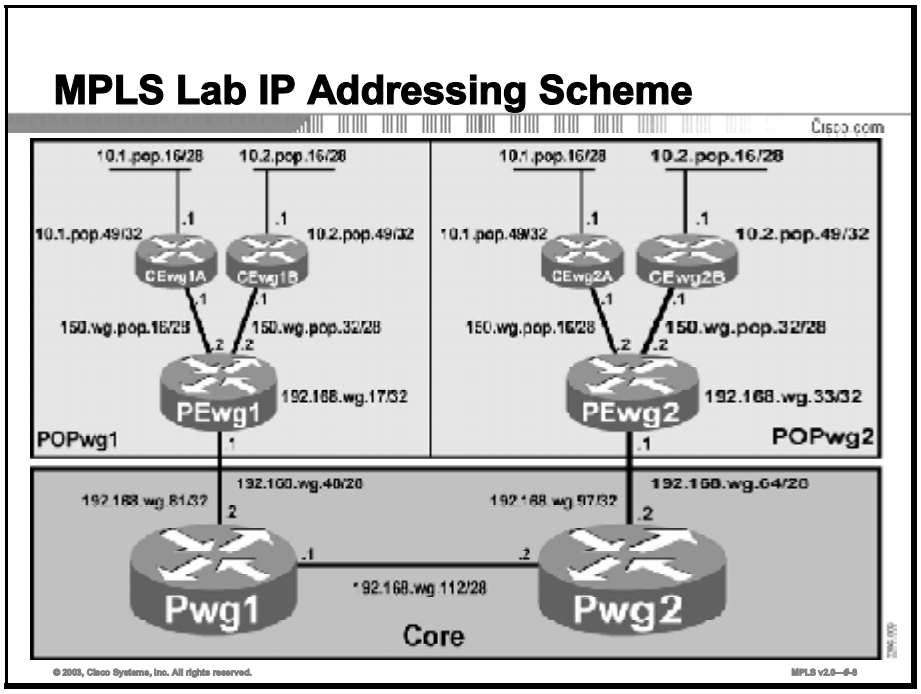

The addressing of customer routers has been performed using the IP allocations scheme detailed in the following table.

Service Provider Address Space

Parameter	Value
PEwg1 (S0/0.101)	150.wg.pop.18/28
PEwg1 (S0/0.102)	150.wg.pop.33/28
PEwg1 (loopback0)	192.168.wg.17/32
PEwg1 (S0/0.111)	192.168.wg.49/28
PEwg2 (S0/0.101)	150.wg.pop.18/28
PEwg1 (S0/0.102)	150.wg.pop.33/28
PEwg2 (loopback0)	192.168.wg.33/32
PEwg1 (S0/0.111)	192.168.wg.65/28
Pwg1 (S0/0.111)	192.168.wg.50/28
Pwg1 (S0/0.112)	192.168.wg.113/28
Pwg2 (S0/0.111)	192.168.wg.66/28
Pwg1 (S0/0.112)	192.168.wg.114/28

Note	The above addressing scheme has been selected for ease of use in the labs. It does not optimize the use of the address space.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

IP, IGP, and Interface Commands

Command	Description
network *network-number [network-mask]* **no network** *network-number [network-mask]*	To specify a list of networks for the Enhanced Interior Gateway Routing Protocol (EIGRP) routing process, use the **network** router configuration command. To remove an entry, use the **no** form of this command.
router eigrp *as-number* **no router eigrp** *as-number*	To configure the EIGRP routing process, use the **router eigrp** global configuration command. To shut down a routing process, use the **no** form of this command.
show frame-relay pvc	To display statistics about PVCs for Frame Relay interfaces, use the **show frame-relay pvc** privileged EXEC command.
show interfaces serial [***slot***/port]	To display information about a serial interface, use the **show interfaces serial** command in privileged EXEC mode. When using Frame Relay encapsulation, use the **show interfaces serial** command in EXEC mode to display information about the multicast DLCI, the DLCIs used on the interface, and the DLCI used for the Local Management Interface (LMI).
show ip protocols	To display the parameters and current state of the active routing protocol process, use the **show ip protocols** EXEC command.
show ip route [*ip-address* [*mask*] [**longer-prefixes**]] | [*protocol* [*process-id*]]	To display the current state of the routing table, use the **show ip route** EXEC command.

Task 1: Verifying the Service Provider IP Interfaces

The physical interfaces, subinterfaces, DLCI assignments, and IP addressing have been preconfigured on all PE and P routers. Your task is to verify the addressing and ensure that the proper interfaces are enabled.

Note The enable password on all routers is "mpls."

Exercise Procedure

Complete these steps:

Step 1 If you are team B, skip to Step 6.

Step 2 On PE*wg*1 verify and record the IP address scheme, DLCI assignment, and interface status of the interfaces indicated in Table 6.

Step 3 On P*wg*1 verify and record the IP address scheme, DLCI assignment, and interface status of the interfaces indicated in Table 6.

Step 4 Ensure that team B has the information that you recorded in Table 6.

Step 5 Proceed to Task 2.

Step 6 On PE*wg*2 verify and record the IP address scheme, DLCI assignment, and interface status of the interfaces indicated in Table 6.

Step 7 On P*wg*2 verify and record the IP address scheme, DLCI assignment, and interface status of the interfaces indicated in Table 6.

Step 8 Ensure that team A has the information that you recorded in Table 6.

Step 9 Proceed to Task 2.

Service Provider Interface Information

Router	Interface	Subnet Address	IP Address	DLCI	Status
PEwg1	S0/0.101				
	S0/0.102				
	S0/0.111				
	Loopback0				
PEwg2	S0/0.101				
	S0/0.102				
	S0/0.111				
	Loopback0				
Pwg1	S0/0.111				
	S0/0.112				
	Loopback0				
Pwg2	S0/0.111				
	S0/0.112				
	Loopback0				

Task 2: Configuring the Service Provider IGP

Your next task is to establish the service provider IGP routing environment. This task will involve enabling the EIGRP routing protocol.

Exercise Procedure

Complete these steps:

Step 1 If you are team B, skip to Step 5.

Step 2 On PEwg1 and Pwg1, enable the EIGRP routing process, using 1 as the autonomous system (AS) number, and ensure that the service provider networks are configured and are being advertised by the EIGRP process.

Step 3 Ensure that team B has completed its configuration tasks.

Step 4 Proceed to the Exercise Verification.

Step 5 On PEwg2 and Pwg2, enable the EIGRP routing process, using 1 as the AS number, and ensure that the service provider networks are configured and are being advertised by the EIGRP process.

Step 6 Ensure that team A has completed its configuration.

Step 7 Proceed to the Exercise Verification.

Exercise Verification

You have completed this exercise when you attain these results:

- On each P and PE router, verify that the EIGRP router process is active.
- On each P and PE router, verify that the EIGRP router process is enabled on all serial interfaces.
- On each P and PE router, verify that the loopback interfaces of all P and PE routers are displayed in the IP routing table.
- On each P and PE router, verify that 192.168.wg.0 subnets of all P and PE routers are displayed in the IP routing table.
- On each PE router, verify that 150.wg.pop.0 subnets of all P and PE routers are displayed in the IP routing table.

Lab Exercise 3-2: Establishing the Core MPLS Environment

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

In this exercise, you will use the tasks and commands necessary to implement MPLS on frame-mode Cisco IOS platforms. After completing this exercise, you will be able to meet these objectives:

- Enable LDP on your PE and P routers
- Disable MPLS TTL propagation
- Configure conditional label distribution

Visual Objective

The figure illustrates what you will accomplish in this exercise.

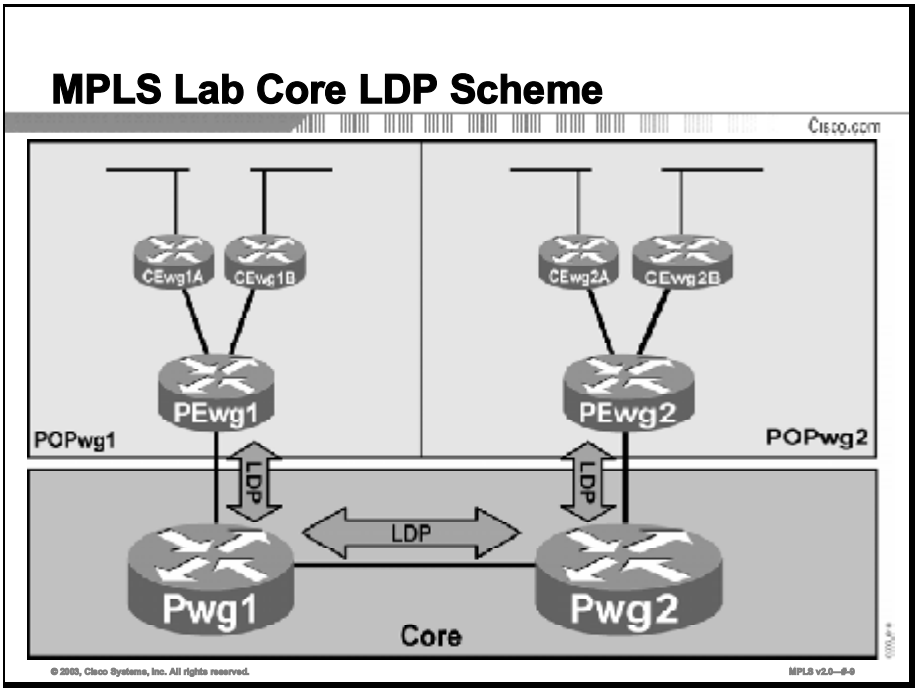

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

MPLS Commands

Command	Description
access-list *access-list-number* {**permit** | **deny**} {*type-code wild-mask* | *address mask*} **no access-list** *access-list-number* {**permit** | **deny**} {*type-code wild-mask* | *address mask*}	To configure the access list mechanism for filtering frames by protocol type or vendor code, use the **access-list** global configuration command. To remove the single specified entry from the access list, use the **no** form of this command.
ip cef	To enable Cisco Express Forwarding (CEF) on the route processor (RP) card, use the **ip cef** command in global configuration mode. To disable CEF, use the **no** form of this command.
mpls ip **no mpls ip**	To enable MPLS forwarding of IP version 4 (IPv4) packets along normally routed paths for the platform, the **mpls ip** command can be used in global configuration mode (for traffic engineering [TE]) but must be used at the interface configuration mode for Label Distribution Protocol (LDP) to become active. To disable this feature, use the **no** form of this command.
mpls ip propagate-ttl **no mpls ip propagate-ttl** [**forwarded** | **local**]	To control the generation of the time-to-live (TTL) field in the MPLS header when labels are first added to an IP packet, use the **mpls ip propagate-ttl** global configuration command. To use a fixed TTL value (255) for the first label of the IP packet, use the **no** form of this command.
mpls label protocol {**ldp** | **tdp** | **both** } [**no**] **mpls label protocol**	To specify the label distribution protocol to be used on a given interface, use the **mpls label protocol** interface configuration command. Use the **no** form of the command to disable this feature.
show mpls interfaces [*interface*] **[detail]**	To display information about one or more interfaces that have been configured for label switching, use the **show mpls interfaces** privileged EXEC command.
show mpls ldp discovery	To display the status of the LDP discovery process, use the **show mpls ldp discovery** privileged EXEC command. This command generates a list of interfaces over which the LDP discovery process is running.
show mpls ldp neighbor [*address* | *interface*] **[detail]**	To display the status of LDP sessions, issue the **show mpls ldp neighbor** privileged EXEC command.
show mpls ldp bindings **[*network* {*mask* | *length*}** [longer-prefixes]] **[**local-label ***label*** **[-*label*]}** [remote-label ***label [- label]*** **]** [neighbor ***address]*** **[**local**]**	To display the contents of the label information base (LIB), use the **show mpls ldp bindings** privileged EXEC command.

Command	Description
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]]	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.

Task 1: Enabling LDP on Your PE and P Routers

Your next task is to establish MPLS within the service provider routing environment. This task will involve enabling CEF and MPLS.

Exercise Procedure

Complete these steps:

Step 1 On your assigned PE router:

- Enable CEF.
- Enable LDP on the subinterface that is connected to your assigned P router.

Step 2 On your assigned P router:

- Enable CEF.
- Enable LDP on the subinterface that is connected to your assigned PE router.
- Enable LDP on the subinterface that is connected to the P router of the other team.

Step 3 Verify that the other team has completed its configuration.

Exercise Verification

You have completed this exercise when you attain these results:

- On each of your routers, verify that the interfaces in question have been configured to use LDP.

```
P11#sh mpls interface
Interface              IP            Tunnel   Operational
Serial0/0.111          Yes (ldp)     No       Yes
Serial0/0.112          Yes (ldp)     No       Yes
```

- On each of your routers, verify that the interface is up and has established an LDP neighbor relationship.

```
P11#show mpls ldp discovery
 Local LDP Identifier:
    192.168.1.81:0
    Discovery Sources:
    Interfaces:
        Serial0/0.111 (ldp): xmit/recv
            LDP Id: 192.168.1.17:0
        Serial0/0.112 (ldp): xmit/recv
            LDP Id: 192.168.1.97:0

P11#show mpls ldp nei
```

```
    Peer LDP Ident: 192.168.1.17:0; Local LDP Ident
192.168.1.81:0
        TCP connection: 192.168.1.17.646 - 192.168.1.81.11000
        State: Oper; Msgs sent/rcvd: 20/23; Downstream
        Up time: 00:08:03
        LDP discovery sources:
          Serial0/0.111, Src IP addr: 192.168.1.49
        Addresses bound to peer LDP Ident:
          192.168.1.17    192.168.1.49    150.wg.11.18
150.wg.11.34
    Peer LDP Ident: 192.168.1.97:0; Local LDP Ident
192.168.1.81:0
        TCP connection: 192.168.1.97.11000 - 192.168.1.81.646
        State: Oper; Msgs sent/rcvd: 18/18; Downstream
        Up time: 00:06:15
        LDP discovery sources:
          Serial0/0.112, Src IP addr: 192.168.1.114
        Addresses bound to peer LDP Ident:
          192.168.1.97    192.168.1.66    192.168.1.114
```

- On each of your routers, verify that LDP has allocated a label for each prefix in its IP routing table.

```
PE11#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     192.168.1.0/24 is variably subnetted, 8 subnets, 3 masks
D       192.168.1.97/32 [90/2809856] via 192.168.1.50,
00:49:50, Serial0/0.111
D       192.168.1.112/28
           [90/2681856] via 192.168.1.50, 00:49:50,
Serial0/0.111
```



```
D       192.168.1.64/28 [90/3193856] via 192.168.1.50,
00:49:50, Serial0/0.111
D       192.168.1.81/32 [90/659968] via 192.168.1.50,
00:49:50, Serial0/0.111
D       192.168.1.33/32 [90/3321856] via 192.168.1.50,
00:47:00, Serial0/0.111
C       192.168.1.48/28 is directly connected, Serial0/0.111
D       192.168.1.0/24 is a summary, 00:49:20, Null0
C       192.168.1.17/32 is directly connected, Loopback0
     150.wg.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       150.wg.11.16/28 is directly connected, Serial0/0.101
D       150.wg.0.0/16 is a summary, 00:49:20, Null0
C       150.wg.11.32/28 is directly connected, Serial0/0.102

P11#sh mpls ldp bindings
  tib entry: 150.wg.0.0/16, rev 16
        local binding:  tag: 20
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 20
  tib entry: 150.wg.11.16/28, rev 18
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.wg.11.32/28, rev 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.0/24, rev 17
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.17/32, rev 14
        local binding:  tag: 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 19
  tib entry: 192.168.1.33/32, rev 10
        local binding:  tag: 18
        remote binding: tsr: 192.168.1.17:0, tag: 20
        remote binding: tsr: 192.168.1.97:0, tag: 17
  tib entry: 192.168.1.48/28, rev 12
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 18
  tib entry: 192.168.1.64/28, rev 6
        local binding:  tag: 17
        remote binding: tsr: 192.168.1.17:0, tag: 18
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
```

```
  tib entry: 192.168.1.81/32, rev 8
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: 19
        remote binding: tsr: 192.168.1.97:0, tag: 16
  tib entry: 192.168.1.97/32, rev 2
        local binding:  tag: 16
        remote binding: tsr: 192.168.1.17:0, tag: 16
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
  tib entry: 192.168.1.112/28, rev 4
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: 17
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
```

- On each of your routers, verify that LDP has received a label of the subnets and loopback interfaces of the other core routers.

```
P11#sh mpls ldp bindings
  tib entry: 150.wg.0.0/16, rev 16
        local binding:  tag: 20
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 20
  tib entry: 150.wg.11.16/28, rev 18
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.wg.11.32/28, rev 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.0/24, rev 17
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.17/32, rev 14
        local binding:  tag: 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 19
  tib entry: 192.168.1.33/32, rev 10
        local binding:  tag: 18
        remote binding: tsr: 192.168.1.17:0, tag: 20
        remote binding: tsr: 192.168.1.97:0, tag: 17
  tib entry: 192.168.1.48/28, rev 12
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 18
  tib entry: 192.168.1.64/28, rev 6
        local binding:  tag: 17
```



```
        remote binding: tsr: 192.168.1.17:0, tag: 18
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
    tib entry: 192.168.1.81/32, rev 8
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: 19
        remote binding: tsr: 192.168.1.97:0, tag: 16
    tib entry: 192.168.1.97/32, rev 2
        local binding:  tag: 16
        remote binding: tsr: 192.168.1.17:0, tag: 16
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
    tib entry: 192.168.1.112/28, rev 4
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: 17
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
```

- Perform a traceroute from your PE router to the loopback address of the PE router of the other team and verify that the results display the associated labels.

```
Tracing the route to 192.168.1.33


  1 192.168.1.50 [MPLS: Label 18 Exp 0] 164 msec 196 msec 200
msec
  2 192.168.1.114 [MPLS: Label 17 Exp 0] 56 msec 56 msec 56
msec
  3 192.168.1.65 40 msec 40 msec
```

Task 2: Disabling TTL Propagation

In this task, you will disable MPLS TTL propagation and verify the results. Team A will configure PE*wg*1 and P*wg*1. Team B will configure PE*wg*2 and P*wg*2.

Exercise Procedure

Complete these steps:

Step 1 On your assigned PE router, disable MPLS TTL propagation.

Step 2 On your assigned P router, disable MPLS TTL propagation.

Step 3 Verify that the other team has completed its configuration.

Exercise Verification

You have successfully completed this exercise when you attain these results:

- Perform a traceroute from your PE router to the loopback address of the PE router of the other team and compare this display to the display obtained in the previous task.

```
PE11#traceroute 192.168.1.33
Type escape sequence to abort.
Tracing the route to 192.168.1.33

  1 192.168.1.65 40 msec 40 msec *
```

Note When you are troubleshooting, it may become necessary to view the core routes when doing traces. If so, it will be necessary to re-enable TTL propagation. Doing so may affect the results of the traces shown in the lab exercise verification because additional hops and labs will be displayed.

Task 3: Configuring Conditional Label Distribution

For the label binding displays that you did in Task 2, you can see that a label is assigned to every prefix that is in the IP routing table of a router. This label assignment results in wasted label space and resources necessary to build unused label switched paths (LSPs). In this task you will use conditional label advertising to restrict the distribution of labels related to the WAN interfaces in the core.

Team A will configure PE*wg*1 and P*wg*1. Team B will configure PE*wg*2 and P*wg*2.

Exercise Procedure

Complete these steps:

Step 1 On your PE router, display the LSPs that are being built.

```
PE11#sh mpls for
Local  Outgoing    Prefix           Bytes tag  Outgoing
Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
16     16          192.168.1.97/32  0          Se0/0.111
point1point
17     Pop tag     192.168.1.112/28 0          Se0/0.111
point1point
18     17          192.168.1.64/28  0          Se0/0.111
point1point
19     Pop tag     192.168.1.81/32  0          Se0/0.111
point1point
20     18          192.168.1.33/32  0          Se0/0.111
point1point
```

Step 2 Note that an LSP has been built to the WAN interface that connects the other PE and P router. This LSP will never be used because traffic will not normally terminate at this point.

Step 3 On your assigned P and PE routers, configure conditional label distribution to allow only the distribution of labels related to the core loopback addresses and the interfaces that provide direct customer support.

Step 4 Verify that the other team has completed its configuration tasks.

Exercise Verification

You have completed this exercise when you attain these results:

- On your PE router, display the LSPs that are being built.

```
PE11#sh mpls f
Local  Outgoing    Prefix            Bytes tag  Outgoing
Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     16          192.168.1.97/32   0          Se0/0.111
point1point
17     Untagged    192.168.1.112/28  0          Se0/0.111
point1point
18     Untagged    192.168.1.64/28   0          Se0/0.111
point1point
19     Pop tag     192.168.1.81/32   0          Se0/0.111
point1point
20     18          192.168.1.33/32   0          Se0/0.111
point1point
```

Note An LSP is no longer built to the WAN interface that connects the other PE and P routers.

- On your P router display the LDP bindings.

```
P11#sh mpls ldp bind
  tib entry: 150.wg.0.0/16, rev 31
        local binding:  tag: 20
        remote binding: tsr: 192.168.1.97:0, tag: 20
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.wg.11.16/28, rev 36
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.wg.11.32/28, rev 37
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.17/32, rev 35
        local binding:  tag: 19
        remote binding: tsr: 192.168.1.97:0, tag: 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.33/32, rev 32
        local binding:  tag: 18
        remote binding: tsr: 192.168.1.97:0, tag: 17
        remote binding: tsr: 192.168.1.17:0, tag: 20
  tib entry: 192.168.1.48/28, rev 26
        local binding:  tag: imp-null
  tib entry: 192.168.1.64/28, rev 27
```



```
      local binding:  tag: 17
tib entry: 192.168.1.81/32, rev 34
      local binding:  tag: imp-null
      remote binding: tsr: 192.168.1.97:0, tag: 16
      remote binding: tsr: 192.168.1.17:0, tag: 19
tib entry: 192.168.1.97/32, rev 33
      local binding:  tag: 16
      remote binding: tsr: 192.168.1.97:0, tag: imp-null
      remote binding: tsr: 192.168.1.17:0, tag: 16
tib entry: 192.168.1.112/28, rev 30
      local binding:  tag: imp-null
```

Note The prefix assigned to the WAN interface connecting the other P and PE routers no longer has a remote label assigned. Further, none of the core WAN interfaces have remote labels assigned. This lessening of assignments results in a reduced label space, which saves memory resources.

Task 4: Removing Conditional Label Distribution

For the conditional label distribution displays that you did in Task 2, you can see that a label is not assigned to every prefix that is in the IP routing table of a router. In this task you will remove conditional label advertising so no restrictions are on the distribution of labels related to the WAN interfaces in the core.

Team A will configure PE*wg*1 and P*wg*1. Team B will configure PE*wg*2 and P*wg*2.

Exercise Procedure

Complete these steps:

Step 1 Remove conditional label distribution.

Step 2 Verify that the other team has completed its configuration task.

Exercise Verification

You have completed this exercise when you attain these results:

- On your PE router, display the LSPs that are being built.

```
PE11#sh mpls for
Local  Outgoing    Prefix           Bytes tag  Outgoing
Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
16     16          192.168.1.97/32  0          Se0/0.111
point1point
17     Pop tag     192.168.1.112/28 0          Se0/0.111
point1point
18     17          192.168.1.64/28  0          Se0/0.111
point1point
19     Pop tag     192.168.1.81/32  0          Se0/0.111
point1point
20     18          192.168.1.33/32  0          Se0/0.111
point1point
```

Next Step

- **Module 4:** MPLS Virtual Private Networks Technology

Lab Exercise 5-1: Initial MPLS VPN Setup

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

The company that you work for is a small service provider. Your workgroup has been tasked with creating two simple VPNs to support two new customers (customer A and customer B) who have just signed with you.

In this exercise, you will create a simple VPN for your customer. After completing this exercise, you will be able to meet these objectives:

- Configure Multiprotocol BGP to establish routing between the PE routers of your workgroup
- Configure the virtual routing and forwarding tables necessary to support your customer and establish your customer RIP routing using a simple VPN

Visual Objective

The figure illustrates what you will accomplish in this exercise.

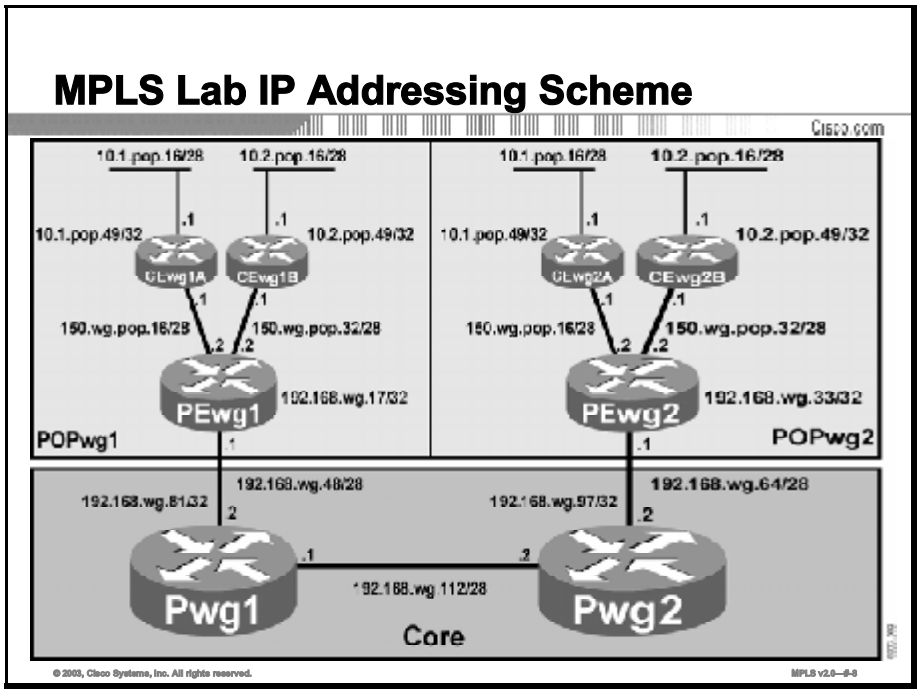

These exercises rely on the "Frame-Mode MPLS Configuration" lab exercises, where you established MPLS connectivity in your backbone. If this is the first set of exercises that you are undertaking, please refer to Task 1 in the lab exercise "Establishing the Service Provider IGP Routing Environment" to familiarize yourself with the IP addressing and routing in your workgroup.

Please also verify that MPLS has been enabled on all core interfaces in your backbone and that it has not been enabled on interfaces toward the customer workgroup routers or other service providers.

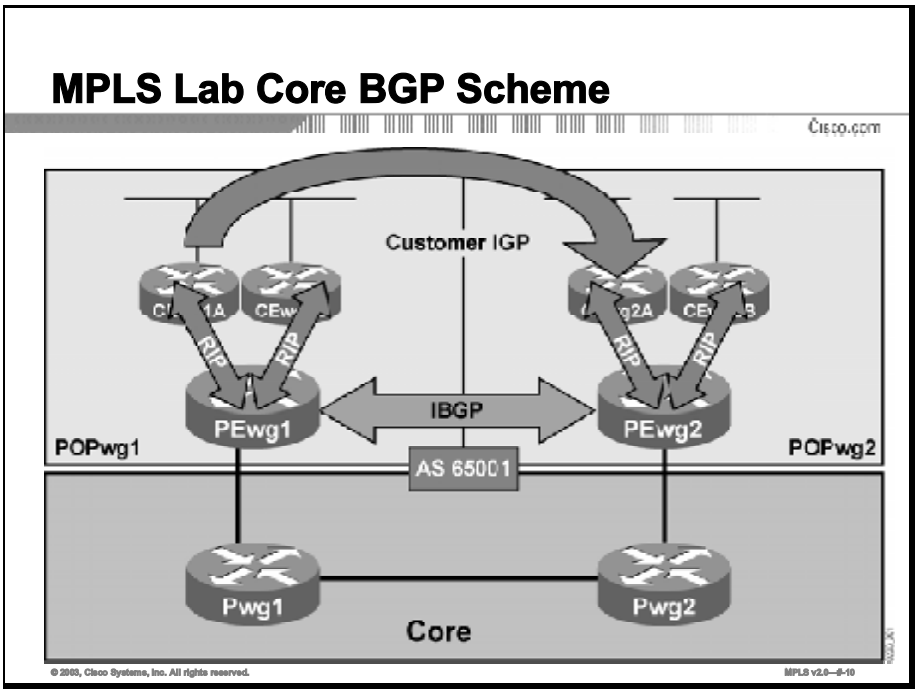

This exercise contains tasks that enable you to configure your core MPLS VPN infrastructure and to establish a simple any-to-any VPN service for a customer.

You will also test various provider edge-customer edge (PE-CE) routing options, ranging from RIP and Open Shortest Path First (OSPF) to running Border Gateway Protocol (BGP) between the PE and the CE routers.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

VPN-Related Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol.
address-family vpnv4	Selects VPNv4 address family configuration.
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF.
ip vrf *vrf-name*	Creates a virtual routing and forwarding table (VRF).
neighbor *ip-address* **activate**	Activates exchange of routes from address family under configuration for specified neighbor.
neighbor *ip-address* **route-reflector-client**	Configures a route reflector client on a route reflector.
neighbor next-hop-self	To configure the router as the next hop for a BGP-speaking neighbor or peer group, use the **neighbor next-hop-self** router configuration command. To disable this feature, use the **no** form of this command.
neighbor remote-as	To add an entry to the BGP or Multiprotocol Border Gateway Protocol (MP-BGP) neighbor table, use the **neighbor remote-as** router configuration command. To remove an entry from the table, use the **no** form of this command.
neighbor send-community	To specify that a communities attribute should be sent to a BGP neighbor, use the **neighbor send-community** command in address family or router configuration mode. To remove the entry, use the **no** form of this command.
neighbor update-source	To have the Cisco IOS software allow Internal Border Gateway Protocol (IBGP) sessions to use any operational interface for TCP connections, use the **neighbor update-source** router configuration command. To restore the interface assignment to the closest interface, which is called the "best local address," use the **no** form of this command.
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF.
rd *value*	Assigns a route distinguisher (RD) to a VRF.
redistribute bgp *as-number* **metric transparent**	Redistributes BGP routes into RIP with propagation of the multi-exit discriminator (MED) into the RIP hop count.
router bgp *as-number*	Selects BGP configuration.
route-target import|export *value*	Assigns a route target (RT) to a VRF.
show ip bgp neighbor	Displays information on global BGP neighbors.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPN IPv4 (VPNv4) routes associated with the specified VRF.
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF.
show ip vrf detail	Displays detailed VRF information.

telnet host /vrf *vrf-name*	Telnets to a CE router connected to the specified VRF.

Task 1: Configuring Multiprotocol BGP

In this section of the exercise, you will configure MP-BGP between the PE routers in your workgroup.

Team A will configure MP-BGP on PE*wg*1, and team B will perform the same task on PE*wg*2.

Exercise Procedure

Complete these steps:

Step 1 Activate the BGP process on your assigned router using AS 65001 as the AS number.

Step 2 Activate VPNv4 BGP sessions between your assigned PE router and the PE router being configured by the other team.

Step 3 Verify that the other team has completed its configuration tasks.

Exercise Verification

You have completed this exercise when you attain these results:

- Display the BGP neighbor information and ensure that BGP sessions have been established between the two PE routers.

Note The following displays are showing *WG 4*. Your display should be the same except where you see a "4" in the host name or in the addresses displayed. That number will be your *WG #*.

```
PEpop#sh ip bgp sum
BGP router identifier 192.168.wg.17, local AS number 65001
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd


192.168.wg.33    4 65001       6       6        1    0    0
00:02:23        0



PEpop#sh ip bgp sum
BGP router identifier 192.168.wg.33, local AS number 65001
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
192.168.wg.17    4 65001       9       9        1    0    0
00:05:24        0
```

```
PEpop#sh bgp nei
BGP neighbor is 192.168.wg.33,  remote AS 65001, internal link
  BGP version 4, remote router ID 192.168.wg.33
  BGP state = Established, up for 00:03:39
  Last read 00:00:39, hold time is 180, keepalive interval is
60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    IPv4 MPLS Label capability:
  Received 7 messages, 0 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0
  0 accepted prefixes consume 0 bytes
  Prefix advertised 0, suppressed 0, withdrawn 0

  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes:
0
Local host: 192.168.wg.17, Local port: 11022
Foreign host: 192.168.wg.33, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0
(0 bytes)

Event Timers (current time is 0xA12E784):
Timer          Starts    Wakeups            Next
Retrans             8          0             0x0
TimeWait            0          0             0x0
AckHold             7          5             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
```



```
iss: 1596106025  snduna: 1596106185  sndnxt: 1596106185
sndwnd:  16225

irs: 2134453172  rcvnxt: 2134453332  rcvwnd:      16225
delrcvwnd:    159


SRTT: 197 ms, RTTO: 984 ms, RTV: 787 ms, KRTT: 0 ms

minRTT: 44 ms, maxRTT: 300 ms, ACK hold: 200 ms

Flags: higher precedence, nagle


Datagrams (max data segment is 536 bytes):

Rcvd: 8 (out of order: 0), with data: 7, total data bytes: 159

Sent: 14 (retransmit: 0, fastretransmit: 0), with data: 7,
total data bytes: 159
```

Task 2: Configuring Virtual Routing and Forwarding Tables

In this task and the following task, you will establish simple VPNs for customer A and customer B. Team A will establish a VPN between CE*wg*1A and CE*wg*2A, and team B will establish a VPN between CE*wg*1B and CE*wg*2B. Each team is responsible for all PE router configurations related to its customer. This division of work between teams applies to all future exercises.

Exercise Procedure

Complete these steps:

Step 1 Design your VPN networks—decide on the RD and the RT numbering. Coordinate your number with the other team.

Note The easiest numbering plan would be to use the same values for the RD and the RT. Use simple values, for example, *wg:*10 for customer A and *wg:*20 for customer B.

Step 2 Create VRFs on the PE routers and associate the PE-CE interfaces into the proper VRFs; use simple yet descriptive VRF names (for example, "CE*wg*A" and "CE*wg*B").

Step 3 Your customer is using RIP as its IGP, so enable RIP for the VRF that you have created.

Step 4 Configure redistribution of RIP into BGP within the **ipv4 vrf** address family.

Step 5 Configure redistribution of BGP into RIP within the **ipv4 vrf** address family.

Step 6 Configure RIP metric propagation through MP-BGP by using the **redistribute bgp metric transparent** command in the RIP process.

Step 7 Ensure that RIP is enabled on all the CE routers. Make sure that all the networks (including loopbacks) are active in the RIP process.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify that you have the proper configuration of your virtual routing and forwarding (VRF) tables with **show ip vrf detail**. You should get a printout similar to the one here:

```
PEpop#sh ip vrf detail
VRF Customer_A; default RD wg:10; default VPNID <not set>
  Interfaces:
    Serial0/0.101
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:wg:10
  Import VPN route-target communities
    RT:wg:10
  No import route-map
```



```
  No export route-map
VRF Customer_B; default RD wg:20; default VPNID <not set>
  Interfaces:
    Serial0/0.102
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:wg:20
  Import VPN route-target communities
    RT:wg:20
  No import route-map
  No export route-map
```

- Check the routing protocols running in your VRF with the **show ip protocol vrf** command. When executed on WG2PE2, it will produce a printout similar to the one here:

```
PEpop#sh ip prot vrf Customer_A
Routing Protocol is "bgp 65001"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Redistributing: rip
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.wg.33         200      15:05:06
  Distance: external 20 internal 200 local 200


Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: bgp 65001, rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial0/0.101         2     2
  Maximum path: 4
  Routing for Networks:
    Interface             Send  Recv  Triggered RIP  Key-chain
`    10.0.0.0
```

```
    150.wg.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.wg.pop.17          120      00:00:27
  Distance: (default is 120)

PEpop#sh ip prot vrf Customer_B
Routing Protocol is "bgp 65001"
 Outgoing update filter list for all interfaces is not set
 Incoming update filter list for all interfaces is not set
 IGP synchronization is disabled
 Automatic route summarization is disabled
 Redistributing: rip
 Maximum path: 1
 Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.wg.33         200      15:04:27
  Distance: external 20 internal 200 local 200

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 20 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: bgp 65001, rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial0/0.102         2     2
  Maximum path: 4
  Routing for Networks:
    Interface             Send  Recv  Triggered RIP  Key-chain
    10.0.0.0
    150.wg.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.wg.pop.33          120      00:00:07
  Distance: (default is 120)
```


- Verify the per-VRF routing table on the PE router with the **show ip route vrf** command. It will produce a printout similar to the one here:

```
PEpop#sh ip route vrf Customer_A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B       10.1.pop.49/32 [200/1] via 192.168.wg.33, 15:10:04
R       10.1.pop.49/32 [120/1] via 150.wg.pop.17, 00:00:24,
Serial0/0.101
B       10.1.pop.16/28 [200/1] via 192.168.wg.33, 15:10:04
R       10.1.pop.16/28 [120/1] via 150.wg.pop.17, 00:00:24,
Serial0/0.101
     150.wg.0.0/28 is subnetted, 2 subnets
B       150.wg.pop.16 [200/0] via 192.168.wg.33, 15:46:04
C       150.wg.pop.16 is directly connected, Serial0/0.101

PEpop#sh ip route vrf Customer_B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
```

```
R     10.2.pop.49/32 [120/1] via 150.wg.pop.33, 00:00:01,
Serial0/0.102
B     10.2.pop.49/32 [200/1] via 192.168.wg.33, 15:09:26
R     10.2.pop.16/28 [120/1] via 150.wg.pop.33, 00:00:01,
Serial0/0.102
B     10.2.pop.16/28 [200/1] via 192.168.wg.33, 15:09:26
    150.wg.0.0/28 is subnetted, 2 subnets
B     150.wg.pop.32 [200/0] via 192.168.wg.33, 15:46:11
C     150.wg.pop.32 is directly connected, Serial0/0.102
```

- Use the **show ip bgp vpnv4 vrf** command to display the BGP routing table associated with a VRF. The printout from the WG2PE4 router is shown here:

```
PEpop#show ip bgp vpnv4 vrf Customer_A
BGP table version is 47, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i -internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*> 10.1.pop.16/28    150.wg.pop.17             1
32768 ?
*> 10.1.pop.49/32    150.wg.pop.17             1
32768 ?
*>i10.1.pop.16/28    192.168.wg.33            1    100      0
?
*>i10.1.pop.49/32    192.168.wg.33            1    100      0
?
*> 150.wg.pop.16/28   0.0.0.0                 0         32768
?
*>i150.wg.pop.16/28   192.168.wg.33            0    100
0 ?


PEpop#show ip bgp vpnv4 vrf Customer_B
BGP table version is 47, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
```



```
Route Distinguisher: wg:20 (default for vrf Customer_B)
*> 10.2.pop.16/28    150.wg.pop.33              1
32768 ?
*> 10.2.pop.49/32    150.wg.pop.33              1
32768 ?
*>i10.2.pop.16/28    192.168.wg.33              1    100
0 ?
*>i10.2.pop.49/32    192.168.wg.33              1    100
0 ?
*> 150.wg.pop.32/28   0.0.0.0                  0
32768 ?
*>i150.wg.pop.32/28  192.168.wg.33                   0    100
0 ?
```

- On a CE router, use the **show ip route** command to verify that the router is receiving all VPN routes. Also verify that no routes from the other customer or the MPLS core are being received. On CE*wg*1A, the printout is similar to the one here:

```
CEpopA#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R       10.1.pop.49/32 [120/2] via 150.wg.pop.18, 00:00:14,
Serial0/0.101
C       10.1.pop.49/32 is directly connected, Loopback0
R       10.1.pop.16/28 [120/2] via 150.wg.pop.18, 00:00:14,
Serial0/0.101
C       10.1.pop.16/28 is directly connected, Ethernet0/0
     150.wg.0.0/28 is subnetted, 2 subnets
R       150.wg.pop.16 [120/1] via 150.wg.pop.18, 00:00:14,
Serial0/0.101
C       150.wg.pop.16 is directly connected, Serial0/0.101
```

```
Use ping and trace on the CE routers to verify connectivity
across the VPN.
CEpopA#traceroute 150.wg.pop.17

Type escape sequence to abort.
Tracing the route to 150.wg.pop.17

  1 150.wg.pop.18 12 msec 12 msec 12 msec
  2 150.wg.pop.18 60 msec 60 msec 60 msec
  3 150.wg.pop.17 77 msec 72 msec *

CEpopA#ping 150.wg.pop.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.17, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
144/146/148 ms
```

- Use the **show ip route** command on the PE routers to verify that the customer routes are not in the global IP routing table.

```
PEpop#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     192.168.wg.0/24 is variably subnetted, 7 subnets, 2 masks
D       192.168.wg.97/32 [90/2809856] via 192.168.wg.50,
19:14:54, Serial0/0.111
D       192.168.wg.112/28
           [90/2681856] via 192.168.wg.50, 19:14:54,
Serial0/0.111
D       192.168.wg.64/28 [90/3193856] via 192.168.wg.50,
19:14:54, Serial0/0.111
```



```
D       192.168.wg.81/32 [90/2297856] via 192.168.wg.50,
19:14:54, Serial0/0.111
D       192.168.wg.33/32 [90/3321856] via 192.168.wg.50,
19:14:54, Serial0/0.111
C       192.168.wg.48/28 is directly connected, Serial0/0.111
C       192.168.wg.17/32 is directly connected, Loopback0
```

- Use **ping** and **trace** on the PE routers to verify that you *cannot* reach your customer networks from global address space.

```
PEpop#ping 150.wg.pop.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.17, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)


PEpop#ping 150.wg.pop.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.33, timeout is 2
seconds:
.....
```

- Use the **ping vrf** command on the PE routers to verify that you can reach your customer networks from global address space.

```
PEpop#ping vrf Customer_A 150.wg.pop.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.17, timeout is 2
seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/31/36 ms


PEpop#ping vrf Customer_B 150.wg.pop.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.33, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/28/32 ms


PEpop#ping vrf Customer_B 150.wg.pop.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.33, timeout is 2
seconds:
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms


PE pop #ping vrf Customer_A 150.wg. pop .17

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 150.wg. pop .17, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/117/120 ms
```


Lab Exercise 5-2: Running EIGRP Between PE and CE Routers

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

Some customers use EIGRP as the routing protocol in their VPN; sometimes it is even combined with RIP or BGP at other sites. In this exercise, the customers of the service provider have decided to migrate some of their sites to EIGRP.

You will deploy EIGRP as the PE-CE routing protocol in the VPN of your customer. After completing this exercise, you will be able to meet these objectives:

- Convert one of each of the customer sites to EIGRP (from RIP) and establish VPN routing using EIGRP. The other site will remain running RIP as the IGP.

Visual Objective

The figure illustrates what you will accomplish in this exercise.

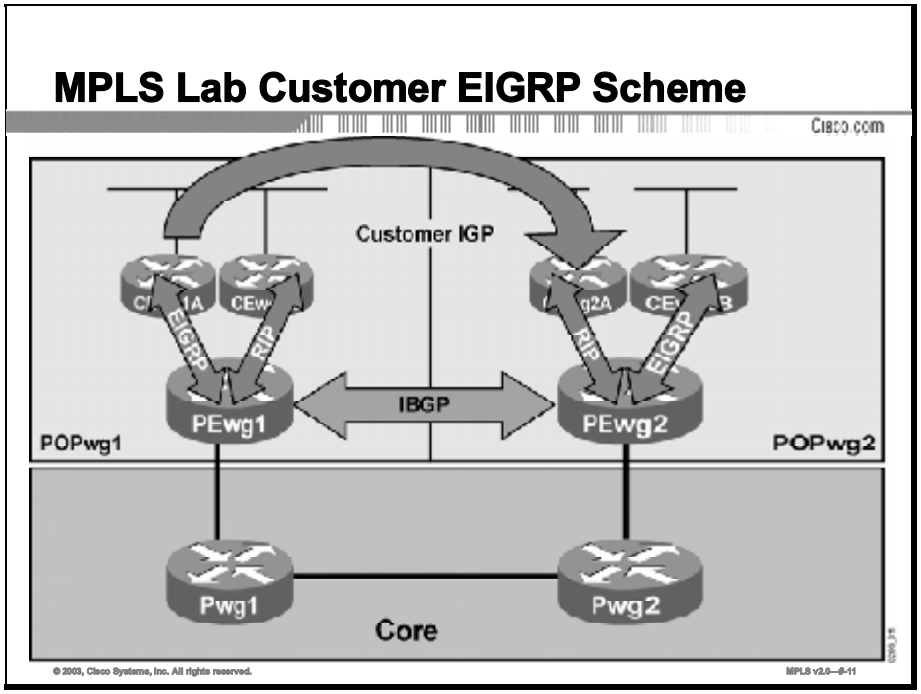

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

OSPF Commands

Command	Description
address-family ipv4 [**multicast** | **unicast** | **vrf** *vrf-name*]	Enters address family configuration mode and creates a VRF. The VRF name (or tag) must match the VRF name that was created in Step 3 from Task 2 above.
network *ip-address network-mask*	Specifies the network for the VRF. The network statement is used to identify which interfaces to include in EIGRP. The VRF must be configured with addresses that fall within the subnet range of the configured network statement.
redistribute protocol [*process-id*] {**level-1** | **level-1-2** | **level-2**} [*as-number*] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-map** *map-name*] [**match** {**internal** | **external 1** | **external 2**}] [**tag** *tag-value*] [**route-map** *map-tag*] [subnets]	Redistributes BGP into the EIGRP. The AS number and metric of the BGP network are configured in this step. BGP must be redistributed into EIGRP in order for the CE site to accept the BGP routes that carry the EIGRP information. A metric must also be specified for the BGP network and is configured in this step.
router eigrp *as-number*	Enters router configuration mode and creates an EIGRP routing process.
show ip eigrp vrf *vrf-name* **interfaces**	Displays EIGRP interfaces that are defined under the specified VRF. If an interface is specified, only that interface is displayed. Otherwise, all interfaces on which EIGRP is running as part of the specified VRF are displayed.
show ip eigrp vrf *vrf-name* **neighbors**	Displays when VRF neighbors become active and inactive. This command can be used to help debug transport problems.
show ip eigrp vrf *vrf-name* **topology**	Displays VRF entries in the EIGRP topology table. This command can be used to determine Diffusing Update Algorithm (DUAL) states and to debug possible DUAL problems.
show ip vrf	Displays the set of defined VRFs and associated interfaces. This command is used to verify that the correct RDs are configured for the VRF.

Task 1: Enabling an EIGRP VPN

In this task your customer has decided to convert its IGP from RIP to EIGRP. Team A will convert the customer A site 1, CE*wg*1A, from RIP to EIGRP and establish a simple VPN. Team B will convert the customer B site 2, CE*wg*2B, from RIP to EIGRP and establish a simple VPN.

Each team is responsible for all PE router configurations related to its customer.

Exercise Procedure

Complete these steps:

Step 1 Disable RIP and configure EIGRP on the router of your customer. Team A will configure CE*wg*1A, and team B will configure CE*wg*2B. Use your *wg#* as the AS number for EIGRP.

Note Do not forget to remove the address family from the RIP routing process. This action will disable the sites still running RIP as the customer edge-provider edge (CE-PE) routing protocol.

Step 2 On your assigned PE router, configure redistribution of EIGRP into BGP within the **ipv4 vrf** address family.

Step 3 On your assigned PE router, configure redistribution of BGP into EIRGP within the **ipv4 vrf** address family.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify that EIGRP has been activated on the proper interfaces.

```
PEpop#sh ip eigrp int
IP-EIGRP interfaces for process 1
                        Xmit Queue   Mean   Pacing Time
Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow
Timer   Routes
Se0/0.111          1        0/0       600       0/15
2991           0
Lo0                0        0/0         0       0/10
0           0
```

- Verify that EIGRP adjacencies have been established between the CE and PE routers.

```
PEpop#sh ip eigrp vrf Customer_A nei
IP-EIGRP neighbors for process 4
H   Address          Interface       Hold Uptime   SRTT   RTO
Q  Seq Type
```

```
                                          (sec)         (ms)
Cnt Num
0   150.wg.pop.17      Se0/0.101         14 00:pop:51  340
2040  0  4


PEpop#sh ip eigrp vrf Customer_B nei
IP-EIGRP neighbors for process 4
H   Address          Interface       Hold Uptime   SRTT   RTO
Q  Seq Type
                                          (sec)         (ms)
Cnt Num
0   150.wg.pop.33      Se0/0.102         14 00:02:29 1050
5000  0  2
```

- Check the EIGRP topology database on the CE routers.

```
PEpop#sh ip eigrp vrf Customer_A topology
IP-EIGRP Topology Table for AS(4)/ID(150.wg.pop.18) Routing
Table: Customer_A


Codes: P - Passive, A - Active, U - Update, Q - Query, R -
Reply,
       r - reply Status, s - sia Status


P 10.1.pop.49/32, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.1.pop.49/32, 1 successors, FD is 2297856
         via 150.wg.pop.17 (2297856/128256), Serial0/0.101
P 10.1.pop.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.1.pop.16/28, 1 successors, FD is 2195456
         via 150.wg.pop.17 (2195456/281600), Serial0/0.101
P 150.wg.pop.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 150.wg.pop.16/28, 1 successors, FD is 2169856
         via Connected, Serial0/0.101


PEpop#sh ip eigrp vrf Customer_B topology
IP-EIGRP Topology Table for AS(4)/ID(150.wg.pop.34) Routing
Table: Customer_B


Codes: P - Passive, A - Active, U - Update, Q - Query, R -
Reply,
       r - reply Status, s - sia Status
```



```
P 10.2.pop.49/32, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.2.pop.49/32, 1 successors, FD is 2297856
         via 150.wg.pop.33 (2297856/128256), Serial0/0.102
P 10.2.pop.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.2.pop.16/28, 1 successors, FD is 2195456
         via 150.wg.pop.33 (2195456/281600), Serial0/0.102
P 150.wg.pop.32/28, 1 successors, FD is 2169856
         via Connected, Serial0/0.102
P 150.wg.pop.32/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
```

- Verify connectivity across the VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers.

```
CEpopB#ping 150.wg.pop.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.33, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
144/147/152 ms


CEpopA#ping 150.wg.pop.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.wg.pop.17, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
144/147/152 ms


CEpopB#trace 150.wg.pop.33
Type escape sequence to abort.
Tracing the route to 150.wg.pop.33
1 150.wg.pop.34 12 msec 12 msec 12 msec
2 150.wg.pop.34 64 msec 60 msec 60 msec
3 150.wg.pop.33 77 msec 76 msec *


CEpopA#trace 150.wg.pop.17
Type escape sequence to abort.
Tracing the route to 150.wg.pop.17
```

```
1 150.wg.pop.18 12 msec 12 msec 12 msec
2 150.wg.pop.18 64 msec 60 msec 64 msec
3 150.wg.pop.17 76 msec 76 msec *

PEpop#ping vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
116/119/120 ms

PEpop#ping vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms

PEpop#ping vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms

PEpop#trace vrf Customer_A 10.2.pop.49
Type escape sequence to abort.
Tracing the route to 10.1.pop.49
  1 150.wg.pop.17 12 msec 12 msec *

PEpop#trace vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
Tracing the route to 10.2.pop.49
  1 150.wg.pop.33 12 msec 12 msec *

PEpop#trace vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
Tracing the route to 10.2.pop.49
  1 150.wg.pop.33 60 msec 60 msec *
```



```
PEpop#trace vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Tracing the route to 10.1.pop.49
  1 150.wg.pop.17 60 msec 60 msec *
```

Lab Exercise 5-3: Running OSPF Between PE and CE Routers

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

Some customers insist on using OSPF as the routing protocol in their VPN, sometimes even combined with RIP or BGP at other sites.

- Convert one of each of the customer sites to OSPF (from RIP) and establish VPN routing using OSPF.
- Complete the OSPF migration.

Visual Objective

The figure illustrates what you will accomplish in this exercise.

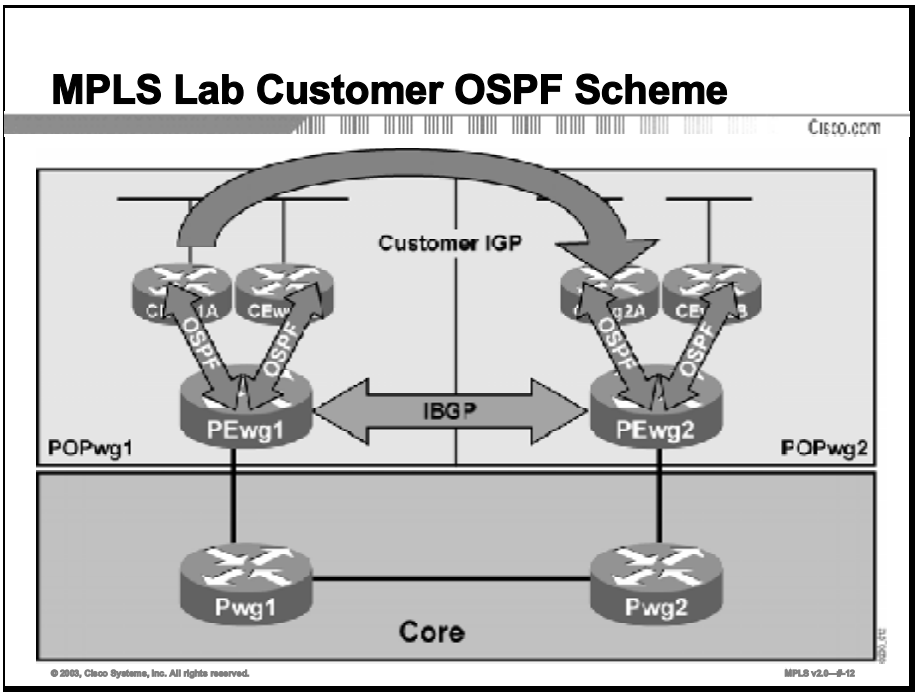

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

OSPF Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol.
default-information originate always	Generates a default route into OSPF.
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF.
ip vrf *vrf-name*	Creates a virtual routing and forwarding table (VRF).
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF.
rd *value*	Assigns an RD to a VRF.
redistribute bgp *as-number* **subnets**	Redistributes BGP routes (including subnet routes) into OSPF.
router bgp *as-number*	Selects BGP configuration.
router ospf ***process*** **vrf** *vrf-name*	Starts an OSPF process within the specified VRF.
route-target import|export *value*	Assigns an RT to a VRF.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPNv4 routes associated with the specified VRF.
show ip ospf database	Displays OSPF database information.
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF.
show ip vrf detail	Displays detailed VRF information.
telnet host /vrf *vrf-name*	Telnets to a CE router connected to the specified VRF.

Task 1: Configuring Virtual Routing and Forwarding Tables

The customer is performing a phase conversion of its IGP to OSPF. In this task your customer has decided to convert the site running RIP to OSPF. Team A will convert the customer A site 2, CE*wg*2A, from RIP to OSPF and establish a simple VPN. Team B will convert the customer B site 1, CE*wg1*B, from RIP to OSPF and establish a simple VPN.

Each team is responsible for all PE router configurations related to its customer.

Exercise Procedure

Complete these steps:

Step 1 Disable RIP on the CE routers of your customer.

Step 2 Configure OSPF on the CE routers of your customer. (Use an OSPF process ID of 1 for team A [CE*wg*2A] and a process ID of 2 for team B [CE*wg*1B]). Use the areas in the CE router as detailed here.

Area	Interface(s)
Area 0	WAN interface toward PE router Loopback 0
Area 1	E0/0

Step 3 Configure OSPF (use an OSPF process ID of 1 for team A and a process ID of 2 for team B) in the VRFs on PE routers using the router ospf vrf command. Use OSPF Area 0 on the PE-CE link.

Step 4 Configure redistribution from OSPF to MP-BGP using the redistribute ospf command inside the VRF address family configuration.

Step 5 Configure redistribution from MP-BGP to OSPF using the redistribute bgp subnets command in the OSPF router configuration.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify the OSPF adjacency on the PE*wg*1 and PE*wg*2 routers using the **show ip ospf neighbor** command.

```
PEpop#sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address
Interface
10.2.pop.49      0   FULL/  -         00:00:39
150.wg.pop.33   Serial0/0.102


PEpop#sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address
Interface
10.1.pop.49      0   FULL/  -         00:00:33
150.wg.pop.17   Serial0/0.101
```

- Check the OSPF topology database on CE*wg*1B and CE*wg*2A. You should see router link states (resulting from OSPF connectivity between the PE and the CE routers) and type 5 external link states (all other VPN routes originated in RIP or BGP), but no summaries. A sample printout from CE*pop*B and CE*pop*A are shown here:

```
CEpopB#sh ip ospf data
            OSPF Router with ID (10.2.pop.49) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.2.pop.49     10.2.pop.49      920         0x8000000D 0xA60D
3
150.wg.pop.34   150.wg.pop.34    1684        0x80000008 0x6449
2

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum
10.2.pop.16     10.2.pop.49     920         0x80000003 0xDDB6

                Router Link States (Area 1)
```

```
Link ID         ADV Router      Age         Seq#
Checksum Link count
10.2.pop.49     10.2.pop.49     920         0x80000003 0x36EC
1

                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
10.2.pop.49     10.2.pop.49      920         0x80000003 0x92DA
150.wg.pop.32   10.2.pop.49      920         0x80000003 0x209F

                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
150.wg.pop.34   10.2.pop.49      921         0x80000003 0x5855

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#
Checksum Tag
10.2.pop.16     150.wg.pop.34     329         0x80000001
0x7BC1   3489725929
10.2.pop.49     150.wg.pop.34     329         0x80000001
0xA6D4   3489725929
150.wg.pop.32   150.wg.pop.34     329         0x80000001
0x1C23   3489725929

CEpopA#sh ip ospf data
            OSPF Router with ID (10.1.pop.49) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.1.pop.49     10.1.pop.49      1941        0x80000005 0xE305
3
150.wg.pop.18   150.wg.pop.18     885         0x80000009
0x32B6   2

                Summary Net Link States (Area 0)
```



```
Link ID         ADV Router      Age         Seq#
Checksum
10.1.pop.16     10.1.pop.49      1941        0x80000002 0xE1B3

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.1.pop.49      10.1.pop.49      1941        0x80000002
0x3BE8   1

                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
10.1.pop.49     10.1.pop.49      1941        0x80000002 0x96D7
150.wg.pop.16   10.1.pop.49      1941        0x80000002 0xB817

                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
150.wg.pop.18   10.1.pop.49     1944        0x80000002 0xF0CC

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#
Checksum Tag
10.1.pop.16     150.wg.pop.18     347         0x80000001
0xEB62   3489725929
10.1.pop.49     150.wg.pop.18     347         0x80000001
0x1775   3489725929
150.wg.pop.16   150.wg.pop.18     347         0x80000001
0x213E   3489725929
```

- Verify connectivity across the VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers.

```
CEpopA#ping 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
!!!!!
CEpopB#ping 10.2.pop.49
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
144/147/152 ms

PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/119/120 ms


PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_B 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/121/132 ms

 PEpop#ping vrf Customer_B 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/28/32 ms
```



```
PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms


PEpop#ping vrf Customer_B 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_B 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms


PEpop#trace vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Tracing the route to 10.1.pop.49


  1 150.wg.pop.17 60 msec 60 msec *
```

Task 2: Completing the OSPF Migration

In this task your customer has decided to have one IGP, OSPF. This decision means that the sites that are running EIGRP will have to be converted to OSPF. Team A will convert the customer A site 1, CE*wg*1A, from EIGRP to OSPF and establish a simple VPN. Team B will convert the customer B site 2, CE*wg*2B, from EIGRP to OSPF and establish a simple VPN.

Each team is responsible for all PE router configurations related to its customer.

Exercise Procedure

Complete these steps:

Step 1 Disable EIGRP and configure OSPF on the CE routers of your customer. Configure OSPF (use an OSPF process ID of 1 for team A and a process ID of 2 for team B) areas in the CE router according to the information here.

Area	Interface(s)
Area 0	WAN interface toward PE router Loopback 0
Area 1	E0/0

Step 2 Configure OSPF (use an OSPF process ID of 1 for team A and a process ID of 2 for team B) in the VRFs on PE routers using the **router ospf vrf** command. Use OSPF Area 0 on the PE-CE link.

Step 3 Configure redistribution from OSPF to MP-BGP using the **redistribute ospf** command inside the VRF address family configuration.

Step 4 Configure redistribution from MP-BGP to OSPF using the **redistribute bgp subnets** command in the OSPF router configuration.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify the OSPF adjacency on PE*wg*1 and PE*wg*2 routers using the **show ip ospf neighbor** command.

```
PEpop#sh ip ospf nei
Neighbor ID     Pri   State           Dead Time   Address
Interface
10.1.pop.49       0   FULL/  -        00:00:36
150.wg.pop.17     Serial0/0.101
10.2.pop.49       0   FULL/  -        00:00:37
150.wg.pop.33     Serial0/0.102


PEpop#sh ip ospf nei
```



```
Neighbor ID     Pri   State           Dead Time   Address
Interface
10.2.pop.49       0   FULL/  -        00:00:30
150.wg.pop.33    Serial0/0.102
10.1.pop.49       0   FULL/  -        00:00:39
150.wg.pop.17    Serial0/0.101
```

- Check the OSPF topology database on CE*wg*1A and CE*wg*2B. You should see router link states (resulting from OSPF connectivity between the PE and the CE routers) and type 5 external link states (all other VPN routes originated in RIP or BGP), but no summaries. A sample printout from CE*pop*A and CE*pop*B are shown here:

```
CEwg2B#sh ip ospf data

            OSPF Router with ID (10.2.pop.49) (Process ID 2)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.2.pop.49      10.2.pop.49      140         0x80000003
0x347A   2
150.4.pop.34     150.4.pop.34     141         0x80000002
0x8925   2

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum
10.2.pop.16      150.4.pop.34     146         0x80000001
0x2F26
10.2.pop.49      150.4.pop.34     146         0x80000001
0xE34A
10.2.pop.16      10.2.pop.49      1129        0x80000001
0xCFC4
150.4.pop.32     150.4.pop.34     146         0x80000001
0x7689

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.2.pop.49      10.2.pop.49      1139        0x80000001
0x31F0   1

                Summary Net Link States (Area 1)
```

```
Link ID         ADV Router      Age         Seq#
Checksum
10.2.pop.16      10.2.pop.49      136         0x80000001
0x67EC
10.2.pop.49      10.2.pop.49      157         0x80000001
0x1C11
150.4.pop.32     10.2.pop.49      157         0x80000001
0xAE50
150.4.pop.32     10.2.pop.49      1160        0x80000001
0x9965

                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
150.4.pop.34     10.2.pop.49      157         0x80000001
0xD11B

CEwg1A#sh ip ospf data

            OSPF Router with ID (10.1.pop.65) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.1.pop.65      10.1.pop.65      965         0x80000003
0xD678   3
150.4.pop.18     150.4.pop.18     964         0x80000002
0xC71D   2

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#
Checksum
10.1.pop.16      10.1.pop.65      1271        0x80000001
0x95F2
10.1.pop.16      150.4.pop.18     969         0x80000001
0x97CE
10.1.pop.49      150.4.pop.18     969         0x80000001
0x4CF2
150.4.pop.16     150.4.pop.18     969         0x80000001
0x73AC
```



```
                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum Link count
10.1.pop.65      10.1.pop.65      1280        0x80000001
0x45C2   1

                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
10.1.pop.49      10.1.pop.65      1281        0x80000001
0x4A17
10.1.pop.16      10.1.pop.65      983         0x80000001
0x172F
10.1.pop.49      10.1.pop.65      983         0x80000001
0xCB53
150.4.pop.16     10.1.pop.65      1303        0x80000001
0xF30E
150.4.pop.16     10.1.pop.65      983         0x80000001
0xF20D

                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#
Checksum
150.4.pop.18     10.1.pop.65      983         0x80000001
0x2CC3
```

- Check the IP routing table on CE*wg*1A and CE*wg*2B and note the OSPF interarea (IA) routes in the routing table.

```
CEwg1A#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set


     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA    10.1.pop.49/32 [110/66] via 150.4.pop.18, 00:31:44,
Serial0/0.101
C       10.1.pop.49/32 is directly connected, Loopback0
O IA    10.1.pop.16/28 [110/75] via 150.4.pop.18, 00:31:44,
Serial0/0.101
C       10.1.pop.16/28 is directly connected, Ethernet0/0
C       10.1.pop.64/28 is directly connected, Loopback1
     150.4.0.0/28 is subnetted, 2 subnets
O IA    150.4.pop.16 [110/2] via 150.4.pop.18, 00:32:14,
Serial0/0.101
C       150.4.pop.16 is directly connected, Serial0/0.101


CEwg2B#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route


Gateway of last resort is not set


     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA    10.2.pop.49/32 [110/66] via 150.4.pop.34, 00:55:59,
Serial0/0.102
C       10.2.pop.49/32 is directly connected, Loopback0
O IA    10.2.pop.16/28 [110/75] via 150.4.pop.34, 00:55:59,
Serial0/0.102
C       10.2.pop.16/28 is directly connected, Ethernet0/0
     150.4.0.0/28 is subnetted, 2 subnets
C       150.4.pop.32 is directly connected, Serial0/0.102
O IA    150.4.pop.32 [110/2] via 150.4.pop.34, 00:55:59,
Serial0/0.102
```


- Verify connectivity across the VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers.

```
CEpopA#ping 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
148/148/149 ms


CEpopB#ping 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
144/148/153 ms


PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/119/120 ms


PEpop#ping vrf Customer_A 10.1.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_B 10.2.pop.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
116/121/132 ms

 PEpop#ping vrf Customer_B 10.2.pop.49
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/28/32 ms


PEpop#ping vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms


PEpop#ping vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms


PEpop#ping vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
116/118/120 ms


PEpop#trace vrf Customer_A 10.1.pop.49
Type escape sequence to abort.
Tracing the route to 10.1.pop.49


  1 150.wg.pop.17 60 msec 60 msec *
```



```
PEpop#trace vrf Customer_B 10.2.pop.49
Type escape sequence to abort.
Tracing the route to 10.2.pop.49

  1 150.wg.pop.33 60 msec 60 msec *
```

Lab Exercise 5-4: Running BGP Between PE and CE Routers

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

Your customer has indicated that it wants to have a backup link for a selected site for redundancy. This addition will produce a multihomed environment. As a result, it is necessary to use BGP as the CE-to-PE routing protocol. The provider has decided to do this conversion in a phased implementation. The existing links will be converted to BGP and then the backup links will be added and activated.

In this exercise, you will convert the CE-to-PE routing protocol of your customer to BGP. After completing this exercise, you will be able to meet these objectives:

- Enable EBGP as the CE-to-PE link routing protocol
- Enable a backup link
- Configure BGP to control the selection of primary and backup links

Visual Objective

Team A will configure BGP between CE*wg*1A and PE*wg*1 and between CE*wg*2A and PE*wg*2. Team B will configure BGP between CE*wg*1B and PE*wg*1 and between CE*wg*2B and PE*wg*2.

The figure illustrates what you will accomplish in this exercise.

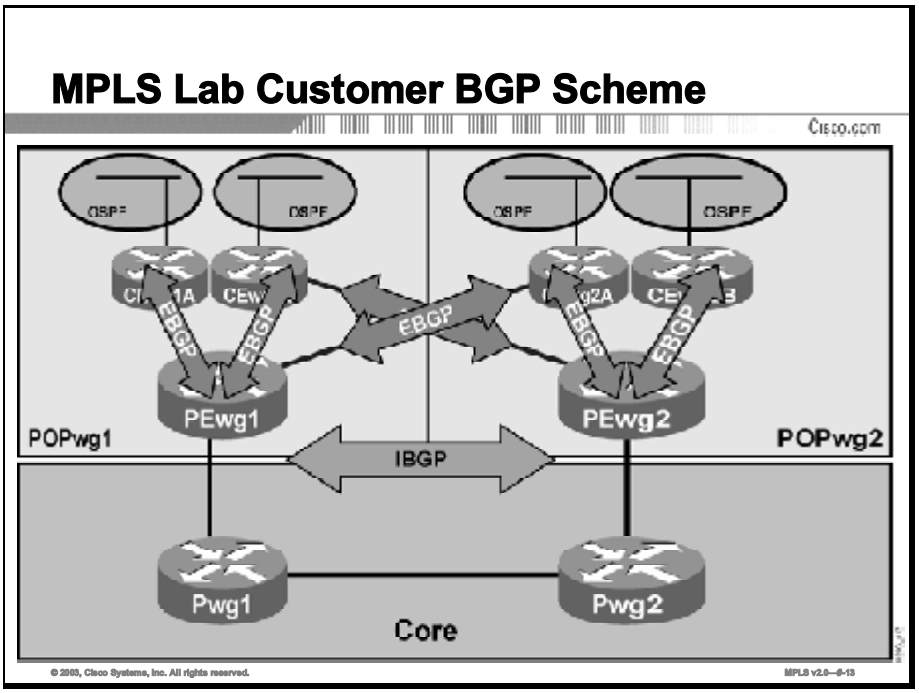

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

BGP Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol.
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF.
ip vrf *vrf-name*	Creates a virtual routing and forwarding table (VRF).
neighbor *ip-address* **as-override**	To configure a PE router to override the AS number of a site with the AS number of a provider, use the **neighbor as-override** command in router configuration mode. To remove VPNv4 prefixes from a specified router, use the **no** form of this command.
neighbor *ip-address* **route-map** *name* **in|out**	Applies a route map to BGP updates received from or sent to the specified neighbor.
no neighbor *ip-address* **shutdown**	Enables a BGP neighbor previously disabled with the **neighbor shutdown** command.
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF.
rd *value*	Assigns an RD to a VRF.
route-map *name* **permit** *seq*	Creates an entry in a route map.
router bgp *as-number*	Selects BGP configuration.
route-target import|export *value*	Assigns an RT to a VRF.
set metric *value*	Sets the BGP MED attribute in a route map.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPNv4 routes associated with the specified VRF.
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF.
telnet host /vrf *vrf-name*	Telnets to a CE router connected to the specified VRF.

Task 1: Configuring BGP as the PE-CE Routing Protocol

In this task you will make BGP the routing protocol between the PE and your customer routers. OSPF will remain the customer IGP. You will need to redistribute from BGP to OSPF and from OSPF to BGP on the routers of your customer. You will establish simple VPNs for customer A and customer B. Team A will establish a VPN between CE*wg*1A and CE*wg*2A, and team B will establish a VPN between CE*wg*1B and CE*wg*2B. Each team is responsible for all PE router configurations related to its customer.

Step 1 Activate the BGP routing process on the CE routers of your customer using AS650wg1 for customer A and AS 650wg2 for customer B.

Step 2 Remove OSPF on the associated PE router and activate the BGP neighbor relationship between each CE router and its associated PE router.

Step 3 Because both of your customer sites are using the same AS number, you will need to re-enable the allowas-in feature.

Exercise Verification

You have completed this exercise when you attain these results:

- Check BGP connectivity with the **show ip bgp summary** command on the CE routers.

```
CEpopA#sh ip bgp sum
BGP router identifier 10.1.pop.49, local AS number 6500wg
BGP table version is 10, main routing table version 10
9 network entries and 9 paths using 1197 bytes of memory
2 BGP path attribute entries using 120 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 9/30 prefixes, 9/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.18     4 650wg1     617     618       10    0    0
09:50:35        3

CEpopA#sh ip bgp
BGP table version is 63, local router ID is 10.1.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
```



```
*> 10.1.41.16/28    0.0.0.0                  0         32768 ?
*> 10.1.41.49/32    0.0.0.0                  0         32768 ?
*> 10.1.42.16/28    150.4.41.18                            0
65001 65001 ?
*> 10.1.42.49/32    150.4.41.18                            0
65001 65001 ?
*> 150.4.41.16/28   0.0.0.0                  0         32768 ?
*> 150.4.42.16/28   150.4.41.18                            0
65001 65001 ?
```

CE*pop*A#**sh ip bgp sum**

```
BGP router identifier 10.1.pop.49, local AS number 6500wg
BGP table version is 22, main routing table version 22
7 network entries and 7 paths using 931 bytes of memory
2 BGP path attribute entries using 120 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 14/33 prefixes, 14/7 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.18     4 65001     638     637       22    0    0
10:10:56        4
```

CE*pop*A#**sh ip bgp**

```
BGP table version is 38, local router ID is 10.1.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
*> 10.1.41.16/28    150.4.42.18                            0
65001 65001 ?
*> 10.1.41.49/32    150.4.42.18                            0
65001 65001 ?
*> 10.1.42.16/28    0.0.0.0                  0         32768 ?
*> 10.1.42.49/32    0.0.0.0                  0         32768 ?
*> 150.4.41.16/28   150.4.42.18                            0
65001 65001 ?
*> 150.4.42.16/28   0.0.0.0                  0         32768 ?
```

PE*pop*#**sh ip bgp vpn all**

```
BGP table version is 145, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*> 10.1.pop.16/28    150.wg.pop.17             0             0
6500wg ?
*> 10.1.pop.49/32    150.wg.pop.17             0             0
6500wg ?
* i10.1.pop.16/28    192.168.wg.33             0    100      0
6500wg ?
* i10.1.pop.49/32    192.168.wg.33             0    100      0
6500wg ?
r> 150.wg.pop.16/28  150.wg.pop.17             0             0
6500wg ?
r i150.wg.pop.48/28  192.168.wg.33             0    100      0
6500wg ?
* i150.wg.pop.16/28  192.168.wg.33             0    100      0
6500wg ?
Route Distinguisher: wg:20 (default for vrf Customer_B)
* i10.0.0.0         192.168.wg.33             0    100      0
6500wg ?
*>                  150.wg.pop.33             0             0
6500wg ?
*>i10.2.pop.16/28   192.168.wg.33             0    100      0
6500wg ?
* i150.wg.0.0       192.168.wg.33             0    100      0
6500wg ?
*>                  150.wg.pop.33             0             0
6500wg ?
*>i150.wg.pop.32/28 192.168.wg.33             0    100      0
6500wg ?

PEpop#sh ip bgp vpnv4 all
BGP table version is 74, local router ID is 192.168.wg.33
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
```



```
Route Distinguisher: wg:10 (default for vrf Customer_A)
*>i10.1.41.16/28    192.168.4.17             0    100      0
65041 ?
*>i10.1.41.49/32    192.168.4.17             0    100      0
65041 ?
*> 10.1.42.16/28    150.4.42.17              0             0
65041 ?
*> 10.1.42.49/32    150.4.42.17              0             0
65041 ?
*>i150.4.41.16/28   192.168.4.17             0    100      0
65041 ?
r> 150.4.42.16/28   150.4.42.17              0             0
65041 ?
Route Distinguisher: 4:20 (default for vrf Customer_B)
*>i10.2.41.16/28    192.168.4.17             0    100      0
65042 ?
*>i10.2.41.49/32    192.168.4.17             0    100      0
65042 ?
*> 10.2.42.16/28    150.4.42.33              0             0
65042 ?
*> 10.2.42.49/32    150.4.42.33              0             0
65042 ?
*>i150.4.41.32/28   192.168.4.17             0    100      0
65042 ?
r> 150.4.42.32/28   150.4.42.33              0             0
65042 ?
```

Task 2: Configuring the Backup PE-CE Link

In this task you will enable the backup links on the PE routers. Team A will establish the link between its CE*wg*2A router and the PE*wg*1 router, and team B will establish the link between its CE*wg*1B router and the PE*wg*2 router. Ensure that the interface is added to the proper VRF and BGP is activated.

Exercise Procedure

Complete these steps:

Step 1 Configure an additional subinterface on the existing serial interfaces on your PE and CE routers.

Step 2 Add the backup link to the appropriate VRF.

Which VRF is CE*wg*1B added to?

_____

Which VRF is CE*wg*2A added to?

_____

Step 3 Configure IP addresses and DLCIs on this interface using the parameters in the table here, where x = PE*wg*2 POP number and y = PE*wg*1 POP number.

Backup Link Configuration Parameters

Source Router	IP Address	DLCI	Destination Router	IP Address	DLCI
CE*wg*1B	150.*wg.x.49/28*	113	PE*wg*2	150.wg.x.50/28	113
CE*wg*2A	150.*wg.y.49/28*	113	PE*wg*1	150.wg.y.50/28	113

Step 4 Activate the BGP neighbor relationship between your CE router and the appropriate PE router.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify point-to-point connectivity over the new subinterface.

```
CEpopB#ping 150.wg.pop.50

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 150.wg.pop.50, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/28/32 ms


PEwg2#ping vrf Customer_B 10.2.41.49

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 150.wg.pop.49, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/28/32 ms


CEpopA#ping 150.wg.pop.50

 Sending 5, 100-byte ICMP Echos to 150.wg.pop.50, timeout is 2
seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
28/29/32 ms



Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 150.wg.pop.49, timeout is 2
seconds:

!!!!!
```

- Check BGP connectivity with the **show ip bgp summary** command on the CE routers.

```
CEpopA#sh ip bgp sum

BGP router identifier 10.1.pop.49, local AS number 650pop

BGP table version is 10, main routing table version 10

9 network entries and 9 paths using 1197 bytes of memory

2 BGP path attribute entries using 120 bytes of memory

1 BGP AS-PATH entries using 24 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory
```

```
BGP activity 9/30 prefixes, 9/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.50     4 65001     606     607       10    0    0
10:01:29        2
150.wg.pop.18     4 65001     617     618       10    0    0
09:50:35        3

CEpopA#sh ip bgp
BGP table version is 63, local router ID is 10.1.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
*  10.1.pop.16/28    150.wg.pop.18
0 65001 650pop ?
*>                   150.wg.pop.50                          0
65001 650pop ?
*  10.1.pop.49/32    150.wg.pop.18
0 65001 650pop ?
*>                   150.wg.pop.50                          0
65001 650pop ?
*> 10.1.pop.16/28    0.0.0.0                  0         32768
?
*> 10.1.pop.49/32    0.0.0.0                  0         32768
?
*  150.wg.pop.16/28   150.wg.pop.18
0 65001 650pop ?
*>                   150.wg.pop.50                          0
65001 650pop ?
*> 150.wg.pop.48/28   0.0.0.0                  0         32768
?
*> 150.wg.pop.16/28   0.0.0.0                  0         32768
?

CEpopA#sh ip bgp sum
BGP router identifier 10.1.pop.49, local AS number 650pop
BGP table version is 22, main routing table version 22
7 network entries and 7 paths using 931 bytes of memory
2 BGP path attribute entries using 120 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
```



```
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 14/33 prefixes, 14/7 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.18     4 65001     638     637       22    0    0
10:10:56        4
```

```
CEpopA#sh ip bgp
BGP table version is 38, local router ID is 10.1.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
*> 10.1.pop.16/28     0.0.0.0                  0         32768
?
*> 10.1.pop.49/32     0.0.0.0                  0         32768
?
*> 10.1.pop.16/28     150.wg.pop.18
0 65001 650pop ?
*> 10.1.pop.49/32     150.wg.pop.18
0 65001 650pop ?
*> 150.wg.pop.16/28   0.0.0.0                  0         32768
?
*> 150.wg.pop.48/28   150.wg.pop.18
0 65001 650pop ?
*> 150.wg.pop.16/28   150.wg.pop.18
0 65001 650pop ?
```

```
PEpop#sh ip bgp vpn all
BGP table version is 145, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf
Weight Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*> 10.1.pop.16/28     150.wg.pop.17            0
0 650pop ?
*> 10.1.pop.49/32     150.wg.pop.17            0
0 650pop ?
```

```
* i10.1.pop.16/28     192.168.wg.33             0    100
0 650pop ?

*>                    150.wg.pop.49             0
0 650pop ?

* i10.1.pop.49/32     192.168.wg.33             0    100
0 650pop ?

*>                    150.wg.pop.49             0
0 650pop ?

r> 150.wg.pop.16/28   150.wg.pop.17             0
0 650pop ?

r i150.wg.pop.48/28   192.168.wg.33             0    100
0 650pop ?

r>                    150.wg.pop.49             0
0 650pop ?

* i150.wg.pop.16/28   192.168.wg.33             0    100
0 650pop ?

*>                    150.wg.pop.49             0
0 650pop ?

Route Distinguisher: wg:20 (default for vrf Customer_B)

* i10.0.0.0           192.168.wg.33             0    100
0 650pop ?

*>                    150.wg.pop.33             0
0 650pop ?

*>i10.2.pop.16/28     192.168.wg.33             0    100
0 650pop ?

*>i10.2.pop.49/32     192.168.wg.33             0    100
0 650pop ?

* i150.wg.0.0         192.168.wg.33             0    100
0 650pop ?

*>                    150.wg.pop.33             0
0 650pop ?

*>i150.wg.pop.32/28   192.168.wg.33             0    100
0 650pop ?


PEpop#sh ip bgp vpnv4 all

BGP table version is 74, local router ID is 192.168.wg.33

Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,

              r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete


   Network          Next Hop            Metric LocPrf Weight
Path

Route Distinguisher: wg:10 (default for vrf Customer_A)

*>i10.1.pop.16/28     192.168.wg.17             0    100
0 650pop ?
```



```
*>i10.1.pop.49/32     192.168.wg.17             0    100
0 650pop ?

*> 10.1.pop.16/28     150.wg.pop.17             0
0 650pop ?

* i                   192.168.wg.17             0    100
0 650pop ?

*> 10.1.pop.49/32     150.wg.pop.17             0
0 650pop ?

* i                   192.168.wg.17             0    100
0 650pop ?

*>i150.wg.pop.16/28   192.168.wg.17             0    100
0 650pop ?

*> 150.wg.pop.48/28   150.wg.pop.17             0
0 650pop ?

* i                   192.168.wg.17             0    100
0 650pop ?

r> 150.wg.pop.16/28   150.wg.pop.17             0
0 650pop ?

r i                   192.168.wg.17             0    100
0 650pop ?

Route Distinguisher: wg:20 (default for vrf Customer_B)

*> 10.0.0.0           150.wg.pop.49             0
0 650pop ?

* i                   192.168.wg.17             0    100
0 650pop ?

*> 10.2.pop.16/28     150.wg.pop.33             0
0 650pop ?

*> 10.2.pop.49/32     150.wg.pop.33             0
0 650pop ?

*> 150.wg.0.0         150.wg.pop.49             0
0 650pop ?

* i                   192.168.wg.17             0    100
0 650pop ?

r> 150.wg.pop.32/28   150.wg.pop.33             0
0 650pop ?
```

Task 3: Selecting the Primary and Backup Link with BGP

It may be necessary to control the BGP selection of the link to establish a primary-backup relationship. In this task, you will use the local preference and MED attributes to control link selection. In this implementation, the new link bypasses the MPLS core. However, because it a high-cost link, it should be considered only as the backup link; the link through the MPLS core is to be used as the primary link.

Step 1 Use BGP local preference on the CE router to select the link to its local PE router (through the MPLS core) as the primary link and the link to the remote PE router (bypass link) as the backup link.

Step 2 Set the MED in outgoing routing updates from your CE router to make sure that the PE routers prefer the link through the MPLS core before using the backup link.

Exercise Verification

You have completed this exercise when you attain these results:

- You may have to issue a **clear ip route** or **clear ip bgp *** on the CE router to propagate routes with the new parameters.
- Verify that the primary link (the link to your local PE router) is being used. Use the **show ip bgp** command to verify. Make sure that the routes received from the primary link are always selected as the best routes.

```
CEpopB#sh ip bgp
BGP table version is 15, local router ID is 10.2.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf
Weight Path
*> 10.0.0.0         0.0.0.0                  0
32768 ?
*> 10.2.pop.16/28   150.wg.pop.34
0 65001 6500wg ?
*                   150.wg.pop.50                 50
0 65001 6500wg ?
*> 10.2.pop.49/32   150.wg.pop.34
0 65001 6500wg ?
*                   150.wg.pop.50                 50
0 65001 6500wg ?
*> 150.wg.0.0       0.0.0.0                  0
32768 ?
*> 150.wg.pop.32/28 150.wg.pop.34
0 65001 6500wg ?
*                   150.wg.pop.50                 50
0 65001 6500wg ?
```


- Verify the proper setting of the MED by using the **show ip bgp vpnv4 vrf** command on the PE routers. Make sure that the PE routers select routes coming from the primary link as the best routes.

```
PEpop#sh ip bgp vpnv4 all
BGP table version is 84, local router ID is 192.168.wg.33
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete


   Network            Next Hop               Metric LocPrf
Weight Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*>i10.1.pop.16/28     192.168.wg.17               0    100
0 6500wg ?
*>i10.1.pop.49/32     192.168.wg.17               0    100
0 6500wg ?
*> 10.1.pop.16/28     150.wg.pop.17               0
0 6500wg ?
* i                   192.168.wg.17               0    100
0 6500wg ?
*> 10.1.pop.49/32     150.wg.pop.17               0
0 6500wg ?
* i                   192.168.wg.17               0    100
0 6500wg ?
*>i150.wg.pop.16/28   192.168.wg.17               0    100
0 6500wg ?
*> 150.wg.pop.48/28   150.wg.pop.17               0
0 6500wg ?
* i                   192.168.wg.17               0    100
0 6500wg ?
r> 150.wg.pop.16/28   150.wg.pop.17               0
0 6500wg ?
r i                   192.168.wg.17               0    100
0 6500wg ?
Route Distinguisher: wg:20 (default for vrf Customer_B)
*>i10.0.0.0           192.168.wg.17               0    100
0 6500wg ?
*                     150.wg.pop.49             200
0 6500wg ?
*> 10.2.pop.16/28     150.wg.pop.33               0
0 6500wg ?
*> 10.2.pop.49/32     150.wg.pop.33               0
0 6500wg ?
*>i150.wg.0.0         192.168.wg.17               0    100
0 6500wg ?
```

```
*                    150.wg.pop.49            200
0 6500wg ?
r> 150.wg.pop.32/28  150.wg.pop.33              0
0 6500wg ?
```

- Shut down the link from the local PE router to the CE router (either CE*wg*B1 or CE*wg*A1).
- Verify that the backup link (the link to your local PE router) is being used. Use the **show ip bgp** command to verify.

```
CEpopB#sh ip bgp
BGP table version is 18, local router ID is 10.2.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
*> 10.0.0.0         0.0.0.0                  0         32768
?
*> 10.2.pop.16/28   150.wg.pop.50                     50
0 65001 6500wg ?
*> 10.2.pop.49/32   150.wg.pop.50                     50
0 65001 6500wg ?
*> 150.wg.0.0       0.0.0.0                  0         32768
?
*> 150.wg.pop.32/28 150.wg.pop.50                     50
0 65001 6500wg ?
```

- Re-enable the subinterface.
- After the BGP session is established with the local PE router, verify that the local link is shown as the preferred link for traffic. Use the **show ip bgp** command to verify.

```
CEpopB#sh ip bgp
BGP table version is 21, local router ID is 10.2.pop.49
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop              Metric LocPrf
Weight Path
*> 10.0.0.0         0.0.0.0                  0         32768
?
*> 10.2.pop.16/28   150.wg.pop.34
0 65001 6500wg ?
*                   150.wg.pop.50                     50
0 65001 6500wg ?
*> 10.2.pop.49/32   150.wg.pop.34
0 65001 6500wg ?
```



```
*                    150.wg.pop.50                    50
0 65001 6500wg ?
*> 150.wg.0.0        0.0.0.0                  0         32768
?
*> 150.wg.pop.32/28  150.wg.pop.34
0 65001 6500wg ?
*                    150.wg.pop.50                    50
0 65001 6500wg ?
```

Next Step

- **Module 6: Complex** MPLS VPNs

Lab Exercise 6-1: Overlapping VPNs

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

Your VPN customers want to exchange data between their central sites. You have decided to implement this request with an overlapping VPN topology.

In this exercise, you will establish overlapping VPNs to support the needs of your customers. After completing this exercise, you will have met these objectives:

- Design a VPN solution
- Remove CE*wg*1A and CE*wg*2B from existing VRFs
- Configure new VRFs for CE*wg*1A and CE*wg*2B

Visual Objective

The figure illustrates what you will accomplish in this exercise.

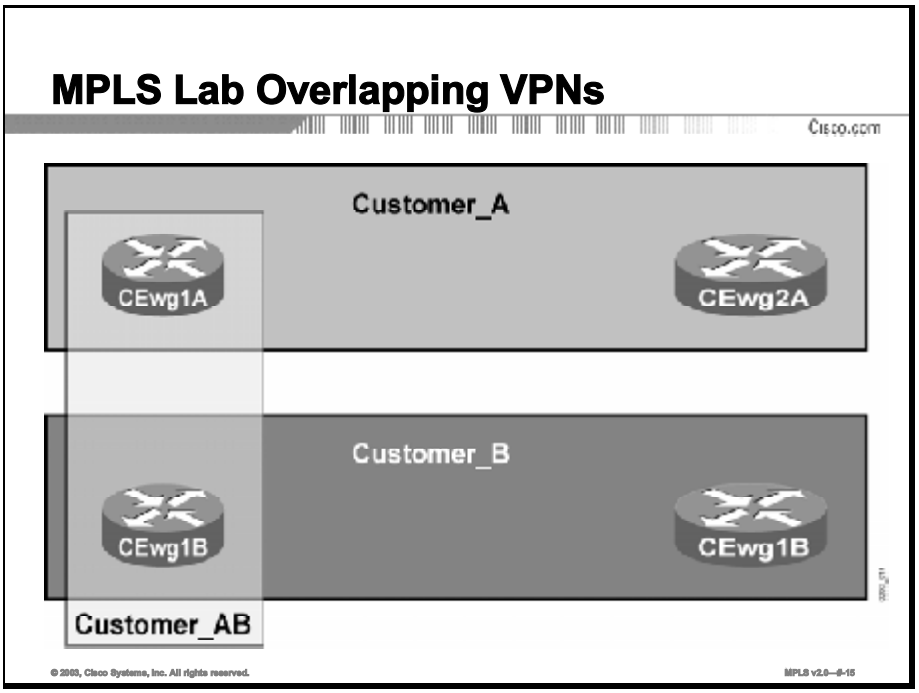

In this laboratory exercise, you will establish overlapping VPNs with the following connectivity goals:

- CE*wg*1A and CE*wg*1B cannot communicate.
- CE*wg*1A and CE*wg*2A can communicate.
- CE*wg*1A and CE*wg*2B can communicate.
- CE*wg*1B and CE*wg*2A cannot communicate.
- CE*wg*1B and CE*wg*2B can communicate.
- CE*wg*2A and CE*wg*2B cannot communicate.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise have been used in previous exercises.

Task 1: Designing Your VPN Solution

Site CE*wg*1A cannot belong to the same VRF as the other *wg*A sites. Similarly, site CE*wg*2B cannot belong to the same VRF as the *wg*B sites. Also, CE*wg*1A and CE*wg*2B cannot share the same VRF.

Exercise Procedure

Complete these steps:

Step 1 Allocate new RDs for VRFs to which CE*wg*1A and CE*wg*2B will be connected.

Step 2 A new RT is needed for the Customer_AB VPN. Coordinate the value of this RT with the other team within your workgroup.

Note You could use *wg*:11 as the RD for VRFs connected to CE*wg*1A and *wg*:21 as the RD for VRFs connected to CE*wg*2B. You could use *wg*:1001 as the RT for the Customer_AB VPN.

Exercise Verification

You have completed this exercise when you attain these results:

- Establish RDs and RTs for the new VRFs

Task 2: Removing CE*wg*1A and CE*wg*2B from Existing VRFs

CE*wg*1A and CE*wg*2B must be migrated to new routing contexts. It is tempting to do this by merely changing the RDs and RTs of their existing VRF. However, this approach is not possible because the other VPN site, connected to the same PE router, is sharing those VRFs.

Note When you enabled the backup link, you connected both CE*wg*1A and CE*wg*2A to PE*wg*1. So if you change the routing context of customer A on PE*wg*1, you will affect both CE*wg*1A and CE*wg*2A. This situation also holds true for CE*wg*1B, CE*wg*1B, and PE*wg*2.

Sites CE*wg*1A and CE*wg*2B have to be migrated to new VRFs. All the references to them must be removed from the existing routing protocol contexts.

In this task, you will remove the references to CE*wg*1A and CE*wg*2B.

Exercise Procedure

Complete these steps:

Step 1 Remove the address family BGP neighbor relationship between CE*wg*1A and CE*wg*2B on their PE router.

Step 2 Check any other references to CE*wg*1A and CE*wg*2B in their PE router configuration and, if required, remove them.

Exercise Verification

You have completed this exercise when you attain these results:

- On the PE router, verify that the interface toward the CE router is no longer in the original VRF by using the **show ip vrf interfaces** command. This action should result in a printout similar to the one here:

```
PEwg1#sh ip vrf int
Interface              IP-Address      VRF
Protocol
Interface              IP-Address      VRF
Protocol
Serial0/0.113          150.wg.41.50     Customer_A
up
Serial0/0.102          150.wg.41.34     Customer_B
up


PEwg2#sh ip vrf int
Interface              IP-Address      VRF
Protocol
Serial0/0.101          150.wg.42.18     Customer_A
up
Serial0/0.113          150.wg.42.50     Customer_B
up
```


- Verify the BGP neighbor relationship has been removed on the PE router with the **show ip bgp vpnv4 vrf summary** command. This action should give you a printout similar to the one here. Check the status of CE*wg*1A and CE*wg*2B in the printout.

```
PEpop#sh ip bgp vpnv4 vrf Customer_A sum
BGP router identifier 192.168.wg.17, local AS number 65001
BGP table version is 34, main routing table version 34
7 network entries using 847 bytes of memory
11 path entries using 704 bytes of memory
7 BGP path attribute entries using pop0 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2139 total bytes of memory
BGP activity 51/29 prefixes, 69/43 paths, scan interval 15
secs
 Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.49     4 6500wg     976     979       34    0    0
00:29:12        4

PEpop#sh ip bgp vpnv4 vrf Customer_B sum
BGP router identifier 192.168.wg.33, local AS number 65001
BGP table version is 33, main routing table version 33
5 network entries using 605 bytes of memory
7 path entries using 448 bytes of memory
7 BGP path attribute entries using pop0 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 16pop total bytes of memory
BGP activity 122/102 prefixes, 160/138 paths, scan interval 15
secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.49     4 6500wg    1477    1479       33    0    0
00:30:26        2
```

Task 3: Configuring New VRFs for CE*wg*1A and CE*wg*2B

In this task, you will create the new VRFs for CE*wg*1A and CE*wg*2B.

Exercise Procedure

Complete these steps:

Step 1 Create the new VRFs for CE*wg*1A and CE*wg*2B on their PE router with the **ip vrf** command.

Step 2 Assign new RDs to the newly created VRFs with the **rd** command.

Step 3 Assign proper import and export RTs to the newly created VRFs with the **route-target** command.

Step 4 Re-establish BGP routing between the PE routers and the CE routers. Please refer to the "Running BGP Between PE and CE Routers" lab exercise if you need more details.

Exercise Verification

You have completed this exercise when you attain these results:

- On the PE router, verify that the interface toward the CE router is in the proper VRF by using the **show ip vrf interfaces** command. This action should result in a printout similar to the one here:

```
PEpop#sh ip vrf int
Interface              IP-Address      VRF
Protocol
Serial0/0.113          150.wg.pop.50     Customer_A         up
Serial0/0.101          150.wg.pop.18     Customer_AB        up
Serial0/0.102          150.wg.pop.34     Customer_B         up

PEpop#sh ip vrf int
Interface              IP-Address      VRF
Protocol
Serial0/0.101          150.wg.pop.18     Customer_A         up
Serial0/0.102          150.wg.pop.34     Customer_AB        up
Serial0/0.113          150.wg.pop.50     Customer_B         up
```

- Verify the BGP neighbors on the PE router with the **show ip bgp vpnv4 vrf summary** command. This should give you a printout similar to the one here. Check the status of CE*wg*1A and CE*wg*2B in the printout.

```
PEpop#sh ip bgp vpnv4 vrf Customer_AB sum
BGP router identifier 192.168.wg.17, local AS number 65001
BGP table version is 49, main routing table version 49
10 network entries using 1210 bytes of memory
10 path entries using 640 bytes of memory
```



```
7 BGP path attribute entries using pop0 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2438 total bytes of memory
BGP activity 57/35 prefixes, 75/49 paths, scan interval 15
secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.17     4 6500wg      53      54       49    0    0
00:48:pop        3
 PEpop#sh ip bgp vpnv4 vrf Customer_AB sum
BGP router identifier 192.168.wg.33, local AS number 65001
BGP table version is 56, main routing table version 56
8 network entries using 968 bytes of memory
8 path entries using 512 bytes of memory
7 BGP path attribute entries using pop0 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2068 total bytes of memory
BGP activity 130/110 prefixes, 168/146 paths, scan interval 15
secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
150.wg.pop.33     4 6500wg       9      10       56    0    0
00:04:17        3
```

- Check the BGP routing table in the new VRF with the **show ip bgp vpnv4 vrf** command. You should see routes from CE*wg*1A or CE*wg*2B as well as routes imported from other VRFs. Use the AS path to work out which routes belong to which CE router. Routes announced by CE*wg*1A should have 650x1 in the AS path, and routes announced by CE*wg*2B should have 650x2 in the AS path.

```
PEpop#sh ip bgp vpnv4 vrf Customer_AB
BGP table version is 49, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
```

```
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network            Next Hop                Metric LocPrf
Weight Path
Route Distinguisher: pop:1 (default for vrf Customer_AB)
*> 10.1.pop.16/28     150.wg.pop.17                0
0 6500wg ?
*> 10.1.pop.49/32     150.wg.pop.17                0
0 6500wg ?
*>i10.1.pop.16/28     192.168.wg.33                0    100
0 6500wg ?
*>i10.1.pop.49/32     192.168.wg.33                0    100
0 6500wg ?
*>i10.2.pop.16/28     192.168.wg.33                0    100
0 6500wg ?
*>i10.2.pop.49/32     192.168.wg.33                0    100
0 6500wg ?
r> 150.wg.pop.16/28   150.wg.pop.17                0
0 6500wg ?
*>i150.wg.pop.48/28   192.168.wg.33                0    100
0 6500wg ?
*>i150.wg.pop.16/28   192.168.wg.33                0    100
0 6500wg ?
*>i150.wg.pop.32/28   192.168.wg.33                0    100
0 6500wg ?

PEpop#sh ip bgp vpnv4 vrf Customer_AB
BGP table version is 56, local router ID is 192.168.wg.33
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network            Next Hop                Metric LocPrf
Weight Path
Route Distinguisher: pop:1 (default for vrf Customer_AB)
*>i10.0.0.0           192.168.wg.17                0    100
0 6500wg ?
*>i10.1.pop.16/28     192.168.wg.17                0    100
0 6500wg ?
*>i10.1.pop.49/32     192.168.wg.17                0    100
0 6500wg ?
*> 10.2.pop.16/28     150.wg.pop.33                0
0 6500wg ?
```



```
*> 10.2.pop.49/32     150.wg.pop.33            0
0 6500wg ?
*>i150.wg.0.0         192.168.wg.17            0    100
0 6500wg ?
*>i150.wg.pop.16/28   192.168.wg.17            0    100
0 6500wg ?
r> 150.wg.pop.32/28   150.wg.pop.33            0
0 6500wg ?
```

- Use the **show ip bgp vpnv4 vrf** *name prefix* command to display details of an individual route and verify that the proper RTs are attached to the route. Your printout should be similar to the one here:

```
PEpop#sh ip bgp vpnv4 vrf Customer_AB 10.2.pop.49
BGP routing table entry for pop:1:10.2.pop.49/32, version 53
Paths: (1 available, best #1, table Customer_AB)
  Advertised to non peer-group peers:
  192.168.100.129
  6500wg
    150.wg.pop.33 from 150.wg.pop.33 (10.2.pop.49)
      Origin incomplete, metric 0, localpref 100, valid,
external, best
Extended Community: RT:pop:1 RT:wg:20

PEpop#sh ip bgp vpnv4 vrf Customer_AB 10.1.pop.49
BGP routing table entry for pop:1:10.1.pop.16/28, version 23
Paths: (1 available, best #1, table Customer_AB)
  Advertised to non peer-group peers:
  192.168.100.129
  6500wg
    150.wg.pop.17 from 150.wg.pop.17 (10.1.pop.49)
      Origin incomplete, metric 0, localpref 100, valid,
external, best
Extended Community: RT:pop:1 RT:wg:10
```

- Connect to CE*wg*A1 and perform **ping** and **trace** to the loopback address of CE*wg*2B (or vice versa). The other router should be reachable. For subgroup B, perform the test in the other direction.

```
CEwg1A#ping 10.1.wg2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.wg2.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
148/148/149 ms
```

```
CEpopAB#trace 10.1.pop.49
Type escape sequence to abort.
Tracing the route to 10.1.pop.49
  1 150.wg.pop.34 16 msec 16 msec 12 msec
2 150.wg.pop.17 [AS 6500wg] 72 msec 76 msec *


CEpopA#ping 10.2.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
148/148/149 ms



CEpopA#trace 10.2.pop.49
Type escape sequence to abort.
Tracing the route to 10.2.pop.49
  1 150.wg.pop.18 16 msec 16 msec 12 msec
  2 150.wg.pop.33 [AS 6500wg] 72 msec 77 msec *
```

- Connect to CE*wg*A2 and try to **ping** CE*wg*2B or CE*wg*1B. Those routers should not be reachable from CE*wg*A2. For subgroup B, **ping** CE*wg*1A and CE*wg*2A from CE*wg*1B.

```
CEpopA#ping 10.2.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)


CEpopA#ping 10.2.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.pop.49, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)


CEpopB#ping 10.1.pop.49
Type escape sequence to abort.
```



```
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)

CEpopB#ping 10.1.pop.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.pop.49, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)
```

Lab Exercise 6-2: Merging Service Providers

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

Your small service provider is merging with several other small service providers. To accomplish this consolidation, a central P router (P1) has been installed and configured, and Frame Relay connectivity has been provided from each local P router to P1. In addition, the core IGP is being converted from EIGRP to Intermediate System-to-Intermediate System (IS-IS).

In this exercise, merge your small service provider with several other small service providers. After completing this exercise, you will be able to meet these objectives:

- Convert the core IGP from EIGRP to IS-IS
- Enable MPLS LDP connectivity with the central P router
- Enable IBGP connectivity between all PE routers

Visual Objective

Team A will configure PE*wg*1/Pwg1, and team B will configure PE*wg*2/P*wg*2. P1 has been preconfigured.

The figure illustrates what you will accomplish in this exercise.

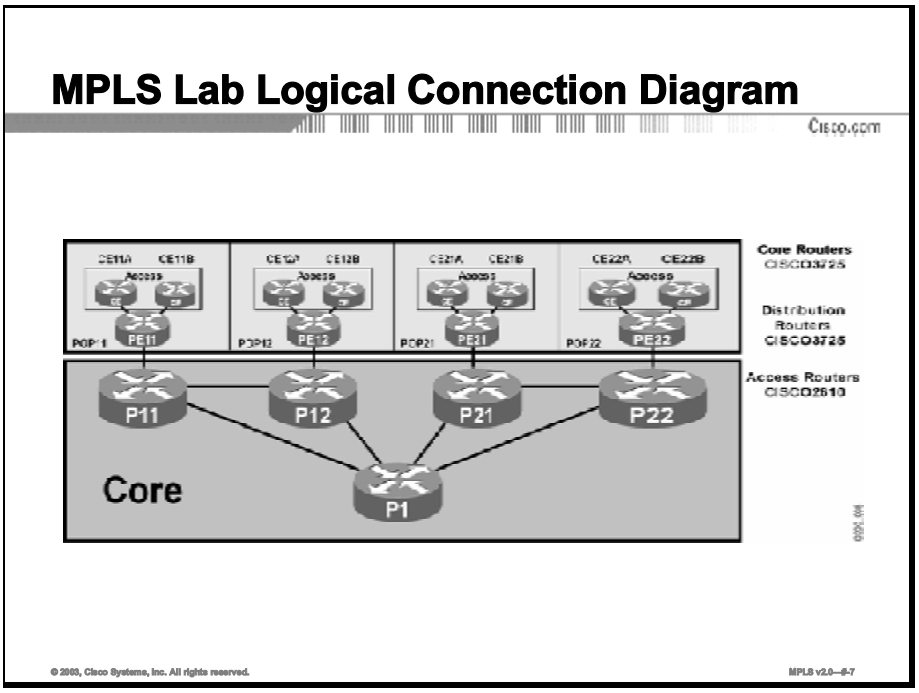

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Commands for Merging Service Providers

Command	Description
router isis area-tag	To enable the IS-IS routing protocol and to specify an IS-IS process, use the **router isis** command in global configuration mode. To disable IS-IS routing, use the **no** form of this command.
net network-entity-title	To configure an IS-IS network entity title (NET) for a Connectionless Network Service (CLNS) routing process, use the **net** command in router configuration mode. To remove a NET, use the **no** form of this command.
isis circuit-type {level-1 | level-1-2 | level-2-only}	To configure the type of adjacency, use the **isis circuit-type** interface configuration command. To reset the circuit type to Level I and Level 2, use the **no** form of this command.
metric-style wide [transition] [level-1 | level-2 | level-1-2]	To configure a router running IS-IS so that it generates and accepts only new-style type, length, and value objects (TLVs), use the **metric-style wide** command in router configuration mode. To disable this function, use the **no** form of this command.

Task 1: Enabling Connectivity with the Central P Router

In this task you will enable the Frame Relay link between your P routers and P1, and then enable LDP connectivity between the two routers.

Exercise Procedure

Complete these steps:

Step 1 Configure IP addresses and DLCIs on this interface using the parameters in the table here.

IP Address and DLCI Configuration Parameters

Router	Subinterface	DLCI	IP Address
P11	S0/0.211	211	192.168.100.10/29
P12	S0/0.212	212	192.168.100.18/29
P21	S0/0.221	221	192.168.100.26/29
P22	S0/0.222	222	192.168.100.34/29

Exercise Verification

You have completed this exercise when you attain these results:

- On your P router, use the **show interface** command to verify that the new interfaces are operational.

Task 2: Migrating the Core to IS-IS

Because a link-state protocol is more scalable than a distance vector protocol, the service provider has decided to migrate the core to IS-IS. The P1 router has already been migrated. Your workgroup is responsible for the migration of all of your assigned routers. Team A will migrate PE*wg*1 and P*wg*1. Team B will migrate PE*wg*2 and P*wg*2.

Exercise Procedure

Complete these steps:

Step 1 Disable EIGRP as the core IGP on your assigned routers.

Step 2 Enable IS-IS as the core IGP using the parameters detailed in the following table.

IS-IS Parameters

Router ID	NET	Remarks
PE*wg*1	net 49.0001.0000.0000.01xx.00	Where xx = the POP number
PE*wg*2	net 49.0001.0000.0000.01xx.00	Where xx = the POP number
P*wg*1	net 49.0001.0000.0000.02xx.00	Where xx = the POP number
P*wg*2	net 49.0001.0000.0000.02xx.00	Where xx = the POP number

Note Ensure that the **metric-style** is set to *wide*, the **is-type** is set to *level-2*-only, and IS-IS has been enabled on the active serial interfaces that are supporting the core MPLS.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **show ip protocol** command to verify that IS-IS is active and enabled on all appropriate interfaces.

```
PEpop#sh ip prot
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    Serial0/0.111
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
```

```
    192.168.1.97         115      00:05:48
    192.168.3.97         115      00:05:48
    192.168.2.97         115      00:05:48
    192.168.4.97         115      00:00:22
    192.168.4.113        115      00:06:31
    192.168.100.58       115      00:00:22
    192.168.1.81         115      00:05:48
    192.168.3.81         115      00:05:48
    192.168.2.81         115      00:05:48
    192.168.1.33         115      00:05:48
    192.168.3.33         115      00:05:48
    192.168.2.33         115      00:05:48
    192.168.4.33         115      00:00:22
    192.168.100.66       115      00:00:28
    192.168.1.17         115      00:05:48
    192.168.3.17         115      00:05:48
    192.168.2.17         115      00:05:48
    192.168.100.129      115      00:05:53    Distance:
(default is 115)

Ppop#sh ip prot
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    Serial0/0.111
    Serial0/0.112
    Serial0/0.2pop
  Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.1.97         115      00:10:18
    192.168.3.97         115      00:10:18
    192.168.2.97         115      00:10:18
    192.168.4.97         115      00:04:53
    192.168.1.81         115      00:10:18
```



```
    192.168.3.81          115       00:10:18
    192.168.2.81          115       00:10:18
    192.168.1.33          115       00:10:18
    192.168.3.33          115       00:10:20
    192.168.2.33          115       00:10:20
    192.168.4.33          115       00:04:55
    192.168.100.66        115       00:05:00
    192.168.1.17          115       00:10:20
    192.168.3.17          115       00:10:20
    192.168.2.17          115       00:10:20
    192.168.4.17          115       00:04:55
    192.168.100.129       115       00:04:55    Distance:
(default is 115)
```

- Use the **show ip route** command and verify that all routers are sending and receiving the appropriate prefixes.

```
PEwg1#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

********** output omitted
**************************************
     192.168.wg.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.wg.97/32 [115/20] via 192.168.wg.66,
Serial0/0.111
i L2    192.168.wg.112/28 [115/20] via 192.168.wg.66,
Serial0/0.111
C       192.168.wg.64/28 is directly connected, Serial0/0.111
i L2    192.168.wg.81/32 [115/30] via 192.168.wg.66,
Serial0/0.111
C       192.168.wg.33/32 is directly connected, Loopback0
```

```
i L2    192.168.wg.48/28 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.wg.17/32 [115/40] via 192.168.wg.66,
Serial0/0.111

     192.168.100.0/24 is variably subnetted, 8 subnets, 2
masks

i L2    192.168.100.24/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.16/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.40/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.32/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.56/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.48/29 [115/30] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.64/29 [115/20] via 192.168.wg.66,
Serial0/0.111

i L2    192.168.100.129/32 [115/30] via 192.168.wg.66,
Serial0/0.111

**************************output omitted
***************************



PE wg2#sh ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area

       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2

       E1 - OSPF external type 1, E2 - OSPF external type 2

       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area

       * - candidate default, U - per-user static route, o -
ODR

       P - periodic downloaded static route


Gateway of last resort is not set


            ***************output omitted
**************************

     192.168.wg.0/24 is variably subnetted, 7 subnets, 2 masks

i L2    192.168.wg.97/32 [115/30] via 192.168.wg.50,
Serial0/0.111
```



```
i L2    192.168.wg.112/28 [115/20] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.wg.64/28 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.wg.81/32 [115/20] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.wg.33/32 [115/40] via 192.168.wg.50,
Serial0/0.111

C       192.168.wg.48/28 is directly connected, Serial0/0.111

C       192.168.wg.17/32 is directly connected, Loopback0

     192.168.100.0/24 is variably subnetted, 8 subnets, 2
masks

i L2    192.168.100.24/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.16/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.40/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.32/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.56/29 [115/20] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.48/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.64/29 [115/30] via 192.168.wg.50,
Serial0/0.111

i L2    192.168.100.129/32 [115/30] via 192.168.wg.50,
Serial0/0.111

**********output omitted **************************

  Pwg1#sh ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area

       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2

       E1 - OSPF external type 1, E2 - OSPF external type 2

       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area

       * - candidate default, U - per-user static route, o -
ODR

       P - periodic downloaded static route


Gateway of last resort is not set


**********output omitted **************************

    192.168.wg.0/24 is variably subnetted, 7 subnets, 2 masks
```

```
C       192.168.wg.97/32 is directly connected, Loopback0
C       192.168.wg.112/28 is directly connected, Serial0/0.112
C       192.168.wg.64/28 is directly connected, Serial0/0.111
i L2    192.168.wg.81/32 [115/20] via 192.168.wg.113,
Serial0/0.112
i L2    192.168.wg.33/32 [115/20] via 192.168.wg.65,
Serial0/0.111
i L2    192.168.wg.48/28 [115/20] via 192.168.wg.113,
Serial0/0.112
i L2    192.168.wg.17/32 [115/30] via 192.168.wg.113,
Serial0/0.112
     192.168.100.0/24 is variably subnetted, 8 subnets, 2
masks
i L2    192.168.100.24/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
i L2    192.168.100.16/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
i L2    192.168.100.40/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
i L2    192.168.100.32/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
i L2    192.168.100.56/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
                          [115/20] via 192.168.wg.113,
Serial0/0.112
i L2    192.168.100.48/29 [115/20] via 192.168.100.65,
Serial0/0.2pop
C       192.168.100.64/29 is directly connected,
Serial0/0.2pop
i L2    192.168.100.129/32 [115/20] via 192.168.100.65,
Serial0/0.2pop
***************output omitted ************************


Pwg2#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
```



```
Gateway of last resort is not set


***************output omitted *************************

          192.168.wg.0/24 is variably subnetted, 7 subnets, 2
masks

 C       192.168.wg.112/28 is directly connected,
Serial0/0.112

i L2    192.168.wg.64/28 [115/20] via 192.168.wg.114,
Serial0/0.112

C       192.168.wg.81/32 is directly connected, Loopback0

i L2    192.168.wg.33/32 [115/30] via 192.168.wg.114,
Serial0/0.112

C       192.168.wg.48/28 is directly connected, Serial0/0.111

i L2    192.168.wg.17/32 [115/20] via 192.168.wg.49,
Serial0/0.111

     192.168.100.0/24 is variably subnetted, 8 subnets, 2
masks

i L2    192.168.100.24/29 [115/20] via 192.168.100.57,
Serial0/0.2pop

i L2    192.168.100.16/29 [115/20] via 192.168.100.57,
Serial0/0.2pop

i L2    192.168.100.40/29 [115/20] via 192.168.100.57,
Serial0/0.2pop

i L2    192.168.100.32/29 [115/20] via 192.168.100.57,
Serial0/0.2pop

C       192.168.100.56/29 is directly connected,
Serial0/0.2pop

i L2    192.168.100.48/29 [115/20] via 192.168.100.57,
Serial0/0.2pop

i L2    192.168.100.64/29 [115/20] via 192.168.wg.114,
Serial0/0.112

                          [115/20] via 192.168.100.57,
Serial0/0.2pop

i L2    192.168.100.129/32 [115/20] via 192.168.100.57,
Serial0/0.2pop

***************output omitted *************************
```

Task 3: Enabling MPLS LDP Connectivity with the Central P Router

In this task you will enable LDP connectivity between your routers and P1.

Exercise Procedure

Complete these steps:

Step 1 Enable LDP on the subinterface that you have created.

Exercise Verification

You have completed this exercise when you attain these results:

- On your P router, verify that an LDP neighbor relationship has been established between your P router and P1.

```
Pwg1#sh mpls ldp nei
    Peer LDP Ident: 192.168.wg.81:0; Local LDP Ident
192.168.wg.97:0
        TCP connection: 192.168.wg.81.646 -
192.168.wg.97.11003
        State: Oper; Msgs sent/rcvd: 6240/6252; Downstream
        Up time: 3d18h
        LDP discovery sources:
          Serial0/0.112, Src IP addr: 192.168.wg.113
        Addresses bound to peer LDP Ident:
          192.168.wg.81    192.168.wg.113   192.168.wg.50
192.168.1.58
    Peer LDP Ident: 192.168.wg.33:0; Local LDP Ident
192.168.wg.97:0
        TCP connection: 192.168.wg.33.646 -
192.168.wg.97.11008
        State: Oper; Msgs sent/rcvd: 1330/1330; Downstream
        Up time: 19:08:59
        LDP discovery sources:
          Serial0/0.111, Src IP addr: 192.168.wg.65
        Addresses bound to peer LDP Ident:
          192.168.wg.33    192.168.wg.65


Pwg2#sh mpls ldp nei
    Peer LDP Ident: 192.168.wg.97:0; Local LDP Ident
192.168.wg.81:0
        TCP connection: 192.168.wg.97.11003 -
192.168.wg.81.646
        State: Oper; Msgs sent/rcvd: 6253/62pop; Downstream
```



```
        Up time: 3d18h
        LDP discovery sources:
          Serial0/0.112, Src IP addr: 192.168.wg.114
        Addresses bound to peer LDP Ident:
          192.168.wg.97    192.168.wg.114   192.168.wg.66
192.168.1.66
    Peer LDP Ident: 192.168.wg.17:0; Local LDP Ident
192.168.wg.81:0
        TCP connection: 192.168.wg.17.646 -
192.168.wg.81.23205
        State: Oper; Msgs sent/rcvd: 1333/1335; Downstream
        Up time: 19:13:02
        LDP discovery sources:
          Serial0/0.111, Src IP addr: 192.168.wg.49
        Addresses bound to peer LDP Ident:
          192.168.wg.17    192.168.wg.49
```

- On your PE router, verify that labels are being received from the other workgroups.

```
PEwg1#sh mpls forwarding
Local  Outgoing    Prefix            Bytes tag  Outgoing
Next Hop
tag    tag or VC   or Tunnel Id       switched    interface
16     Pop tag     192.168.wg.81/32    0          Se0/0.111
point1point
17     Untagged    192.168.wg.112/28   0          Se0/0.111
point1point
18     18          192.168.wg.97/32    0          Se0/0.111
point1point
19     Untagged    192.168.wg.64/28    0          Se0/0.111
point1point
20     19          192.168.wg.33/32    0          Se0/0.111
point1point
21     Untagged    10.1.pop.16/28[V]   0          Se0/0.101
point1point
22     Untagged    10.1.pop.49/32[V]   1752       Se0/0.101
point1point
23     Aggregate   150.wg.pop.16/28[V] 7872
24     Untagged    10.0.0.0/8[V]       0          Se0/0.102
point1point
25     Untagged    150.wg.0.0/16[V]    0          Se0/0.102
point1point
30     Untagged    10.1.pop.16/28[V]   0          Se0/0.113
point1point
31     Untagged    10.1.pop.49/32[V]   0          Se0/0.113
point1point
```

```
32     Aggregate   150.wg.pop.48/28[V] 0
33     Untagged    150.wg.pop.16/28[V] 660        Se0/0.113
point1point

PEwg2#sh mpls for
Local  Outgoing    Prefix            Bytes tag  Outgoing
Next Hop
tag    tag or VC   or Tunnel Id      switched     interface
16     Untagged    192.168.wg.112/28   0          Se0/0.111
point1point
17     Pop tag     192.168.wg.97/32    0          Se0/0.111
point1point
18     18          192.168.wg.81/32    0          Se0/0.111
point1point
19     Untagged    192.168.wg.48/28    0          Se0/0.111
point1point
20     19          192.168.wg.17/32    0          Se0/0.111
point1point
25     Untagged    10.2.pop.16/28[V]   0          Se0/0.102
point1point
26     Untagged    10.2.pop.49/32[V]   20056      Se0/0.102
point1point
27     Aggregate   150.wg.pop.32/28[V] 0
28     Untagged    10.1.pop.16/28[V]   0          Se0/0.101
point1point
31     Untagged    10.1.pop.49/32[V]   0          Se0/0.101
point1point
32     Untagged    150.wg.pop.48/28[V] 0          Se0/0.101
point1point
33     Aggregate   150.wg.pop.16/28[V] 0
```


Task 4: Enabling IBGP Connectivity for All PE Routers

At this point, you have established LDP connectivity for all of the P routers in your new service provider environment, but you have not yet established BGP connectivity. You now need to establish IBGP connectivity for your PE routers.

There are two methods that you can implement. The first is to use the **bgp neighbor** command to add a neighbor relationship between each of the routers, but this approach would entail a substantial configuration effort.

The second method is to implement route reflectors. To this end, P1 has been configured as a BGP route reflector. However, to take advantage of this fact, you will need to remove the neighbor relationship between your two PE routers and make them clients of P1.

Note The loopback address for P1 is 192.168.100.129.

Team A will configure PE*wg*1, and team B will configure PE*wg*2.

Step 1 Remove the neighbor relationship between your PE router and the remote PE router in your workgroup.

Step 2 Activate your PE router as a client of P1.

Exercise Verification

You have completed this exercise when you attain these results:

- On your PE routers, check BGP connectivity to all workgroups with the **show ip bgp summary** and **show ip bgp neighbor** commands on CE routers.

```
PEpop#sh ip bgp sum
BGP router identifier 192.168.wg.33, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
192.168.100.129 4 65001      25      25        1    0    0
00:17:25        0
```

Note The following command output has been edited.

```
PEwg1#sh ip bgp nei
BGP neighbor is 150.wg.pop.17,  vrf Customer_A,  remote AS
6500wg, external link
  BGP version 4, remote router ID 10.1.pop.49
  BGP state = Established, up for 00:23:29
  Last read 00:00:29, hold time is 180, keepalive interval is
60 seconds
```

```
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Default minimum time between advertisement runs is 30
seconds


 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF Customer_A
  BGP table version 31, neighbor version 31
  Index 1, Offset 0, Mask 0x2
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               3          4 (Consumes 256
bytes)
    Prefixes Total:                 3          4
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          4
    Used as multipath:            n/a          0
                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    AS_PATH loop:                       n/a          3
    Bestpath from this peer:              8        n/a
    Total:                                8          3
  Number of NLRIs in the update sent: max 4, min 0


  Connections established 5; dropped 4
  Last reset 00:24:09, due to User reset
Connection state is ESTAB, I/O status: 1, unread input bytes:
0
Local host: 150.wg.pop.18, Local port: 179
Foreign host: 150.wg.pop.17, Foreign port: 11079


BGP neighbor is 150.wg.pop.33,  vrf Customer_B,  remote AS
6500wg, external link
  BGP version 4, remote router ID 10.2.pop.49
  BGP state = Established, up for 00:23:32
  Last read 00:00:32, hold time is 180, keepalive interval is
60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
```



```
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  4          4
    Notifications:          0          0
    Updates:              pop          4
    Keepalives:          1318       1318
    Route Refresh:          0          0
    Total:               1363       1326
  Default minimum time between advertisement runs is 30
seconds


 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF Customer_B
  BGP table version 31, neighbor version 31
  Index 2, Offset 0, Mask 0x4
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               2          3 (Consumes 192
bytes)
    Prefixes Total:                 8          3
    Implicit Withdraw:              6          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          3
    Used as multipath:            n/a          0
                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:              6        n/a
    Total:                                6          0
  Number of NLRIs in the update sent: max 2, min 0
   Connections established 4; dropped 3
  Last reset 00:24:15, due to User reset
Connection state is ESTAB, I/O status: 1, unread input bytes:
0
Local host: 150.wg.pop.34, Local port: 11074
Foreign host: 150.wg.pop.33, Foreign port: 179


BGP neighbor is 150.wg.pop.49,  vrf Customer_B,  remote AS
6500wg, external link
  BGP version 4, remote router ID 10.2.pop.49
```

```
  BGP state = Established, up for 00:23:39
  Last read 00:00:39, hold time is 180, keepalive interval is
60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:               17         14
    Keepalives:           381        382
    Route Refresh:          0          0
    Total:                401        399
  Default minimum time between advertisement runs is 30
seconds

 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF Customer_B
  BGP table version 31, neighbor version 31
  Index 4, Offset 0, Mask 0x10
                                Sent       Rcvd
  Prefix activity:              ----       ----
    Prefixes Current:              5          2 (Consumes 128
bytes)
    Prefixes Total:               10          2
    Implicit Withdraw:             5          0
    Explicit Withdraw:             0          0
    Used as bestpath:            n/a          0
    Used as multipath:           n/a          0
                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    AS_PATH loop:                      n/a          3
    Bestpath from this peer:             4        n/a
    Total:                               4          3
  Number of NLRIs in the update sent: max 3, min 0

  Connections established 3; dropped 2
  Last reset 00:24:22, due to User reset
```



```
Connection state is ESTAB, I/O status: 1, unread input bytes:
0
Local host: 150.wg.pop.50, Local port: 179
Foreign host: 150.wg.pop.49, Foreign port: 11086
 BGP neighbor is 192.168.100.129,  remote AS 65001, internal
link
  BGP version 4, remote router ID 192.168.100.129
  BGP state = Established, up for 00:19:02
  Last read 00:00:01, hold time is 180, keepalive interval is
60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                4          4
    Keepalives:            22         22
    Route Refresh:          0          0
    Total:                 27         27
  Default minimum time between advertisement runs is 5 seconds
  For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               0          0
    Prefixes Total:                 0          0
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          0
    Used as multipath:            n/a          0

                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Total:                                0          0
  Number of NLRIs in the update sent: max 0, min 0
```

```
 For address family: VPNv4 Unicast
  BGP table version 31, neighbor version 31
  Index 3, Offset 0, Mask 0x8
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               7          5 (Consumes 320
bytes)
    Prefixes Total:                 9          5
    Implicit Withdraw:              0          0
    Explicit Withdraw:              2          0
    Used as bestpath:             n/a          5
    Used as multipath:            n/a          0

                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:             10        n/a
    Total:                               10          0
  Number of NLRIs in the update sent: max 4, min 0

  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes:
0
Local host: 192.168.wg.33, Local port: 179
Foreign host: 192.168.100.129, Foreign port: 11pop5

PEpop#sh ip bgp sum
BGP router identifier 192.168.wg.17, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
192.168.100.129 4 65001      55      57        1    0    0
00:27:16        0
```


- On your PE router, verify that you are seeing the BGP routes from the other workgroups using the **show ip route bgp** command.

```
PEpop#sh ip bgp vpnv4 vrf Customer_A
BGP table version is 35, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*> 10.1.pop.16/28    150.wg.pop.17             0             0 6500wg ?
*> 10.1.pop.49/32    150.wg.pop.17             0             0 6500wg ?
*>i10.1.pop.16/28    192.168.wg.33             0    100      0 6500wg ?
*                    150.wg.pop.49           200             0 6500wg ?
*>i10.1.pop.49/32    192.168.wg.33             0    100      0 6500wg ?
*                    150.wg.pop.49           200             0 6500wg ?
r> 150.wg.pop.16/28   150.wg.pop.17            0             0 6500wg ?
r>i150.wg.pop.48/28   192.168.wg.33            0    100      0 6500wg ?
r                    150.wg.pop.49           200             0 6500wg ?
*>i150.wg.pop.16/28   192.168.wg.33            0    100      0 6500wg ?
*                    150.wg.pop.49           200             0 6500wg ?
 PEpop#sh ip bgp vpnv4 vrf Customer_B
BGP table version is 35, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: wg:20 (default for vrf Customer_B)
*> 10.0.0.0         150.wg.pop.33             0             0 6500wg ?
```

```
*>i10.2.pop.16/28    192.168.wg.33            0    100      0
6500wg ?
*>i10.2.pop.49/32    192.168.wg.33            0    100      0
6500wg ?
*> 150.wg.0.0        150.wg.pop.33            0             0
6500wg ?
*>i150.wg.pop.32/28   192.168.wg.33           0    100      0
6500wg ?
```

- Verify the per-VRF BGP table for your customer on your PE routers with the **show ip bgp vpnv4 vrf** command. You should still see that the BGP routes coming from the CE routers are being selected as the best routes for those destinations.

```
PEpop#sh ip bgp vpnv4 vrf Customer_A
BGP table version is 31, local router ID is 192.168.wg.33
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*>i10.1.pop.16/28    192.168.wg.17            0    100      0
6500wg ?
*>i10.1.pop.49/32    192.168.wg.17            0    100      0
6500wg ?
*> 10.1.pop.16/28    150.wg.pop.17            0             0
6500wg ?
*> 10.1.pop.49/32    150.wg.pop.17            0             0
6500wg ?
*>i150.wg.pop.16/28   192.168.wg.17           0    100      0
6500wg ?
*> 150.wg.pop.48/28   150.wg.pop.17           0             0
6500wg ?
r> 150.wg.pop.16/28   150.wg.pop.17           0             0
6500wg ?

PEpop#sh ip bgp vpnv4 vrf Customer_B
BGP table version is 31, local router ID is 192.168.wg.33
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight
Path
Route Distinguisher: wg:20 (default for vrf Customer_B)
```



```
*>i10.0.0.0          192.168.wg.17            0    100      0
6500wg ?
*                    150.wg.pop.49          200             0
6500wg ?
*> 10.2.pop.16/28    150.wg.pop.33            0             0
6500wg ?
*> 10.2.pop.49/32    150.wg.pop.33            0             0
6500wg ?
*>i150.wg.0.0        192.168.wg.17            0    100      0
6500wg ?
*                    150.wg.pop.49          200             0
6500wg ?
r> 150.wg.pop.32/28   150.wg.pop.33           0             0
6500wg ?
```

- Verify the per-VRF table for your customer on your PE routers with the **show ip route vrf** command. You should still see only the routes coming from the CE routers being selected.

```
PEpop#sh ip route vrf Customer_A
Routing Table: Customer_A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
Gateway of last resort is not set
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B       10.1.pop.49/32 [20/0] via 150.wg.pop.17, 00:35:25
B       10.1.pop.49/32 [200/0] via 192.168.wg.17, 00:30:28
B       10.1.pop.16/28 [20/0] via 150.wg.pop.17, 00:35:25
B       10.1.pop.16/28 [200/0] via 192.168.wg.17, 00:30:28
     150.wg.0.0/28 is subnetted, 3 subnets
B       150.wg.pop.48 [20/0] via 150.wg.pop.17, 00:35:25
C       150.wg.pop.16 is directly connected, Serial0/0.101
B       150.wg.pop.16 [200/0] via 192.168.wg.17, 00:30:28
 PEpop#sh ip route vrf Customer_B
Routing Table: Customer_B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
Gateway of last resort is not set
     10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
B       10.0.0.0/8 [200/0] via 192.168.wg.17, 00:30:48
B       10.2.pop.49/32 [20/0] via 150.wg.pop.33, 00:35:33
B       10.2.pop.16/28 [20/0] via 150.wg.pop.33, 00:35:33
     150.wg.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       150.wg.pop.48/28 is directly connected, Serial0/0.113
B       150.wg.0.0/16 [200/0] via 192.168.wg.17, 00:30:48
C       150.wg.pop.32/28 is directly connected, Serial0/0.102

PEpop#sh ip route vrf Customer_A
Routing Table: Customer_A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
Gateway of last resort is not set
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B       10.1.pop.49/32 [200/0] via 192.168.wg.33, 00:31:53
B       10.1.pop.49/32 [20/0] via 150.wg.pop.17, 01:12:09
B       10.1.pop.16/28 [200/0] via 192.168.wg.33, 00:31:53
B       10.1.pop.16/28 [20/0] via 150.wg.pop.17, 01:12:09
     150.wg.0.0/28 is subnetted, 3 subnets
C       150.wg.pop.48 is directly connected, Serial0/0.113
B       150.wg.pop.16 [200/0] via 192.168.wg.33, 00:31:53
C       150.wg.pop.16 is directly connected, Serial0/0.101
```



```
PEpop#sh ip route vrf Customer_B
Routing Table: Customer_B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route
Gateway of last resort is not set
     10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
B       10.0.0.0/8 [20/0] via 150.wg.pop.33, 01:12:18
B       10.2.pop.49/32 [200/0] via 192.168.wg.33, 00:31:51
B       10.2.pop.16/28 [200/0] via 192.168.wg.33, 00:31:51
     150.wg.0.0/16 is variably subnetted, 3 subnets, 2 masks
B       150.wg.0.0/16 [20/0] via 150.wg.pop.33, 01:12:18
B       150.wg.pop.32/28 [200/0] via 192.168.wg.33, 00:31:51
C       150.wg.pop.32/28 is directly connected, Serial0/0.102
```

Lab Exercise 6-3: Common Services VPN

The new MPLS VPN infrastructure can be used to implement a new approach to managed CE router services, where the central network management system (NMS) can monitor all CE routers through a dedicated VPN.

The NMS VPN should provide connectivity only between the NMS and a single IP address on the CE router that is used for network management purposes.

In this exercise, your service provider has established a network management center using a VPN between the loopback interfaces of the CE routers and the NMS router. You will establish connectivity only between the NMS and the CE router loopback interfaces with a /32 subnet mask.

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objectives

In this exercise, you will establish a network management VPN between the loopback interfaces of the CE routers and the NMS router. After completing this exercise, you will be able to meet these objectives:

- Design a network management VPN
- Establish connectivity between the management VRF and customer VRFs by configuring proper route targets

Visual Objective

The figure illustrates what you will accomplish in this exercise.

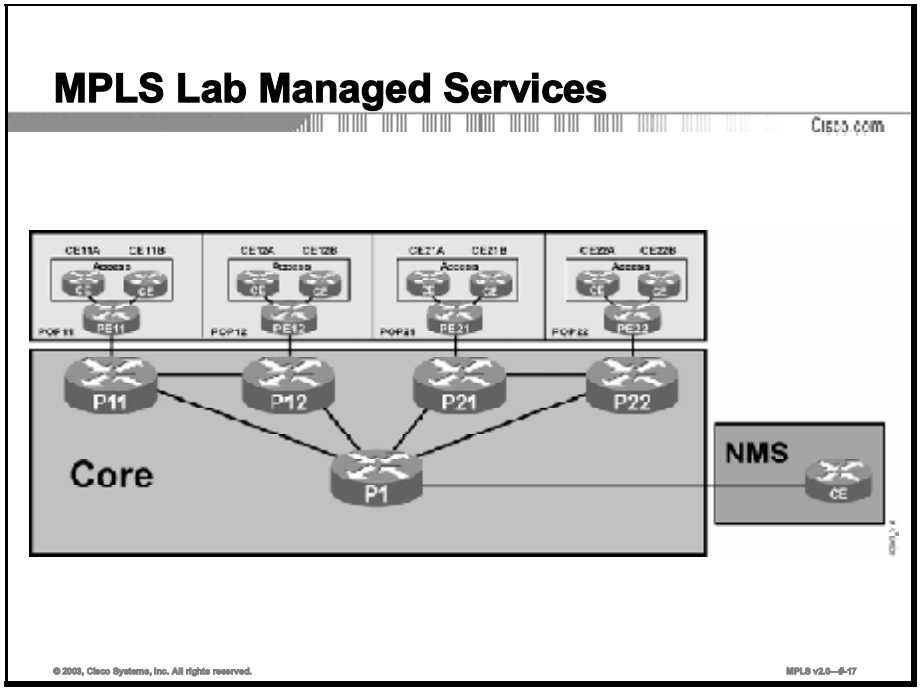

Note The NMS routers are shared between workgroups and are not configurable.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Network Management VPN Commands

Command	Description
export map *name*	Specifies a VRF export route map.
ip prefix-list *name* **permit** *address mask* **ge** *len*	Creates an IP prefix list that matches all prefixes in a specified address space with a subnet mask longer or equal to the specified value.
match ip address prefix-list *list*	Matches a prefix in a route map with a specified IP prefix list.
route-map *name* **permit** *seq*	Creates a route map entry.
set extcommunity rt *value* **additive**	Appends the specified RT to a route matched with the **match** command.

Task 1: Designing a Network Management VPN

The network management VPN is a “common services” VPN; therefore, two RTs are needed for the VPN: the server RT and the client RT. On the PE supporting the NMS, a VRF for the network management VPN and associated RD are also needed. Here are the relevant parts of the configuration on the NMS PE router:

Note	The following configuration resides on the P1 router and in this exercise serves as a PE router.

```
! Create the NMS VRF
!
ip vrf NMS
 rd 101:500
 route-target export 101:500
 route-target import 101:500
 route-target import 101:501
!
! Insert the NMS interface into the VRF
!
interface Serial0/0.201 point-to-point
 description link to simulate NMS CE
 ip vrf forwarding NMS
 ip address 150.wg0.1.17 255.255.255.240
 no cdp enable
 frame-relay interface-dlci 201
router bgp 65001
 no synchronization
 bgp log-neighbor-changes
 neighbor 192.168.1.17 activate
 neighbor 192.168.1.17 route-reflector-client
 neighbor 192.168.1.17 send-community extended
        .
        .
 no auto-summary
 !
 address-family vpnv4
 neighbor 192.168.1.17 activate
 neighbor 192.168.1.17 route-reflector-client
 neighbor 192.168.1.17 send-community extended
        .
        .
```

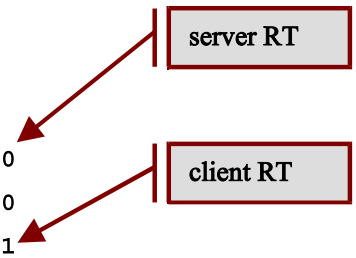


```
no auto-summary
exit-address-family
!
address-family ipv4 vrf NMS
neighbor 150.wg0.1.18 remote-as 65101
neighbor 150.wg0.1.18 activate
neighbor 150.wg0.1.18 send-community both
no auto-summary
no synchronization
exit-address-family
```

Note	If you were implementing a common services VPN from scratch, you would need to configure the supporting PE router using the VRF and routing commands used in previous exercises. In this implementation, the NMS VPN is already configured on the central service PE router, so you will need only to configure the VRF of your customer to match the RT used by the NMS VPN.

Task 2: Establishing Connectivity Between the NMS VRF and Other VRFs

To establish connectivity between the NMS VRF and the customer VRF, you must attach the *client* RT to routes toward CE router loopback addresses when they are exported from the customer VRF. You also need to import routes toward the NMS router into all customer VRFs.

Step 1 Create an IP access list that will match CE router loopback addresses.

Step 2 Create a route map that will match the CE router loopback addresses with the prefix list and append the client RT to those routes.

Step 3 Apply the route map to routes exported from the customer VRF with the **export route-map** command.

Step 4 Import NMS routes into the customer VRF by specifying the proper import RT.

Exercise Verification

- Verify that the proper RTs are appended to the routes toward CE router loopback addresses by using the **show ip bgp vpnv4 vrf** *name prefix* command. This action should result in a printout similar to the one here:

```
PEpop#sh ip bgp vpnv4 vrf Customer_B 10.2.pop.49
BGP routing table entry for wg:20:10.2.pop.49/32, version 46
Paths: (1 available, best #1, table Customer_B)
  Advertised to non peer-group peers:
  192.168.100.129
  6500wg
  150.wg.pop.33 from 150.wg.pop.33 (10.2.pop.49)
  Origin incomplete, metric 0, localpref 100, valid, external,
best
  Extended Community: RT:101:501
 PEpop#sh ip bgp vpnv4 vrf Customer_A 10.1.pop.49
BGP routing table entry for wg:10:10.1.pop.49/32, version 33
Paths: (1 available, best #1, table Customer_A)
Advertised to non peer-group peers:
192.168.100.129
6500wg
150.wg.pop.17 from 150.wg.pop.17 (10.1.pop.49)
Origin incomplete, metric 0, localpref 100, valid, external,
best
    Extended Community: RT:101:501
```

- Using an extended ping, verify that you can ping from the loopback address of the managed CE router to the loopback address of the NMS CE router (10.10.10.49).
- Using an extended ping, verify that you cannot ping from the Ethernet address of the managed CE router to the loopback address of the NMS CE router (10.10.10.49).

- Verify that your CE router is seeing only prefixes within your VPN and that no prefixes are being leaked from other VPNs.

```
PEpop#sh ip bgp vpnv4 vrf Customer_A
BGP table version is 44, local router ID is 192.168.wg.33
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop               Metric LocPrf
Weight Path
Route Distinguisher: wg:10 (default for vrf Customer_A)
*>i10.1.pop.16/28     192.168.wg.17             0    100
0 6500wg ?
*>i10.1.pop.49/32     192.168.wg.17             0    100
0 6500wg ?
*> 10.1.pop.16/28     150.wg.pop.17             0
0 6500wg ?
*> 10.1.pop.49/32     150.wg.pop.17             0
0 6500wg ?
*>i10.10.10.49/32     192.168.100.129           0    100
0 65101 ?
*>i150.wg.pop.16/28   192.168.wg.17             0    100
0 6500wg ?
*> 150.wg.pop.48/28   150.wg.pop.17             0
0 6500wg ?
r> 150.wg.pop.16/28   150.wg.pop.17             0
0 6500wg ?

PEpop#sh ip bgp vpnv4 vrf Customer_B
BGP table version is 85, local router ID is 192.168.wg.17
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
    Network            Next Hop               Metric LocPrf
Weight Path
Route Distinguisher: wg:20 (default for vrf Customer_B)
*> 10.2.pop.16/28     150.wg.pop.33              0
0 6500wg ?
*> 10.2.pop.49/32     150.wg.pop.33              0
0 6500wg ?
*>i10.2.pop.16/28     192.168.wg.33              0    100
0 6500wg ?
```

```
*>i10.2.pop.49/32     192.168.wg.33             0    100
0 6500wg ?
*>i10.10.10.49/32     192.168.100.129           0    100
0 65101 ?
r> 150.wg.pop.32/28   150.wg.pop.33             0
0 6500wg ?
*>i150.wg.pop.32/28   192.168.wg.33             0    100
0 6500wg ?
*> 150.wg.pop.48/28   150.wg.pop.33             0
0 6500wg ?
```

Next Step

- **Module 7:** Internet Access from an MPLS VPN

Lab Exercise 7-1: Separate Interface for Internet Connectivity

In many cases customers may want to retain the traditional Internet access model with a firewall between the customer VPN and the global Internet. This request is usually implemented by using dedicated VPN and Internet subinterfaces on the physical PE-CE link.

Exercise Objectives

After completing this exercise, you will be able to meet these objectives:

- Establish CE-PE connectivity for Internet access
- Establish routing between the customer and the Internet

Visual Objective

You will configure additional virtual links (emphasized in the figure here) between the central site CE routers (CE*wg*1A and CE*wg*2B) and their PE routers. These circuits will be in the global routing table, and you will configure static routing between the PE and CE routers. The remote sites (CE*wg*1B and CE*wg*2A) will access the Internet using the MPLS VPN connection to its central site.

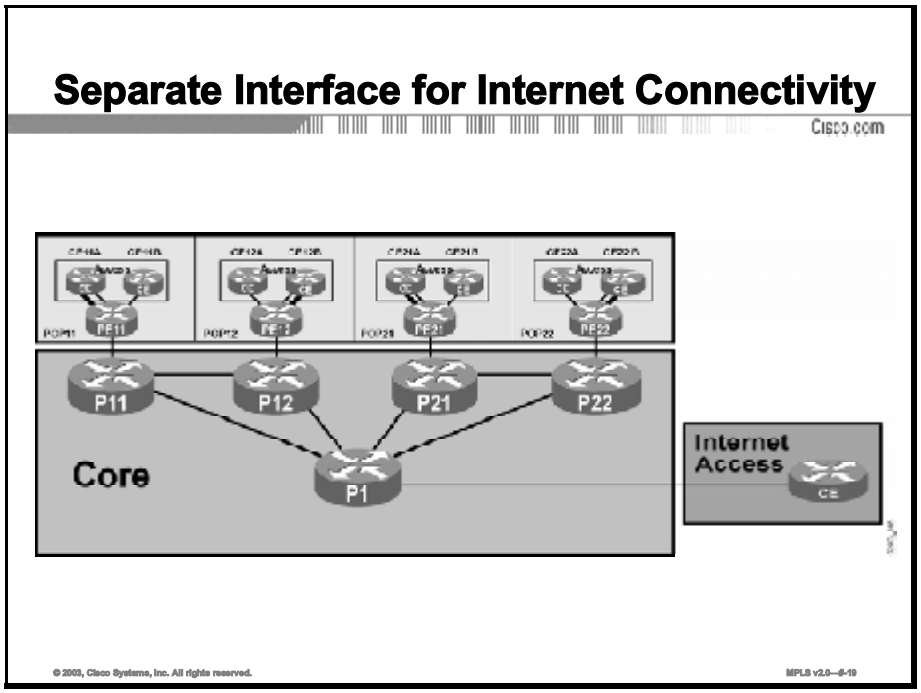

Note In this lab, the customer addressing scheme is in the private addressing range. In an actual implementation, a Network Address Translation (NAT) service would need to be provided at the customer interface to the Internet access point. Because NAT is outside the scope of this course, this function is omitted and the lab has been set up to ensure that the customer addressing does not overlap.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Internet Access Commands

Command	Description
ip route *prefix mask* **null 0**	Creates a summary route in the IP routing table.

Task 1: Establishing CE-PE Connectivity for Internet Access

In this task you will add a new subinterface to support Internet access on the central site router.

Exercise Procedure

Complete these steps:

Step 1 Create a separate subinterface (S0/0.114) on CE*wg*1A (team A) and CE*wg*2B (team B) using DLCI 114 and IP address 150.*wg.pop*.66/28.

Step 2 Activate the new interface in the IGP routing process and make the interface passive.

Step 3 Create a separate subinterface on PE*wg*1 (team A) and PE*wg*2 (team B) using DLCI 114 and IP address 150.*wg.pop*.65/28.

Step 4 Activate the new interface in the IGP routing process and make the interface passive.

Note Global routing between your PE router and P1 was established in the "Merging Service Providers" lab exercise.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **show ip interface** command to verify the status of the new interfaces.

```
CE***#sh ip int s0/0.114
Serial0/0.114 is up, line protocol is up
  Internet address is 150.wg.pop.65/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is enabled
```

```
  IP Flow switching is disabled
  IP Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect inbound is disabled
  WCCP Redirect exclude is disabled
  BGP Policy Mapping is disabled


PEpop#sh ip int s0/0.114
Serial0/0.114 is up, line protocol is up
  Internet address is 150.wg.pop.66/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is enabled
  IP Flow switching is disabled
  IP CEF switching is enabled
```



```
IP CEF VPN Fast switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Policy routing is disabled
Network address translation is disabled
WCCP Redirect outbound is disabled
WCCP Redirect inbound is disabled
WCCP Redirect exclude is disabled
BGP Policy Mapping is disabled
```

Task 2: Establishing Routing Between the Customer and the Internet

In this solution, the customer and the service provider have decided to use static routing, for the PE-CE Internet routing protocol. In this task you will enable a static route on the CE router that points to the Internet and a static route on the PE router that points to the customer public address range.

Exercise Procedure

Complete these steps:

Step 1 On the PE router that is supporting your CE router, create a static route that points to the customer address range.

Note Your first choice for the static route would most likely be 10.1.0.0/16 for customer A and 10.2.0.0/16 for customer B. However, if you examine the addressing scheme used in these labs, you will notice that customer A on all pods uses the same 10.1.0.0 address range. The same is true for customer B, which uses 10.2.0.0 on all pods. To ensure that your static routes do not overlap with the other pods, you will need a statement for each customer site.

Step 2 Redistribute this route into BGP so that it will be advertised to the Internet access point.

Step 3 On your CE router, create a default route that will point all unknown routes to the Internet interface.

Step 4 This static route will be used by both the local and the remote VPN sites. Because of this shared use, you will need to interject the route into both the local and remote routing tables. You can accomplish this task by adding a network statement to the BGP process that enables network 0.0.0.0.

Note For security reasons, the customer never wants packets that originate in its network or that are addressed to its network to be sent out to the Internet. Creating a default route that points all unroutable customer packets to the null interface will address this issue.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify the static route on the PE router.

```
PEpop#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
```



```
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

**** output omitted *******

     10.0.0.0/24 is subnetted, 2 subnets
S       10.1.pop.0 [1/0] via 150.wg.11.66
S       10.1.pop.0 [1/0] via 150.wg.11.66

***** outout omitted ******



Verify the static routes on the CE routers
CE***#sh ip rou
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is 150.wg.pop.66 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
S       10.1.0.0/24 is directly connected, Null0
C       **** output omitted *****
S*   0.0.0.0/0 [1/0] via 150.wg.pop.65
```

- Use an extended ping to verify that host addresses with the customer network can reach the Internet.

```
CE***#ping
Protocol [ip]:
Target IP address: 201.202.26.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.wg.pop.49
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 201.202.26.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
100/135/193 ms
```


Lab Exercise 7-2: Multisite Internet Access

To provide optimum routing, the service provider has convinced the customer to provide Internet access to each site. Because of the multisite access, the routing will have to be converted from static to BGP.

Note	This conversion will require additional firewall and NAT services that are not addressed by this lab exercise.

Exercise Objectives

After completing this exercise, you will be able to meet these objectives:

- Establish remote site CE-PE connectivity for Internet access
- Establish remote site routing between the customer and the Internet

Visual Objective

You will configure additional virtual links (emphasized in the figure here) between the central site CE routers (CE*wg*1B and CE*wg*2A) and their PE routers. You will put these circuits and those created in the previous lab in the global routing table. You will also configure a global BGP session between PE routers and CE routers to exchange Internet routes between the service provider and the customer.

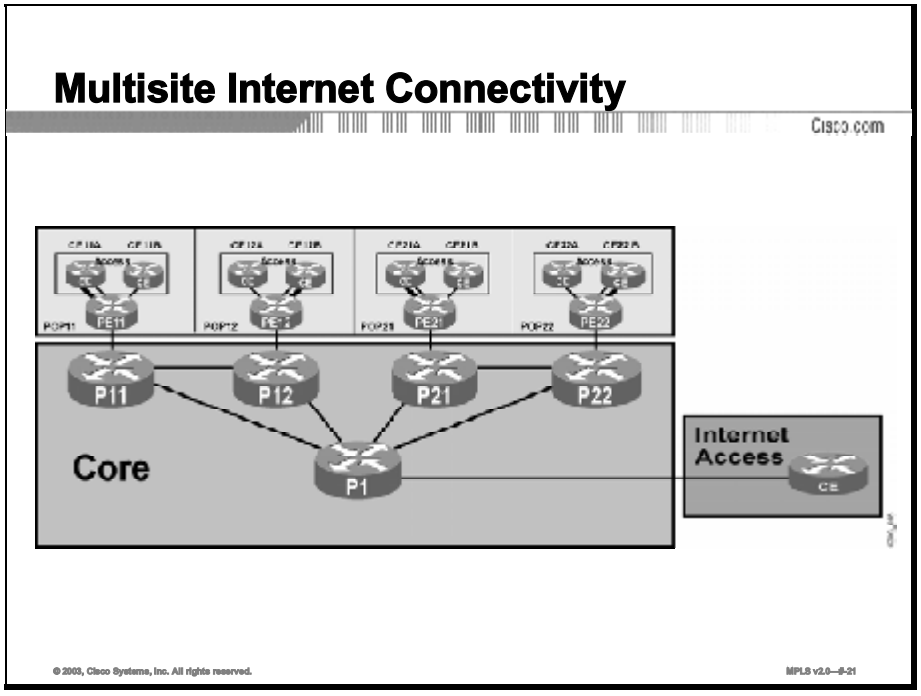

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Multisite Internet Access Commands

Command	Description
ip route *prefix mask* **null 0**	Creates a summary route in the IP routing table.

Task 1: Establishing CE-PE Connectivity for Internet Access

Your service provider has already created a VPN to carry Internet traffic. You will need to join this VPN.

Exercise Procedure

Complete these steps:

Step 1 Create a separate subinterface (S0/0.115) on CE*wg*1B (team B) and CE*wg*2A (team A) using DLCI 115 and IP address 150.*wg.pop*.130/28.

Step 2 Create a separate subinterface on PE*wg*1 (team A) and PE*wg*2 (team B) using DLCI 115 and IP address 150.*wg.pop*.129/28.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify the status of the interface.

```
CE***#sh ip int s0/0.115
Serial0/0.115 is up, line protocol is up
  Internet address is 150.wg.pop.130/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
```



```
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is enabled
  IP Flow switching is disabled
  IP Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect inbound is disabled
  WCCP Redirect exclude is disabled
BGP Policy Mapping is disabled

PE12#sh ip int s0/0.115
Serial0/0.115 is up, line protocol is up
  Internet address is 150.wg.pop.129/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is enabled
```

```
IP Flow switching is disabled
IP CEF switching is enabled
IP CEF VPN Fast switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Policy routing is disabled
Network address translation is disabled
WCCP Redirect outbound is disabled
WCCP Redirect inbound is disabled
WCCP Redirect exclude is disabled
BGP Policy Mapping is disabled
```


Task 2: Establishing Routing Between the Customer and the Internet

The next task is to convert the interface created in the "Separate Interface for Internet Connectivity" lab exercise from static routing to an EBGP session. You then need to enable an EBGP session on the new interface.

Exercise Procedure

Complete these steps:

Step 1 On your assigned CE router (CE*wg*1A – team A and CE*wg*2B – team B), remove the network statement and passive interface command related to the WAN interface from the customer IGP process.

Step 2 Remove the network statement that refers to network 0.0.0.0 from BGP.

Step 3 Remove the 0.0.0.0 static route.

Step 4 Add the associated PE router as a BGP neighbor.

Step 5 On the associated PE router, add the associated CE router as a BGP neighbor.

Step 6 On your assigned CE router (CE*wg*2A – team A and CE*wg*1B – team B), add the associated PE router as a BGP neighbor.

Step 7 On the associated PE router, add the associated CE router as a BGP neighbor.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify the status of the BGP neighbors.

```
PEpop#sh ip bgp sum
BGP router identifier 192.168.1.17, local AS number 65001
BGP table version is 41, main routing table version 41
36 network entries using 3636 bytes of memory
36 path entries using 1728 bytes of memory
18 BGP path attribute entries using 1080 bytes of memory
2 BGP rrinfo entries using 48 bytes of memory
3 BGP AS-PATH entries using 72 bytes of memory
9 BGP extended community entries using 320 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6884 total bytes of memory
BGP activity 96/9 prefixes, 107/13 paths, scan interval 60
secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ
Up/Down  State/PfxRcd
```

```
150.wg.pop.66   4 650xx      10        7       41    0    0
00:01:57        4
150.wg.pop.130  4 650xx      10        7       41    0    0
00:01:57        4
192.168.100.129 4 65001     129     114       41    0    0
01:30:45       32
```

```
CE***#ping
Protocol [ip]:
Target IP address: 201.202.26.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.wg.pop.49
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 201.202.26.1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
100/135/193 ms
```


Lab Exercise 7-3: Internet Connectivity in an MPLS VPN

Internet connectivity in MPLS VPN-based networks can be achieved through a dedicated Internet VPN. The dedicated Internet VPN approach gives you better security as it completely isolates the service provider core (P routers) from the Internet. On the other hand, it is also less scalable—for example, you cannot transport full Internet routing in an Internet VPN.

Exercise Objectives

After completing this exercise, you will be able to meet these objectives:

- Establish central site CE-PE connectivity for Internet access
- Establish remote site CE-PE connectivity for Internet access

Visual Objective

In this lab you will create a VPN (VRF) that will carry all Internet traffic, and then create connectivity between that VPN and the customer site. Each team will be responsible for performing the configuration tasks on its PE router.

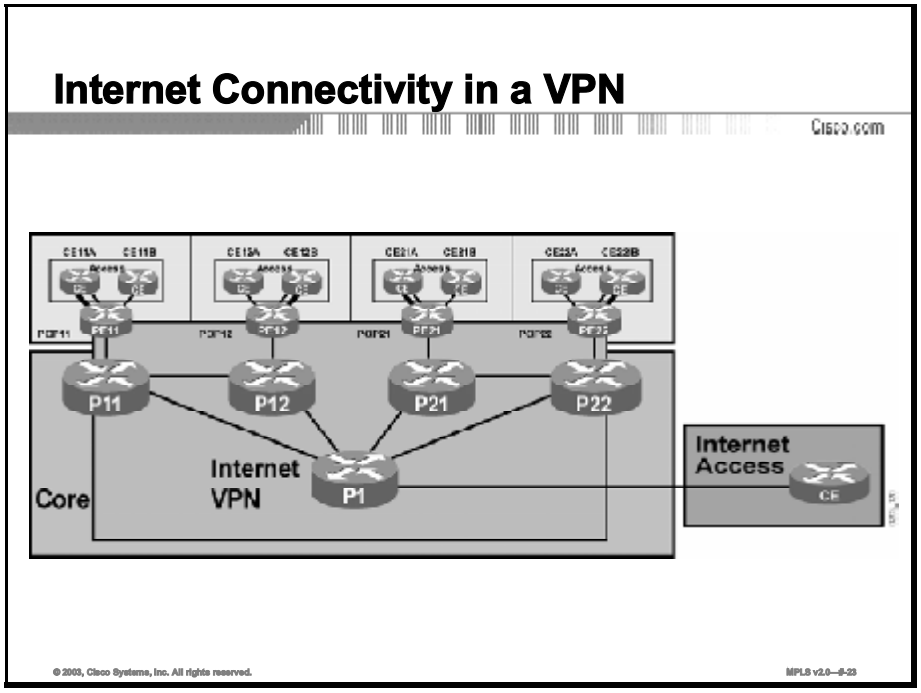

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

All commands used in this lab have been used in previous labs.

Task 1: Establishing Central Site Connectivity for Internet Access

Your service provider has already created a VPN to carry the Internet traffic. You will need to join this VPN.

Exercise Procedure

Complete these steps:

Step 1 On your assigned PE router (PE*wg*1 – team A and PE*wg*2 – team B), create a new Internet VPN VRF. The service provider has assigned an RT of 100:600 and an RD of 100:600 for all Internet-related VRFs.

Step 2 Place the interface (114) that is supporting the central site CE router (CE*wg*1A – team A and CE*wg*2B – team B) into the VRF.

Step 3 Remove the central site router neighbor statement for the unicast (global) address family.

Step 4 Add the central site router neighbor statement to the IPv4 VRF address family for the Internet VRF.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify that the Internet routes being received by the central site CE route are coming from its PE neighbor.

```
CEwg1A#sh ip rou
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set
```



```
B    201.202.20.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    202.100.36.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    207.69.48.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    201.202.21.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    202.100.37.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    207.69.49.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    201.202.22.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    202.100.38.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    201.202.23.0/24 [20/0] via 150.wg.wg1.65, 01:54:31
B    202.100.39.0/24 [20/0] via 150.wg.wg1.65, 01:54:32
B    202.100.32.0/24 [20/0] via 150.wg.wg1.65, 01:54:32
****************** output omitted ************************
```

Task 2: Establishing Remote Site Connectivity for Internet Access

Your service provider has already created a VPN to carry Internet traffic. You will need to join this VPN.

Exercise Procedure

Complete these steps:

Step 1 On the PE router (PE*wg*2 – team A and PE*wg*1 – team B) that supports your remote CE router, place the interface (115) that is supporting the remote site CE router (CE*wg*2A – team A and CE*wg*1B – team B) into the VRF.

Step 2 Remove the remote site router neighbor statement for the unicast (global) address family.

Step 3 Add the remote site router neighbor statement to the IPv4 VRF address family for the Internet VRF.

Exercise Verification

You have completed this exercise when you attain these results:

- Verify that the Internet routes being received by the central site CE route are coming from its PE neighbor.

```
CEwgA#sh ip rou
Codes: C - connected, S - static, I - IGRP, R - RIP, M -
mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF
inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E
- EGP
```

```
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia -
IS-IS inter area
       * - candidate default, U - per-user static route, o -
ODR
       P - periodic downloaded static route

Gateway of last resort is not set

B    201.202.20.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    202.100.36.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    207.69.48.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    201.202.21.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    202.100.37.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    207.69.49.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    201.202.22.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    202.100.38.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    201.202.23.0/24 [20/0] via 150.wg.wg2.129, 01:44:20
B    202.100.39.0/24 [20/0] via 150.wg.wg2.129, 01:44:21
B    202.100.32.0/24 [20/0] via 150.wg.wg2.129, 01:44:21
B    202.100.33.0/24 [20/0] via 150.wg.wg2.129, 01:44:21
```


Lab Exercise Answer Key

Lab Exercise 3-1: Establishing the Service Provider IGP Routing Environment

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 2: Configuring the Service Provider IGP

Configuration steps on PE*wg*1:

```
PEwg1(config)#router eigrp 1
PEwg1(config)#network 150.wg#.0.0 (optional)
PEwg1(config)#network 192.168.wg#.0
```

Configuration steps on P*wg*1:

```
Pwg1(config)#router eigrp 1
Pwg1(config)#network 192.168.wg#.0
```

Configuration steps on P*wg*2:

```
Pwg2(config)#router eigrp 1
Pwg2(config)#network 192.168.wg#.0
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#router eigrp 1
PEwg2(config)#network 150.wg#.0.0 (optional)
PEwg2(config)#network 192.168.wg#..0
```

Lab Exercise 3-2: Establishing the Core MPLS Environment

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Enabling LDP on Your PE and P Routers

Configuration steps on PE*wg*1:

```
PEwg1(config)#ip cef
PEwg1(config)#interface Serial0/0.111
PEwg1(config-subif)#mpls label protocol ldp
PEwg1(config-subif)#mpls ip
```

Configuration steps on P*wg*1:

```
Pwg1(config)#ip cef
Pwg1(config)#interface Serial0/0.111
Pwg1(config-subif)#mpls label protocol ldp
Pwg1(config-subif)#mpls ip
```

```
Pwg1(config)#interface Serial0/0.112
Pwg1(config-subif)#mpls label protocol ldp
Pwg1(config-subif)#mpls ip
```

Configuration steps on P*wg2*:

```
Pwg2(config)#ip cef
Pwg2(config)#interface Serial0/0.111
Pwg2(config-subif)#mpls label protocol ldp
Pwg2(config-subif)#mpls ip
Pwg2(config)#interface Serial0/0.112
Pwg2(config-subif)#mpls label protocol ldp
Pwg2(config-subif)#mpls ip
```

Configuration steps on PE*wg2*:

```
PEwg2(config)#ip cef
PEwg2(config)#interface Serial0/0.111
PEwg2(config-subif)#mpls label protocol ldp
PEwg2(config-subif)#mpls ip
```

Note **mpls label protocol ldp** can be issued at the global configuration level.

Note **mpls ip** is a command issued to enable MPLS on an interface but will be displayed in the configuration (sh run) as **tag-switching ip**.

Task 2: Disabling TTL Propagation

Configuration steps on PE*wg1*:

```
PEwg1(config)# no tag-switching ip propagate-ttl
```

Configuration steps on P*wg1*:

```
Pwg1(config)# no tag-switching ip propagate-ttl
```

Configuration steps on P*wg2*:

```
Pwg2(config)# no tag-switching ip propagate-ttl
```

Configuration steps on PE*wg2*:

```
PEwg2(config)# no tag-switching ip propagate-ttl
```


Task 3: Configuring Conditional Label Distribution

Note	There are different ways to construct an access list to accomplish the desired result. This is one way. The key, however, is to meet the task objective.

Configuration steps on PE*wg*1:

```
PEwg1(config)#no tag-switching advertise-tags
PEwg1(config)#tag-switching advertise-tags for 90
PEwg1(config)#access-list 90 permit 150.wg.0.0 0.0.255.255
PEwg1(config)#access-list 90 permit 192.168.wg.16 0.0.0.15
PEwg1(config)#access-list 90 permit 192.168.wg.32 0.0.0.15
PEwg1(config)#access-list 90 permit 192.168.wg.80 0.0.0.15
PEwg1(config)#access-list 90 permit 192.168.wg.96 0.0.0.15
```

Configuration steps on P*wg*1:

```
Pwg1(config)#no tag-switching advertise-tags
Pwg1(config)#tag-switching advertise-tags for 90
Pwg1(config)#access-list 90 permit 150.wg.0.0 0.0.255.255
Pwg1(config)#access-list 90 permit 192.168.wg.16 0.0.0.15
Pwg1(config)#access-list 90 permit 192.168.wg.32 0.0.0.15
Pwg1(config)#access-list 90 permit 192.168.wg.80 0.0.0.15
Pwg1(config)#access-list 90 permit 192.168.wg.96 0.0.0.15
```

Configuration steps on P*wg*2:

```
Pwg2(config)#no tag-switching advertise-tags
Pwg2(config)#tag-switching advertise-tags for 90
Pwg2(config)#access-list 90 permit 150.wg.0.0 0.0.255.255
Pwg2(config)#access-list 90 permit 192.168.wg.16 0.0.0.15
Pwg2(config)#access-list 90 permit 192.168.wg.32 0.0.0.15
Pwg2(config)#access-list 90 permit 192.168.wg.80 0.0.0.15
Pwg2(config)#access-list 90 permit 192.168.wg.96 0.0.0.15
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#no tag-switching advertise-tags
PEwg2(config)#tag-switching advertise-tags for 90
PEwg2(config)#access-list 90 permit 150.wg.0.0 0.0.255.255
PEwg2(config)#access-list 90 permit 192.168.wg.16 0.0.0.15
PEwg2(config)#access-list 90 permit 192.168.wg.32 0.0.0.15
PEwg2(config)#access-list 90 permit 192.168.wg.80 0.0.0.15
PEwg2(config)#access-list 90 permit 192.168.wg.96 0.0.0.15
```

Task 4: Removing Conditional Label Distribution

Configuration steps on P*wg*2:

```
Pwg2(config)#tag-switching advertise-tags
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#tag-switching advertise-tags
```

Configuration steps on PE*wg*1:

```
PEwg1(config)#tag-switching advertise-tags
```

Configuration steps on P*wg*1:

```
Pwg1(config)#tag-switching advertise-tags
```

Lab Exercise 5-1: Initial MPLS VPN Setup

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configuring Multiprotocol BGP

Configuration steps on PE*wg*1:

```
PEwg1(config)#router bgp 65001
PEwg1(config-router)#nei 192.168.wg.33 remote-as 65001
PEwg1(config-router)#nei 192.168.wg.33 update-source l00
PEwg1(config-router)#address-family vpnv4
PEwg1(config-router-af)#nei 192.168.wg.33 activate
PEwg1(config-router-af)#nei 192.168.wg.33 next-hop-self
PEwg1(config-router-af)#nei 192.168.wg.33 send-community both
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#router bgp 65001
PEwg2(config-router)#nei 192.168.wg.17 remote-as 65001
PEwg2(config-router)#nei 192.168.wg.17 update-source l00
PEwg2(config-router)#address-family vpnv4
PEwg2(config-router-af)#nei 192.168.wg.17 activate
PEwg2(config-router-af)#nei 192.168.wg.17 next-hop-self
PEwg2(config-router-af)#nei 192.168.wg.17 send-community both
```

Note In the two commands **`nei 192.168.`*wg*`.33 update-source l00`** and **nei 192.168.*wg*.17 update-source l00**, the *l00 is the loopback* and not 100.

Task 2: Configuring Virtual Routing and Forwarding Tables

Configuration steps on PE*wg*1:

```
PEwg1(config)#ip vrf Customer_A
PEwg1(config-vrf)#rd wg:10
PEwg1(config-vrf)#route-target both wg:10
PEwg1(config)#ip vrf Customer_B
PEwg1(config-vrf)#rd wg:20
PEwg1(config-vrf)#route-target both wg:20
PEwg1(config)#int s0/0.101
PEwg1(config-subif)#ip vrf forwarding Customer_A
PEwg1(config-subif)#ip address 150.wg.pop.18 255.255.255.240
PEwg1(config)#int s0/0.102
PEwg1(config-subif)#ip vrf forwarding Customer_B
PEwg1(config-subif)#ip address 150.wg.pop.34 255.255.255.240
PEwg1(config)#router rip
PEwg1(config-router)#version 2
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#network 150.wg.0.0
PEwg1(config-router-af)#redistribute bgp 65001 metric
transparent
PEwg1(config-router)#address-family ipv4 vrf Customer_B
PEwg1(config-router-af)#network 150.wg.0.0
PEwg1(config-router-af)#redistribute bgp 65001 metric
transparent
PEwg1(config-router)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#redistribute rip
PEwg1(config-router-af)#exit
PEwg1(config-router)#address-family ipv4 vrf Customer_B
PEwg1(config-router-af)#redistribute rip
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#ip vrf Customer_A
PEwg2(config-vrf)#rd wg:10
PEwg2(config-vrf)#route-target both wg:10
PEwg2(config)#ip vrf Customer_B
PEwg2(config-vrf)#rd wg:20
PEwg2(config-vrf)#route-target both wg:20
PEwg2(config)#int s0/0.101
PEwg2(config-subif)#ip vrf forwarding Customer_A
PEwg2(config-subif)#ip address 150.wg.pop.18 255.255.255.240
```

```
PEwg2(config)#int s0/0.102
PEwg2(config-subif)#ip vrf forwarding Customer_B
PEwg2(config-subif)# ip address 150.wg.pop.34 255.255.255.240
PEwg2(config)#router rip
PEwg2(config-router)#version 2
PEwg2(config-router)#address-family ipv4 vrf Customer_A
PEwg2(config-router-af)#network 10.0.0.0
PEwg2(config-router-af)#network 150.wg.0.0
PEwg2(config-router-af)#redistribute bgp 65001 metric
transparent
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#network 10.0.0.0
PEwg2(config-router-af)# network 150.wg.0.0
PEwg2(config-router-af)#redistribute bgp 65001 metric
transparent
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_A
PEwg2(config-router-af)#redistribute rip
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#redistribute rip
```

Lab Exercise 5-2: Running EIGRP Between PE and CE Routers

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Enabling an EIGRP VPN

Configuration steps on CE*wg*1A:

```
CEwg1A(config)#no router rip
CEwg1A(config)#router eigrp wg
CEwg1A(config-router)#network 10.0.0.0
CEwg1A(config-router)#network 150.wg.0.0
CEwg1A(config-router)#no auto-summary
```

Configuration steps on CE*wg*2B:

```
CEwg2B(config)#no router rip
CEwg2B(config)#router eigrp wg
CEwg2B(config-router)#network 10.0.0.0
CEwg2B(config-router)#network 150.wg.0.0
CEwg2B(config-router)#no auto-summary
```


Configuration steps on PE*wg*1:

```
PEwg1(config)#router rip
PEwg1(config-router)#no address-family ipv4 vrf Customer_A
PEwg1(config)#router eigrp 1
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#autonomous-system wg
PEwg1(config-router-af)#network 10.0.0.0
PEwg1(config-router-af)#network 150.wg.0.0
PEwg1(config-router-af)#redistribute bgp 65001 metric 10000
100 255 1 1500
PEwg1(config-router-af)#exit
PEwg1(config-router)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#no redistribute rip
PEwg1(config-router-af)#redistribute eigrp wg
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#router rip
PEwg2(config-router)#no address-family ipv4 vrf Customer_B
PEwg2(config-router)#router eigrp 1
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#autonomous-system wg
PEwg2(config-router-af)#network 10.0.0.0
PEwg2(config-router-af)#network 150.wg.0.0
PEwg2(config-router-af)#redistribute bgp 65001 metric 10000
100 255 1 1500
PEwg2(config-router-af)#exit
PEwg2(config-router)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#no redistribute rip
PEwg2(config-router-af)#redistribute eigrp wg metric 1
```

Lab Exercise 5-3: Running OSPF Between PE and CE Routers

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configuring Virtual Routing and Forwarding Tables

Configuration steps on CE*wg*1B:

```
CEwg1B(config)#no router rip
CEwg1B(config)#router ospf 2
CEwg1B(config-router)#network 150.wg.0.0 0.0.255.255 area 0
CEwg1B(config-router)#network 10.2.pop.49 0.0.0.0 area 0
CEwg1B(config-router)#network 10.2.pop.16 0.0.0.15 area 1
```

Configuration steps on CE*wg*2A:

```
CEwg2A(config)#no router rip
CEwg2A(config)#router ospf 1
CEwg2A(config-router)# network 150.wg.0.0 0.0.255.255 area 0
CEwg2A(config-router)# network 10.1.pop.49 0.0.0.0 area 0
CEwg2A(config-router)# network 10.1.pop.16 0.0.0.15 area 1
```

Configuration steps on PE*wg*1:

```
PEwg1(config)#no router rip
PEwg1(config)#router ospf 2 vrf Customer_B
PEwg1(config-router)#network 150.wg.0.0 0.0.255.255 area 0
PEwg1(config-router)#redistribute bgp 65001 subnets
PEwg1(config-router)#exit
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_B
PEwg1(config-router)# no redistribute rip
PEwg1(config-router-af)#redistribute ospf 2
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#no router rip
PEwg2(config)#router ospf 1 vrf Customer_A
PEwg2(config-router)# network 150.wg.0.0 0.0.255.255 area 0
PEwg2(config-router)#redistribute bgp 65001 subnets
PEwg2(config-router)#exit
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_A
PEwg2(config-router)# no redistribute rip
PEwg2(config-router-af)#redistribute ospf 1
```


Task 2: Completing the OSPF Migration

Configuration steps on CE*wg*1A:

```
CEwg1A(config)#no router eigrp wg
CEwg1A(config)#router ospf 1
CEwg1A(config-router)#network 150.wg.0.0 0.0.255.255 area 0
CEwg1A(config-router)#network 10.1.pop.49 0.0.0.0 area 0
CEwg1A(config-router)#network 10.1.pop.16 0.0.0.15 area 1
```

Configuration steps on CE*wg*2B:

```
CEwg2B(config)#no router eigrp wg
CEwg2B(config)#router ospf 2
CEwg2B(config-router)# network 150.wg.0.0 0.0.255.255 area 0
CEwg2B(config-router)# network 10.2.pop.49 0.0.0.0 area 0
CEwg2B(config-router)# network 10.2.pop.16 0.0.0.15 area 1
```

Configuration steps on PE*wg*1:

```
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#no redistribute eigrp wg
PEwg1(config-router-af)#exit
PEwg1(config-router)#exit
PEwg1(config)#router eigrp 1
PEwg1(config-router)#no address-family ipv4 vrf Customer_A
PEwg1(config)#router ospf 1 vrf Customer_A
PEwg1(config-router)#network 150.wg.0.0 0.0.255.255 area 0
PEwg1(config-router)#redistribute bgp 65001 subnets
PEwg1(config-router)#exit
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#redistribute ospf 1
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#no redistribute eigrp wg
PEwg2(config-router-af)#exit
PEwg2(config-router)#exit
PEwg2(config)#router eigrp 1
PEwg2(config-router)#no address-family ipv4 vrf Customer_B
PEwg2(config)#router ospf 2 vrf Customer_B
PEwg2(config-router)#network 150.wg.0.0 0.0.255.255 area 0
```

```
PEwg2(config-router)#redistribute bgp 65001 subnets
PEwg2(config-router)#exit
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#redistribute ospf 2
```

Lab Exercise 5-4: Running BGP Between PE and CE Routers

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task1: Configuring BGP as the PE-CE Routing Protocol

Configuration steps on CE*wg*1A:

```
CEwg1A(config)#router bgp 650wg1
CEwg1A(config-router)#nei 150.wg.pop.18 remote-as 65001
CEwg1A(config-router)#no auto-summary
CEwg1A(config-router)#redistribute ospf 1
CEwg1A(config)#router ospf 1
CEwg1A(config-router)#redistribute bgp 650wg1 subnets
```

Configuration steps on CE*wg*1B:

```
CEwg1B(config)#router bgp 650wg2
CEwg1B(config-router)#nei 150.wg.pop.34 remote-as 65001
CEwg1B(config-router)#no auto-summary
CEwg1B(config-router)#redistribute ospf 2
CEwg1B(config-router)#router ospf 2
CEwg1B(config-router)#redistribute bgp 650wg2 subnets
```

Configuration steps on CE*wg*2A:

```
CEwg2A(config)#router bgp 650wg1
CEwg2A(config-router)#nei 150.wg.pop.18 remote-as 65001
CEwg2A(config-router)#no auto-summary
CEwg2A(config-router)#redistribute ospf 1
CEwg2A(config-router)#router ospf 1
CEwg2A(config-router)#redistribute bgp 650wg1 subnets
```

Configuration steps on CE*wg2*B:

```
CEwg2B(config)#router bgp 650wg2
CEwg2B(config-router)#nei 150.wg.pop.34 remote-as 65001
CEwg2B(config-router)#no auto-summary
CEwg2B(config-router)#redistribute ospf 2
CEwg2B(config-router)#router ospf 2
CEwg2B(config-router)#redistribute bgp 650wg2 subnets
```


Configuration steps on PE*wg*1:

```
!******* Team A *************
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#no redistribute ospf 1
PEwg1(config)#no router ospf 1 vrf Customer_A
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#nei 150.wg.pop.17 remote-as 650wg1
PEwg1(config-router-af)#nei 150.wg.pop.17 activate
PEwg1(config-router-af)#nei 150.wg.pop.17 as-override


!******* Team B **************
PEwg1(config)#router bgp 65001

PEwg1(config-router-af)#address-family ipv4 vrf Customer_B
PEwg1(config-router-af)#no redistribute ospf 2
PEwg1(config)#no router ospf 2 vrf Customer_B
PEwg1(config)#router bgp 65001
PEwg1(config-router-af)#address-family ipv4 vrf Customer_B
PEwg1(config-router-af)#nei 150.wg.pop.33 remote-as 650wg2
PEwg1(config-router-af)#nei 150.wg.pop.33 activate
PEwg1(config-router-af)#nei 150.wg.pop.33 as-override
```

Configuration steps on PE*wg*2:

```
!******* Team A *******
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_A
PEwg2(config-router-af)#no redistribute ospf 1
PEwg2(config)#no router ospf 1 vrf Customer_A
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_A
PEwg2(config-router-af)#nei 150.wg.pop.17 remote-as 650wg1
PEwg2(config-router-af)#nei 150.wg.pop.17 activate
PEwg2(config-router-af)#nei 150.wg.pop.17 as-override

!******* Team B *******
PEwg2(config-router-af)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#no redistribute ospf 2
```

```
PEwg2(config)#no router ospf 2 vrf Customer_B
PEwg2(config)#router bgp 65001
PEwg2(config-router-af)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#nei 150.wg.pop.33 remote-as 650wg2
PEwg2(config-router-af)#nei 150.wg.pop.33 activate
PEwg2(config-router-af)#nei 150.wg.pop.33 as-override
```

Task 2: Configuring the Backup PE-CE Link

Configuration steps on CE*wg*1B:

```
CEwg1B(config)#int s0/0.113 point-to-point
CEwg1B(config-subif)# ip address 150.wg.x.49 255.255.255.240
CEwg1B(config-subif)#frame-relay interface-dlci 113
CEwg1B(config-fr-dlci)#no shut
CEwg1B(config)#router bgp 6500wg
CEwg1B(config-router)#nei 150.wg.wg2.50 remote-AS 65001
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#interface s0/0.113 point-to-point
PEwg2(config-subif)#ip vrf forwarding Customer_B
PEwg2(config-subif)#ip address 150. wg.pop.50 255.255.255.240
PEwg2(config-subif)#frame-relay interface-dlci 113
PEwg2(config-fr-dlci)#no shut
PEwg2(config)#router bgp 65001
PEwg2(config-router-af)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#nei 150.wg.pop.49 remote-as 650wg2
PEwg2(config-router-af)#nei 150.wg.pop.49 activate
PEwg2(config-router-af)#nei 150.wg.pop.49 as-override
```

Configuration steps on CE*wg*2A:

```
CEwg2A(config)#int s0/0.113 point-to-point
CEwg2A(config-subif)# ip address 150.wg.y.49 255.255.255.240
CEwg2A(config-subif)#frame-relay interface-dlci 113
CEwg2A(config-fr-dlci)#no shut
CEwg2A(config)#router bgp 650wg1
CEwg2A(config-router)#nei 150.wg.wg1.50 remote-as 65001
```

Configuration steps on PE*wg*1:

```
PEwg1(config)#interface s0/0.113 point-to-point
PEwg1(config-subif)#ip vrf forwarding Customer_A
PEwg1(config-subif)#ip address 150.wg.y.50 255.255.255.240
PEwg1(config-subif)#frame-relay interface-dlci 113
```



```
PEwg1(config-fr-dlci)#no shut
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipvr vrf Customer_A
PEwg1(config-router-af)#neighbor 150. wg.y.49 remote-as 650wg1
PEwg1(config-router-af)#neighbor 150. wg.y.49 activate
PEwg1(config-router-af)#neighbor 150. wg.y.49 as-override
```

Task 3: Selecting the Primary and Backup Link with BGP

Configuration steps on CE*wg*1B:

```
CEwg1B(config)#route-map setLP permit 10
CEwg1B(config-route-map)#set local-preference 50
CEwg1B(config-route-map)#route-map setMED permit 10
CEwg1B(config-route-map)#set metric 200
CEwg1B(config-route-map)#router bgp 650wg2
CEwg1B(config-router)#nei 150.wg.pop.50 route-map setLP in
CEwg1B(config-router)#nei 150.wg.pop.50 route-map setMED out
```

Configuration steps on CE*wg*2A:

```
CEwg2A(config)#route-map setLP permit 10
CEwg2A(config-route-map)#set local-preference 50
CEwg2A(config-route-map)#route-map setMED permit 10
CEwg2A(config-route-map)#set metric 200
CEwg2A(config-route-map)#router bgp 650wg1
CEwg2A(config-router)#nei 150.wg.y.50 route-map setLP in
CEwg2A(config-router)#nei 150.wg.y.50 route-map setMED out
```

Lab Exercise 6-1: Overlapping VPNs

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Designing Your VPN Solution

Note No configuration steps are required for this task.

Task 2: Removing CE*wg*1A and CE*wg*2B from Existing VRFs

Configuration steps on PE*wg*1:

```
PEwg1(config)#router bgp 65001
PEwg1(config-router)#address-family ipv4 vrf Customer_A
PEwg1(config-router-af)#no neighbor 150.wg.pop.17
PEwg1(config-vrf)#int s0/0.101
PEwg1(config-subif)#no ip vrf forwarding Customer_A
```

Note After removing the interface from the VRF, the following message will appear:
% Interface Serial0/0.101 IP address 150.wg.*pop*.18 removed due to disabling VRF Customer_A

Configuration steps on PE*wg*2:

```
PEwg2(config)#router bgp 65001
PEwg2(config-router)#address-family ipv4 vrf Customer_B
PEwg2(config-router-af)#no neig 150.wg.pop.33
PEwg2(config-vrf)#int s0/0.102
PEwg2(config-subif)#no ip vrf forwarding Customer_B
```

Note After removing the interface from the VRF, the following message will appear:
% Interface Serial0/0.102 IP address 150.wg.*pop*.34 removed due to disabling VRF Customer_B

Task 3: Configuring New VRFs for CE*wg*1A and CE*wg*2B

Note RDs and RTs listed in the following results may or may not match what you have used in this lab task.

Configuration steps on PE*wg*1:

```
PEwg1(config)#ip vrf A_Central
PEwg1(config-vrf)#rd wg:11
PEwg1(config-vrf)#route-target both wg:10
PEwg1(config-vrf)#route-target both wg:1001
```



```
PEwg1(config-vrf)#int s0/0.101
PEwg1(config-subif)#ip vrf forwarding A_Central
PEwg1(config-subif)#ip address 150.wg.pop.18 255.255.255.240
PEwg1(config)#router bgp 65001
PEwg1(config-router-af)#address-family ipv4 vrf A_Central
PEwg1(config-router-af)#nei 150.wg.pop.17 remote-as 650wg1
PEwg1(config-router-af)#nei 150.wg.pop.17 activate
```

Configuration steps on PE*wg2*:

```
PEwg2(config)#ip vrf B_Central
PEwg2(config-vrf)#rd wg:21
PEwg2(config-vrf)#route-target both wg:20
PEwg2(config-vrf)#route-target both wg:1001
PEwg1(config-vrf)#int s0/0.102
PEwg2(config-subif)#ip vrf forwarding B_Central
PEwg2(config-subif)#ip address 150.wg.pop.34 255.255.255.240
PEwg2(config)#router bgp 65001
PEwg2(config-router-af)#address-family ipv4 vrf B_Central
PEwg2(config-router-af)#nei 150.wg.pop.33 remote-as 650wg2
PEwg2(config-router-af)#nei 150.wg.pop.33 activate
```

Lab Exercise 6-2: Merging Service Providers

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Enabling Connectivity with the Central P Router

Note	The subinterface number and DLCI number in the following configurations will match with each other, and are determined by the instructions for this task.

Configuration steps on P*wg1*:

```
Pwg1(config)#int s0/0.2pop point-to-point
Pwg1(config-subif)#ip address 192.168.100.xx 255.255.255.248
Pwg1(config-subif)#frame-relay interface-dlci 2pop
Pwg1(config-fr-dlci)#no shut
```

Configuration steps on P*wg2*:

```
Pwg2(config)#int s0/0.2pop point-to-point
Pwg2(config-subif)#ip address 192.168.100.xx 255.255.255.248
Pwg2(config-subif)#frame-relay interface-dlci 2pop
Pwg2(config-fr-dlci)#no shut
```

Task 2: Migrating the Core to IS-IS

Configuration steps on PE*wg*1:

```
PE pop(config)#no router eigrp 1
```

Note	Depending on which router has issued the **no router eigrp** command, you will see the following messages appear.

```
*Mar  6 14:59:15.110: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 192.168.wg.65 (Serial0/0.111) is down: interface down

*Mar  6 14:59:15.110: destroy peer: 192.168.wg.65

*Mar  6 14:59:15.110: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 192.168.wg.113 (Serial0/0.112) is down: interface
down

*Mar  6 14:59:15.110: destroy peer: 192.168.wg.113

*Mar  6 14:59:15.110: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 192.168.100.65

(Serial0/0.2pop) is down: interface down

*Mar  6 14:59:15.126: destroy peer: 192.168.100.65


PEwg1(config)#router isis

PEwg1(config-router)#net 49.0001.0000.0000.01wg1.00

PEwg1(config-router)#is level-2-only

PEwg1(config-router)#metric-style wide

PEwg1(config-router)#int s0/0.111

PEwg1(config-subif)#ip router isis

PEwg1(config)#int Loopback0

PEwg1(config-subif)#ip router isis
```

Configuration steps on PE*wg*2:

```
Pwg2(config)#no router eigrp 1
```

Note	Depending on which router has issued the **no router eigrp** command, you will see the following messages appear.

```
*Mar  6 14:59:15.110: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 192.168.wg.65 (Serial0/0.111) is down: interface down

 *Mar  6 14:59:15.110: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 192.168.wg.113 (Serial0/0.112) is down: interface
down


PEwg2(config)#router isis

PEwg2(config-router)#net 49.0001.0000.0000.01wg2.00
```



```
PEwg2(config-router)#is level-2-only
PEwg2(config-router)#metric-style wide
PEwg2(config)#int s0/0.111
PEwg2(config-subif)#ip router isis
PEwg2(config)#int Loopback0
PEwg2(config-subif)#ip router isis
```

Configuration steps on P*wg1*:

```
Pwg1(config)#no router eigrp 1
Pwg1(config)#router isis
Pwg1(config-router)#net 49.0001.0000.0000.02wg1.00
Pwg1(config-router)#is level-2-only
Pwg1(config-router)#metric-style wide
Pwg1(config-router)#int s0/0.111
Pwg1(config-subif)#ip router isis
Pwg1(config-router)#int s0/0.112
Pwg1(config-subif)#ip router isis
Pwg1(config-router)#int s0/0.2wg1
Pwg1(config-subif)#ip router isis
PEwg2(config)#int Loopback0
PEwg2(config-subif)#ip router isis
Pwg1(config)#int Loopback0
Pwg1(config-subif)#ip router isis
```

Configuration steps on P*wg2*:

```
Pwg2(config)#no router eigrp 1
Pwg2(config)#router isis
Pwg2(config-router)#net 49.0001.0000.0000.02wg2.00
Pwg2(config-router)#is level-2-only
Pwg2(config-router)#metric-style wide
Pwg2(config)#int s0/0.111
Pwg2(config-subif)#ip router isis
Pwg1(config-router)#int s0/0.112
Pwg1(config-subif)#ip router isis
Pwg1(config-router)#int s0/0.2wg1
Pwg1(config-subif)#ip router isis
Pwg2(config)#int Loopback0
Pwg2(config-subif)#ip router isis
```

Task 3: Enabling MPLS LDP Connectivity with the Central P Router

Note	The subinterface number and DLCI number in the following configurations will match with each other, and are determined by the instructions for this task.

Configuration steps on P*wg*1:

```
Pwg1(config)#int s0/0.2pop
Pwg1(config-subif)#mpls ip
Pwg1(config-subif)#mpls label protocol ldp
Pwg1(config)#router eigrp 1
Pwg1(config-router)#network 192.168.100.0
Pwg1(config-router)#no auto-sum
```

Configuration steps on P*wg*2:

```
Pwg2(config)#int s0/0.2pop
Pwg2(config-subif)#mpls ip
Pwg2(config-subif)#mpls label protocol ldp
Pwg2(config)#router eigrp 1
Pwg2(config-router)#no auto-sum
Pwg2(config-router)#network 192.168.100.0
```

Task 4: Enabling IBGP Connectivity for All PE Routers

Configuration steps on PE*wg*1:

```
PEwg1(config)#router bgp 65001
PEwg1(config-router)#no neighbor 192.168.wg.33 remote-as 65001
PEwg1(config-router)#nei 192.168.100.129 remote-as 65001
PEwg1(config-router)#nei 192.168.100.129 update-source
loopback0
PEwg1(config-router)#address-family vpnv4
PEwg1(config-router-af)#nei 192.168.100.129 activate
PEwg1(config-router-af)#nei 192.168.100.129 send-community
both
PEwg1(config-router-af)#nei 192.168.100.129  next-hop-self
```

Configuration steps on PE*wg*2:

```
PEwg2(config)#router bgp 65001
PEwg2(config-router)#no nei 192.168.wg.17 remote-as 65001
PEwg2(config-router)#nei 192.168.100.129 remote-as 65001
PEwg2(config-router)#nei 192.168.100.129 update-source
loopback0
PEwg2(config-router)#address-family vpnv4
PEwg2(config-router-af)#nei 192.168.100.129 act
```



```
PEwg2(config-router-af)#nei 192.168.100.129 send-community ext
PEwg2(config-router-af)#nei 192.168.100.129  next-hop-self
```

Lab Exercise 6-3: Common Services VPN

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Designing a Network Management VPN

Note No configuration steps are required for this task.

Task 2: Establishing Connectivity Between the NMS VRF and Other VRFs

Configuration steps on PE*wg*1 for Customer A:

```
PEwg1(config)#ip vrf Customer_A
PEwg1(config-vrf)#export map NMS_Cus_A
PEwg1(config-vrf)#route-target import 101:500
PEwg1(config)#ip vrf A_Central
PEwg1(config-vrf)#export map NMS_Cus_A
PEwg1(config-vrf)#route-target import 101:500
PEwg1(config)#route-map NMS_Cus_A permit 10
PEwg1(config-route-map)#match ip address access-list 10
PEwg1(config route-map)#set extcommunity rt 101:501 add
PEwg1(config) access-list 10 permit host 10.1.41.49
PEwg1(config) access-list 10 permit host 10.1.42.49
```

Configuration steps on PE*wg*2 for Customer A:

```
PEwg2(config)#ip vrf Customer_A
PEwg2(config-vrf)#export map NMS_Cus_A
PEwg2(config-vrf)#route-target import 101:500
PEwg2(config)#route-map NMS_Cus_A permit 10
PEwg2(config-route-map)#match ip address 10
PEwg2(config route-map)#set extcommunity rt 101:501 add
PEwg2(config) access-list 10 permit host 10.1.41.49
PEwg2(config) access-list 10 permit host 10.1.42.49
```

Configuration steps on PE*wg*1 for Customer B:

```
PEwg1(config)#ip vrf Customer_B
PEwg1(config-vrf)#export map NMS_Cus_B
PEwg1(config-vrf)#route-target import 101:500
PEwg1(config)#route-map NMS_Cus_B permit 10
PEwg1(config-route-map)#match ip address 20
PEwg1(config route-map)#set extcommunity rt 101:501 add
PEwg1(config) access-list 20 permit host 10.2.41.49
```

```
PEwg1(config) access-list 20 permit host 10.2.42.49
```

Configuration steps on PE*wg*2 for Customer B:

```
PEwg2(config)#ip vrf Customer_B
PEwg2(config-vrf)#export map NMS_Cus_B
PEwg2(config-vrf)#route-target import 101:500
PEwg2(config)#ip vrf B_Central
PEwg2(config-vrf)#export map NMS_Cus_B
PEwg2(config-vrf)#route-target import 101:500
PEwg2(config)#route-map NMS_Cus_B permit 10
PEwg2(config-route-map)#match ip address 20
PEwg2(config route-map)#set extcommunity rt 101:501 add
PEwg2(config) access-list 20 permit host 10.2.41.49
PEwg2(config) access-list 20 permit host 10.2.42.49
```

Lab Exercise 7-1: Separate Interface for Internet Connectivity

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Establishing CE-PE Connectivity for Internet Access

Configuration steps on CE routers:

```
CE***(config)#int s0/0.114 point-to-point
CE***(config-subif)#ip add 150.wg.pop.66 255.255.255.240
CE***(config-subif)#frame-relay interface-dlci 114
CE***(config-subif)router ospf 1
CE***(config-router)#network 150.wg.0.0 0.0.255.255 area 0
CE***(config-router)#passive-interface s0/0.114
```

Configuration steps on PE routers:

```
PEpop(config)#int s0/0.114 point-to-point
PEpop(config-subif)#ip add 150.wg.pop.65 255.255.255.240
PEpop(config-subif)#frame-relay interface-dlci 114
PEpop(config-subif)#ip router isis
PEpop(config-subif)#router isis
PEpop(config-router)#passive-interface s0/0.114
```

Task 2: Establishing Routing Between the Customer and the Internet

```
Configuration steps on PEwg1 routers:PEwg1(config)#ip route
10.1.wg1.0 255.255.255.0 150.wg.wg1.66
PEwg1(config)#ip route 10.1.wg2.0 255.255.255.0 150.wg.wg1.66
PEwg1(config)#router bgp 65001
PEwg1(config-router)#red static
```


Configuration steps on PE*wg*2 routers:

```
PEwg1(config)#ip route 10.2.wg1.0 255.255.255.0 150.wg.wg2.66
PEwg1(config)#ip route 10.2.wg2.0 255.255.255.0 150.wg.wg2.66
PEwg1(config)#router bgp 65001
PEwg1(config-router)#red static
Configuration steps on CEwg1A router: CE11A(config)#ip route
0.0.0.0 0.0.0.0 s0/0.114
CE11A(config)router bgp 650wg1
CE11A(config-router)#net 0.0.0.0
```

Configuration steps on CEwg2B router:

```
CE11A(config)#ip route 0.0.0.0 0.0.0.0 s0/0.114
CE11A(config)router bgp 650wg2
CE11A(config-router)#net 0.0.0.0
```

Lab Exercise 7-2: Multisite Internet Access

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Establishing CE-PE Connectivity for Internet Access

Configuration steps on CE routers:

```
CE***(config)#int s0/0.115 point-to-point
CE***(config-subif)#ip add 150.wg.pop.130 255.255.255.240
CE***(config-subif)#frame-relay interface-dlci 115
```

Configuration steps on PE routers:

```
PEpop(config)#int s0/0.115 point-to-point
PEpop(config-subif)#ip add 150.wg.pop.129 255.255.255.240
PEpop(config-subif)#frame-relay interface-dlci 115
```

Task 2: Establishing Routing Between the Customer and the Internet

***** Team A ******

Configuration steps on CE*wg*1A router:

```
CEwg1A(config)#router ospf 1
CEwg1A(config-router)#no passive-interface Serial0/0.114
CEwg1A(config-router)#no network 150.wg.0.0 0.0.255.255 area 0
CEwg1A(config-router)#router bgp 650wg1
CEwg1A(config-router)#no net 0.0.0.0
CEwg1A(config-router)#nei 150.wg.wg1.65 remote 65001
CEwg1A(config-router)#no ip route 0.0.0.0 0.0.0.0
Serial0/0.114
```

Configuration steps on PE*wg*1 routers:

```
PEwg1(config)#no ip route 10.1.wg1.0 255.255.255.0
150.wg.wg1.66
PEwg1(config)#no ip route 10.1.wg2.0 255.255.255.0
150.wg.wg1.66
PEwg1(config)#router bgp 65001
PEwg1(config-router)#nei 150.wg.wg1.66 remote 650wg1
```

Configuration steps on CE*wg*2A router:

```
CEwg2A(config)#router bgp 650wg1
CEwg2A(config-router)#nei 150.wg.wg2.129 remote 65001
```

Configuration steps on PE*wg*2 routers:

```
PEwg2(config)#router bgp 65001
PEwg2(config-router)#nei 150.wg.wg2.130 remote 650wg1
```

***** Team B ******

Configuration steps on CE*wg*2B router:

```
CEwg2B(config)#router ospf 1
CEwg2B(config-router)#no passive-interface Serial0/0.114
CEwg2B(config-router)#no network 150.wg.0.0 0.0.255.255 area 0
CEwg2B(config-router)#router bgp 650wg2
CEwg2B(config-router)#no net 0.0.0.0
CEwg2B(config-router)#nei 150.wg.wg2.65 remote 65001
CEwg2B(config-router)#no ip route 0.0.0.0 0.0.0.0
Serial0/0.114
Configuration steps on PEwg2 routers: PEwg2(config)#no ip
route 10.2.wg1.0 255.255.255.0 150.wg.wg2.66
PEwg2(config)#no ip route 10.2.wg2.0 255.255.255.0
150.wg.wg2.66
PEwg2(config)#router bgp 65001
PEwg2(config-router)#nei 150.wg.wg2.66 remote 650wg2
Configuration steps on CEwg1B router: CEwg1B(config)#router
bgp 650wg1
CEwg1B(config-router)#nei 150.wg.wg1.129 remote 65001
Configuration steps on PEwg1 routers: PEwg1(config)#router bgp
65001
PEwg1(config-router)#nei 150.wg.wg1.130 remote 650wg2
```


Lab Exercise 7-3: Internet Connectivity in an MPLS VPN

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Establishing Central Site Connectivity for Internet Access

Configuration steps on PE routers:

****** Team A **********

```
PEwg1(config)#ip vrf Internet
PEwg1(config-vrf)#route-target 100:600
PEwg1(config-vrf)#rd 100:600
PEwg1(config)#int s0/0.114
PEwg1(config-subif)#ip vrf forwarding Internet
% Interface Serial0/0.114 IP address 150.wg.wg1.65 removed due
to enabling VRF Internet
PEwg1(config-subif)#ip add 150.wg.wg1.65 255.255.255.240
PEwg1(config)#router bgp 65001
PEwg1(config-router)#no neighbor 150.wg.wg1.66 remote-as
650wg1
PEwg1(config-router)#address-family ipv4 vrf Internet
PEwg1(config-router-af)#nei 150.wg.wg1.66 remote 650wg1
PEwg1(config-router-af)#nie 150.wg.wg1.66 activate
PEwg1(config-router-af)#nei 150.wg.wg1.66 activate
```

****** Team B **********

```
PEwg2(config)#ip vrf Internet
PEwg2(config-vrf)#route-target 100:600
PEwg2(config-vrf)#rd 100:600
PEwg2(config)#int s0/0.114
PEwg2(config-subif)#ip vrf forwarding Internet
% Interface Serial0/0.114 IP address 150.wg.wg2.65 removed due
to enabling VRF Internet
PEwg2(config-subif)#ip add 150.wg.wg2.65 255.255.255.240
PEwg2(config)#router bgp 65001
PEwg2(config-router)#no neighbor 150.wg.wg2.66 remote-as
650wg2
PEwg2(config-router)#address-family ipv4 vrf Internet
PEwg2(config-router-af)#nei 150.wg.wg2.66 remote 650wg2
PEwg2(config-router-af)#nie 150.wg.wg2.66 activate
PEwg2(config-router-af)#nei 150.wg.wg2.66 activate
```

Task 2: Establishing Remote Site CE-PE Connectivity for Internet Access

***** Team A ******

Configuration steps on PE*wg2* routers: PE*wg2*(config-vrf)#int s0/0.115

PE*wg2*(config-subif)#ip vrf forward Internet

% Interface Serial0/0.115 IP address 150.wg.wg2.129 removed due to enabling VRF Internet

PE*wg2*(config-subif)#ip add 150.wg.wg2.129 255.255.255.240

PE*wg2*(config-subif)#router bgp 65001

PE*wg2*(config-router)#no nei 150.wg.wg2.130

PE*wg2*(config-router)#address-family ipv4 vrf Internet

PE*wg2*(config-router-af)#nei 150.wg.wg2.130 remote 650wg1

PE*wg2*(config-router-af)#nei 150.wg.wg2.130 act

***** Team B ******

Configuration steps on PE*wg1* routers: PE*wg1*(config-vrf)#int s0/0.115

PE*wg1*(config-subif)#ip vrf forward Internet

% Interface Serial0/0.115 IP address 150.wg.12.129 removed due to enabling VRF Internet

PE*wg1*(config-subif)#ip add 150.wg.wg1.129 255.255.255.240

PE*wg1*(config-subif)#router bgp 65001

PE*wg1*(config-router)#no nei 150.wg.wg1.130

PE*wg1*(config-router)#address-family ipv4 vrf Internet

PE*wg1*(config-router-af)#nei 150.wg.wg1.130 remote 650wg2

PE*wg1*(config-router-af)#nei 150.wg.wg1.130 act

