MPLS VPN Technology

Overview

This module introduces Virtual Private Networks (VPN) and two major VPN design options – overlay VPN and peer-to-peer VPN. VPN terminology and topologies are introduced.

The module then describes MPLS VPN architecture, operations and terminology. It details CE-PE routing from various perspectives and BGP extensions (route targets, and extended community attributes) that allow I-BGP to transport customer routes over a provider network. The MPLS VPN forwarding model is also covered together with its integration with core routing protocols.

Upon completion of this module, the learner will be able to perform the following tasks:

- Identify major Virtual Private network topologies, their characteristics and usage scenarios
- Describe the differences between overlay VPN and peer-to-peer VPN
- List major technologies supporting overlay VPNs and peer-to-peer VPNs
- Position MPLS VPN in comparison with other peer-to-peer VPN implementations
- Describe major architectural blocks of MPLS VPN
- Describe MPLS VPN routing model and packet forwarding
- Describe the model for MPLS VPNs to span more than one autonomous system

Outline

This module contains these lessons:

- Introduction to Virtual Private Networks
- Overlay and Peer-to-Peer VPN

- Major VPN Topologies
- MPLS VPN Architecture
- MPLS VPN Routing Model
- MPLS VPN Packet Forwarding
- MPLS VPN Spanning more than One AS

Introduction to Virtual Private Networks

Overview

This lesson describes the concept of VPN and introduces some VPN terminology.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the concept of VPN
- Explain VPN terminology as defined by MPLS VPN architecture

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- Virtual Private Network Concept
- Business-Needs Based VPN Classification
- VPN Terminology as Used in MPLS VPN Architecture
- Summary

Virtual Private Network Concept

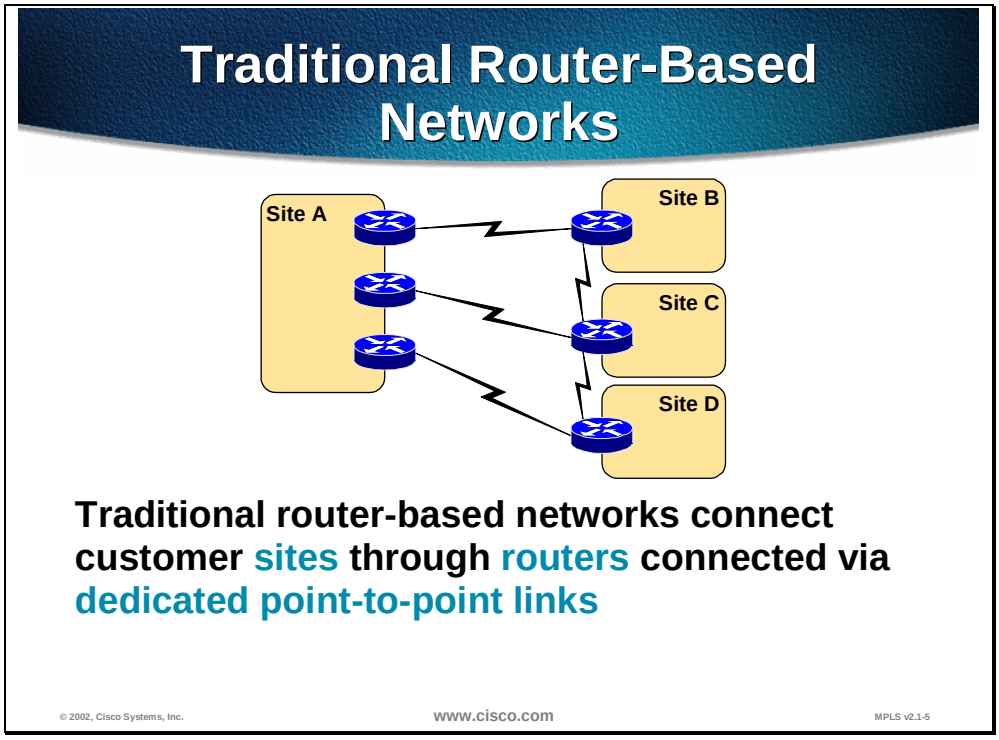

Traditional router-based networks were implemented with dedicated point-to-point links connecting customer sites. The cost of such an approach was comparatively high for a number of reasons:

- 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-customer
- Every link required a dedicated port on a router, resulting in high equipment costs.

Practice

Q1) What were traditional router-based networks implemented with?

A) Point-to-multipoint links.

B) Virtual circuits.

C) Emulated point-to-point links.

D) Dedicated point-to-point links.

Business-Needs Based VPN Classification

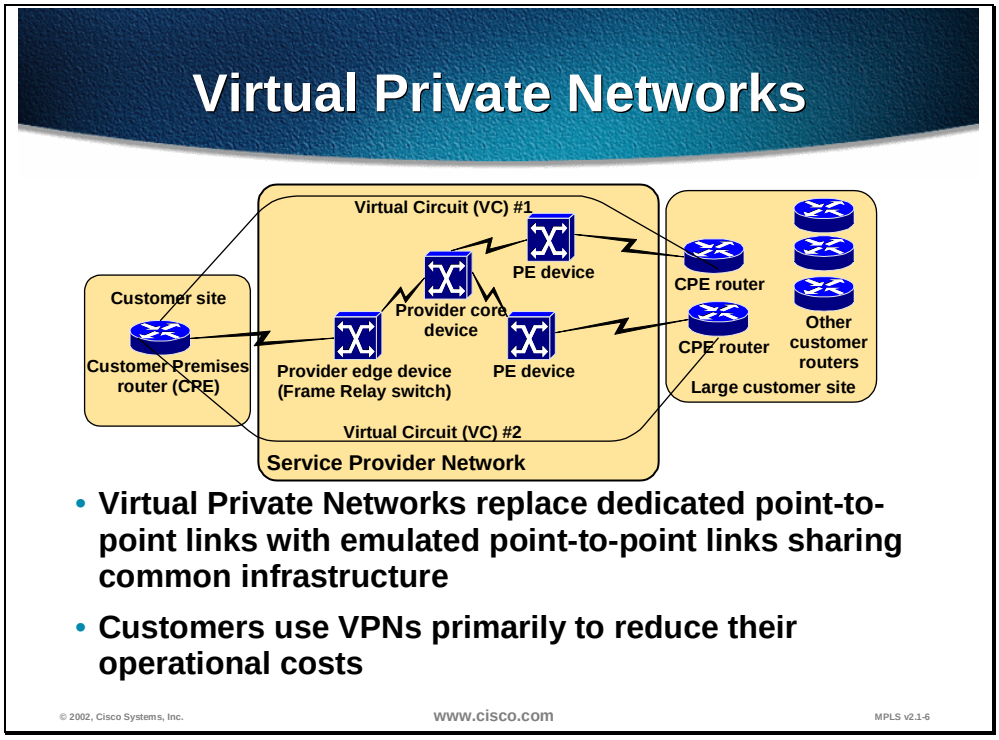

Virtual Private Networks (VPNs) were introduced very early in the history of data communications with technologies like X.25 and Frame Relay, which use virtual circuits to establish the end-to-end connection over a shared service provider infrastructure. These technologies, although sometimes considered legacy and obsolete, still share the basic business assumptions with the modern VPN approaches:

- The dedicated links are replaced with common infrastructure that emulates point-to-point links for the customer, resulting in statistical sharing of Service Provider infrastructure
- Statistical sharing of infrastructure enables the service provider to offer the connectivity for lower price, resulting in lower operational costs for the end customers.

The statistical sharing is illustrated in the graphic, where you can see the CPE router on the left has one physical connection to the service provider with two virtual circuits provisioned. Virtual Circuit 1 (VC # 1) provides connectivity to the top CPE router on the right. Virtual Circuit 2 (VC #2) provides the connectivity to the bottom CPE router on the right.

Practice

Q1) Why are customers interested in Virtual Private Networks?

A) VPNs use point-to-multipoint links which enables more connections for customers.

B) VPNs reduce customers' connectivity costs.

C) VPNs are easiest to configure.

VPN Terminology as Used in MPLS VPN Architecture

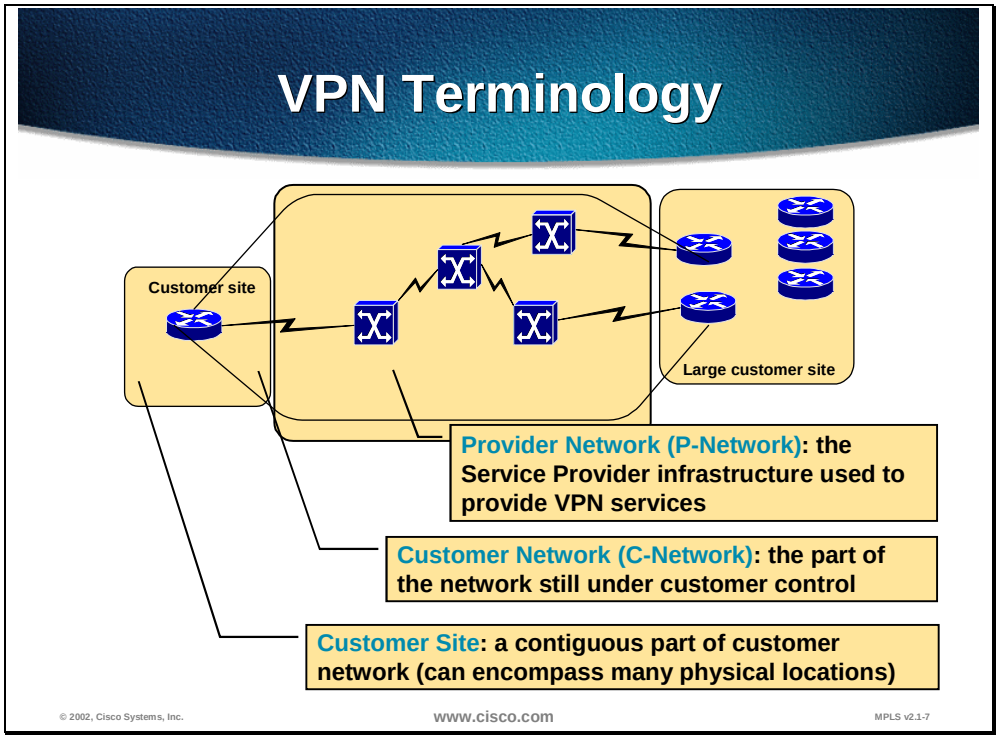

There are many conceptual models and terminologies describing various Virtual Private Network technologies and implementations. In this lesson we'll focus on the terminology introduced by MPLS VPN architecture. As you'll see, the terminology is generic enough to cover any VPN technology or implementation and is thus extremely versatile.

The major parts of an overall VPN solution are always:

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

A typical customer network implemented with any VPN technology would contain islands of connectivity completely under customer control (customer **sites**) connected together via the Service Provider infrastructure (**P-network**).

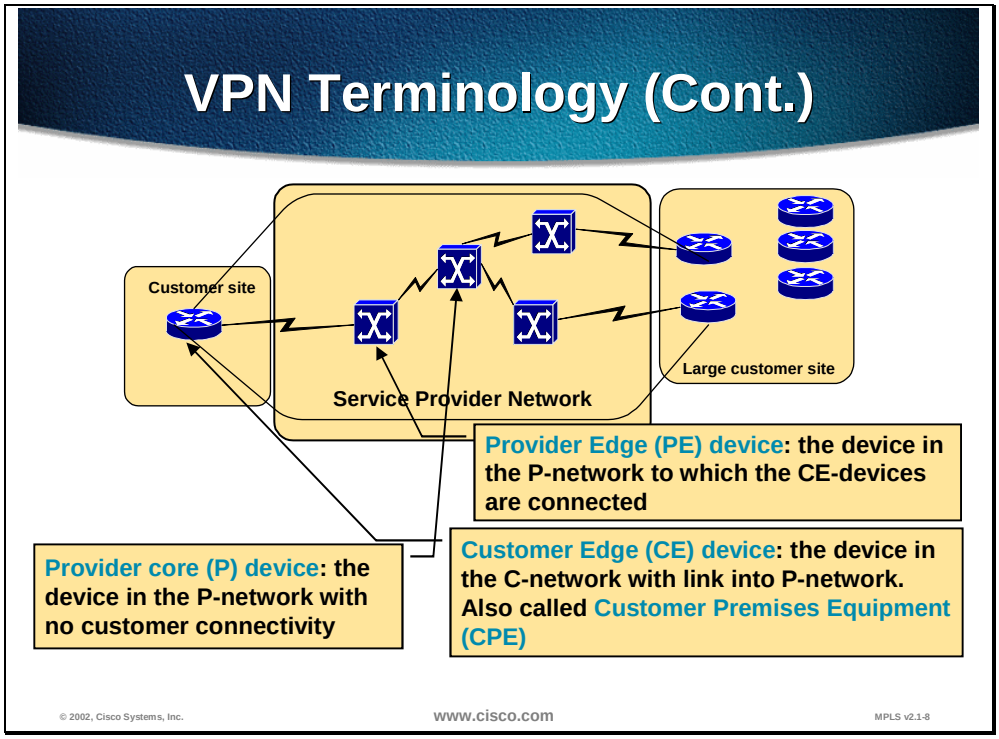

The devices that enable the overall VPN solution are named based on their position in the network:

- Customer router that connected the customer site to the Service Provider network is called a Customer Edge router (**CE-router**). Traditionally this device is called Customer Premises Equipment (**CPE**).

Note	If the CE device is not a router, but, for example, a Packet Assembly and Disassembly (PAD) device, we can still use a generic term CE-device.

- Service Provider devices where the customer devices are attached are called Provider Edge (**PE**) devices. In traditional switched Wide Area Network (WAN) implementations, these devices would be Frame Relay or X.25 edge switches.
- Service Provider devices that only provide 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.

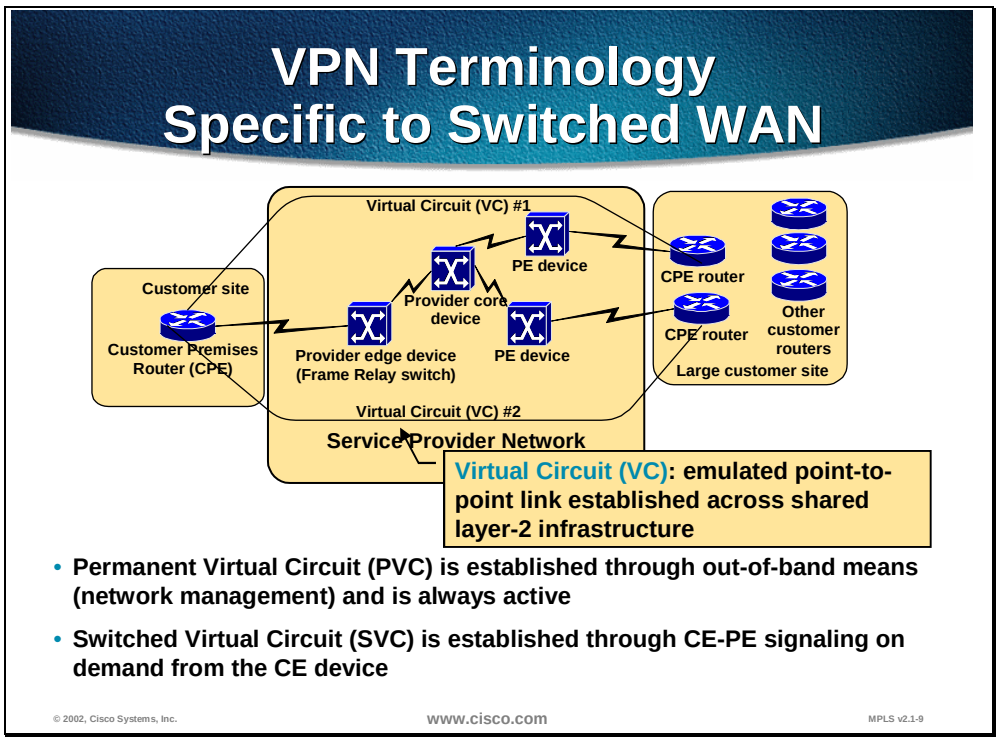

Switched WAN technologies introduced a term **Virtual Circuit (VC)**, which is an emulated point-to-point link established across layer-2 infrastructure (for example, Frame Relay network). The virtual circuits are further differentiated into **Permanent Virtual Circuits (PVC)** which are pre-established by means of network management or manual configuration and **Switched Virtual Circuits (SVC)** which are established on demand through a call setup request from the CE device.

Practice

Q1) What is a customer site?

A) It is a collection of routers and networks that constitute customer's hub.

B) It is a router on customer's premises that connects to the MPLS/VPN backbone.

C) It is an interface on a PE router that connects a specific customer.

D) It is a contiguous part of the C-network.

Summary

After completing this lesson, you should be able to perform the following tasks:

- Describe the concept of VPN
- Explain VPN terminology as defined by MPLS VPN architecture

Next Steps

After completing this lesson, go to:

- Overlay and Peer-to-Peer VPN

Lesson Review

Instructions

Answer the following questions:

1. Why are customers interested in Virtual Private Networks?
2. What is the main role of a VPN?
3. What is a C-network?
4. What is a customer site?
5. What is a CE-router?
6. What is a P-network?
7. What is the difference between a PE-device and a P-device?

Overlay and Peer-to-Peer VPN

Overview

The lesson describes two major VPN implementation options and identifies major differences between them.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the differences between overlay and peer-to-peer VPN
- Describe the benefits and drawbacks of each VPN implementation option
- List major technologies supporting overlay VPNs
- Describe traditional peer-to-peer VPN implementation options

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- Overlay VPN Implementation
- Technologies Supporting Overlay VPN
- Peer-to-Peer VPN Concept
- Peer-to-Peer VPN Implemented with IP Packet Filters
- Peer-to-Peer VPN Implemented with Controlled Route Distribution
- Benefits and Drawbacks of VPN Implementation Options
- Summary

Overlay VPN Implementation

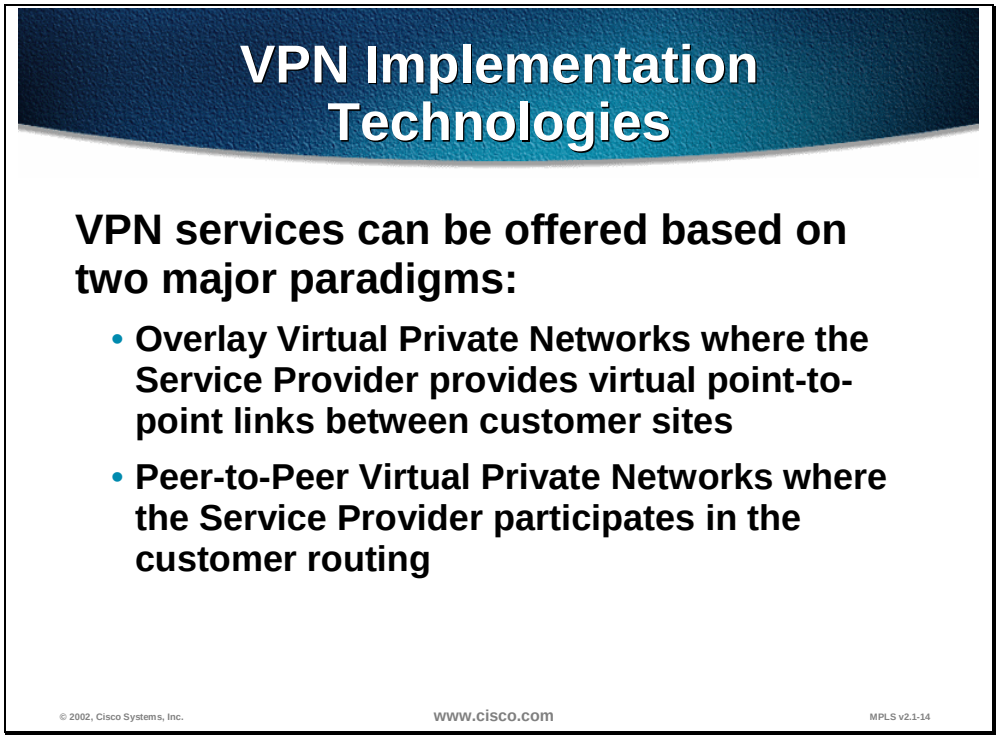

Traditional VPN implementations were all based on the **overlay** paradigm – the Service Provider sells **virtual circuits** between customer sites as a replacement for dedicated point-to-point links. The overlay paradigm has a number of drawbacks that will be identified in this lesson. To overcome these drawbacks (particularly in IP-based customer networks), a new paradigm called **peer-to-peer** VPN was introduced where the Service Provider actively participates in customer routing.

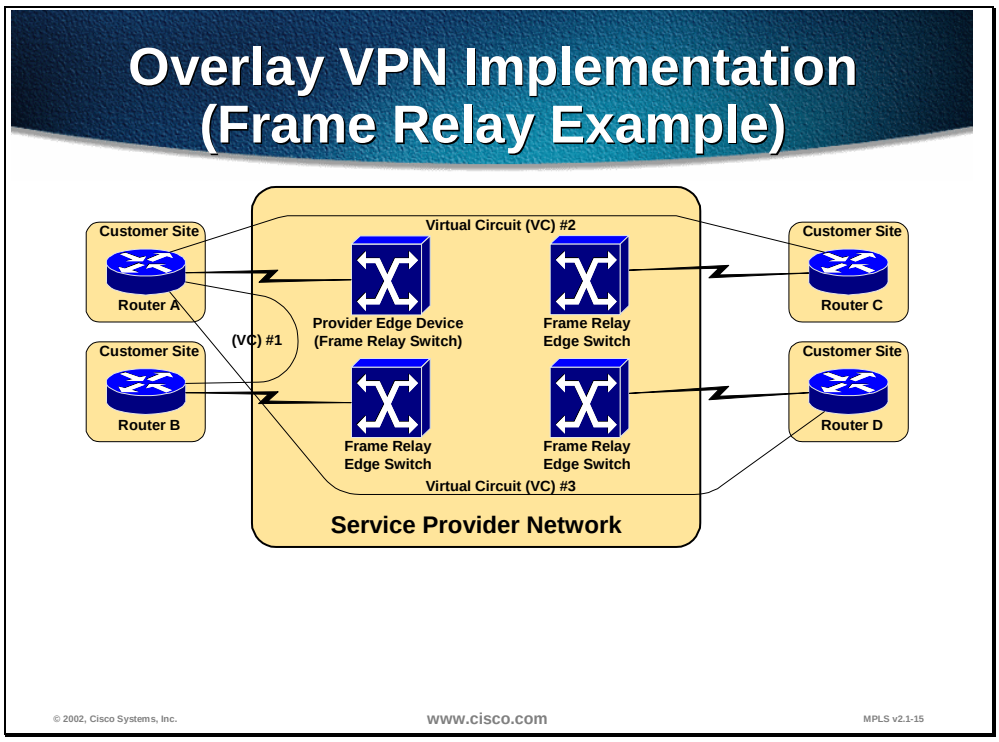

The diagram above shows a typical overlay VPN, implemented by a Frame Relay network. The customer needs to connect three sites (site Alpha being the central site – the hub) and orders connectivity between Alpha (Hub) and Beta (Spoke) and between Alpha (Hub) and Gamma (Spoke). The Service Provider implements this request by providing two PVCs across the Frame Relay network.

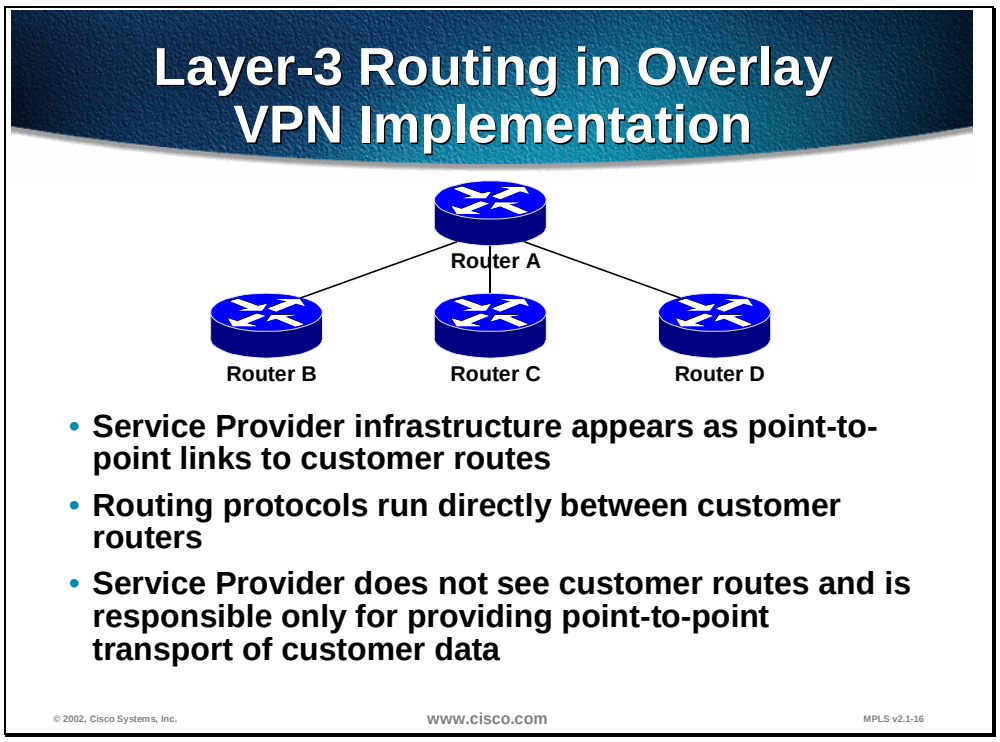

From the layer-3 perspective, the Service Provider network is invisible – the customer routers are linked with emulated point-to-point links. The routing protocol is run directly between customer routers that establish routing adjacencies and exchange routing information.

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.

Practice

Q1) What is an overlay VPN?

A) It is a VPN based on IPsec.

B) It is a VPN providing virtual circuits or emulated point-to-point links/tunnels between the customers routers.

C) It is a VPN providing IP over IP connectivity to customers.

D) It is a VPN providing Frame Relay connectivity between PE-devices.

Technologies Supporting Overlay VPN

There are a number of different overlay VPN implementations, ranging from traditional Time Division Multiplexing (TDM) to highly complex technologies running across IP backbones. In the following slides, we'll introduce major VPN technologies and implementations.

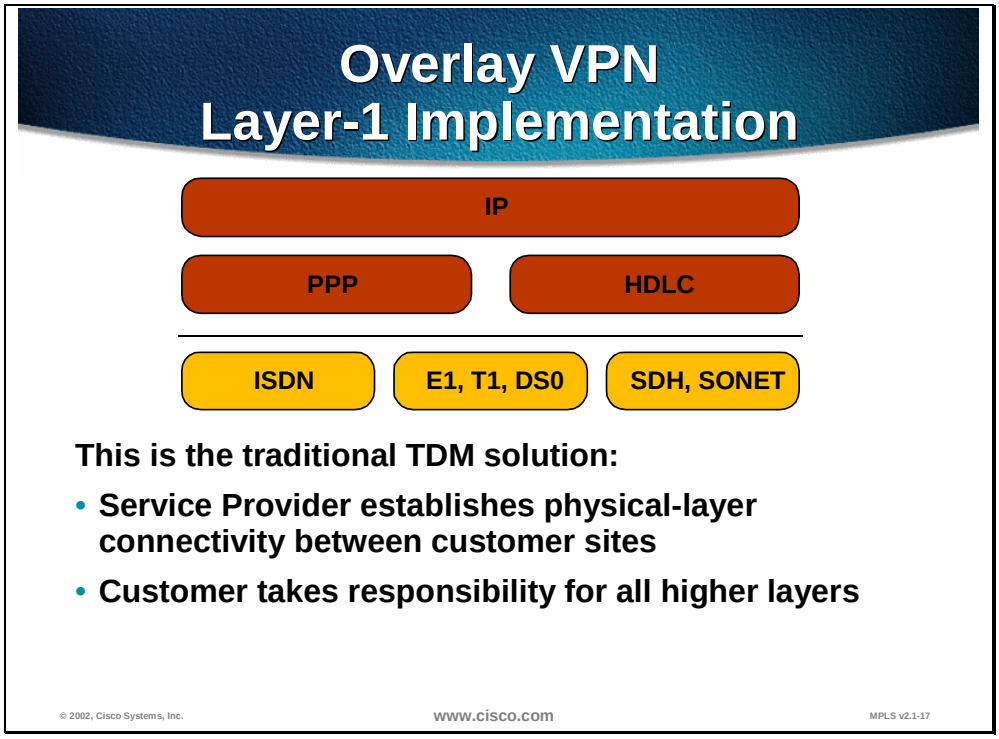

In layer-1 overlay VPN implementation, the Service Provider sells layer-1 circuits (bit pipes) implemented with technologies like ISDN, DS0, E1, T1, SDH or SONET. The customer takes responsibility for layer-2 encapsulation between customer devices and the transport of IP data across the infrastructure.

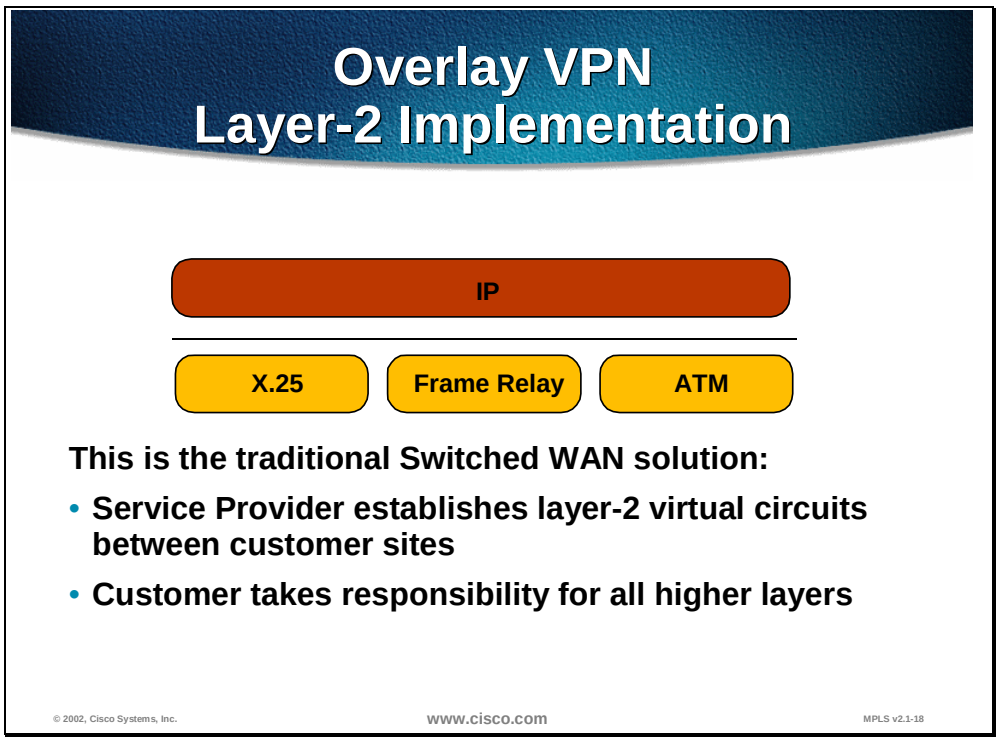

Layer-2 VPN implementation is the traditional switched WAN model, implemented with technologies like X.25, Frame Relay, ATM or SMDS. The Service Provider is responsible for transport of layer-2 frames between customer sites and the customer takes responsibility for all higher layers.

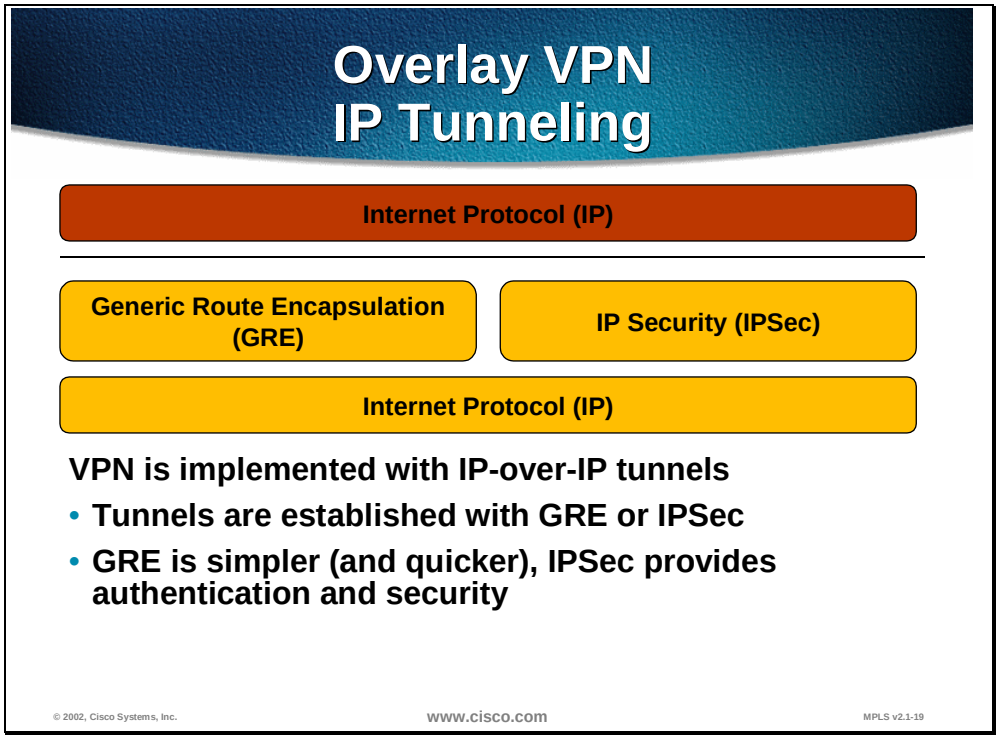

With the success of Internet Protocol (IP) and associated technologies, some Service Providers started to implement pure IP backbones to offer VPN services based on IP. In other cases, the customers want to take advantage of low cost and universal availability of Internet to build low-cost private networks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

Whatever the business reasons behind it, overlay Layer 3 VPN implementation over IP backbone always involves **tunneling** (encapsulation of protocol units at a certain layer of OSI model into protocol units at the same or higher layer of OSI model).

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

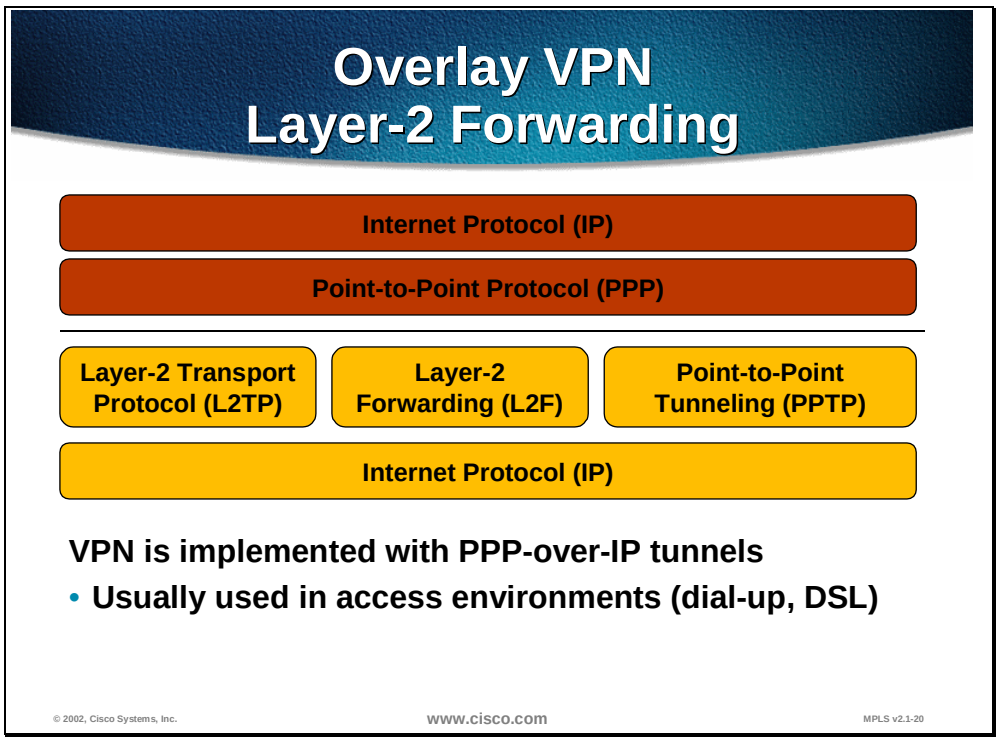

Yet another tunneling technique that was first implemented in dial-up networks, where the Service Providers wanted to tunnel customer dial-up data encapsulated in point-to-point protocol (PPP) frames over an IP backbone to the customer's central site. To make the Service Provider transport transparent to the customer, PPP frames are exchanged between the customer sites (usually a dial-up 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 (L2F)
- Layer 2 Transport Protocol (L2TP)
- Point-to-Point Tunneling Protocol (PPTP)

Practice

Q1) Which of the following are IP-based overlay VPN technologies? (Select all that apply)

A) IP Security (IPsec)

B) Generic Route Encapsulation (GRE) tunnels

C) Internet Protocol (IP)

D) Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH)

E) Integrated Services Digital Network (ISDN)

F) PPP forwarding

Peer-to-Peer VPN Concept

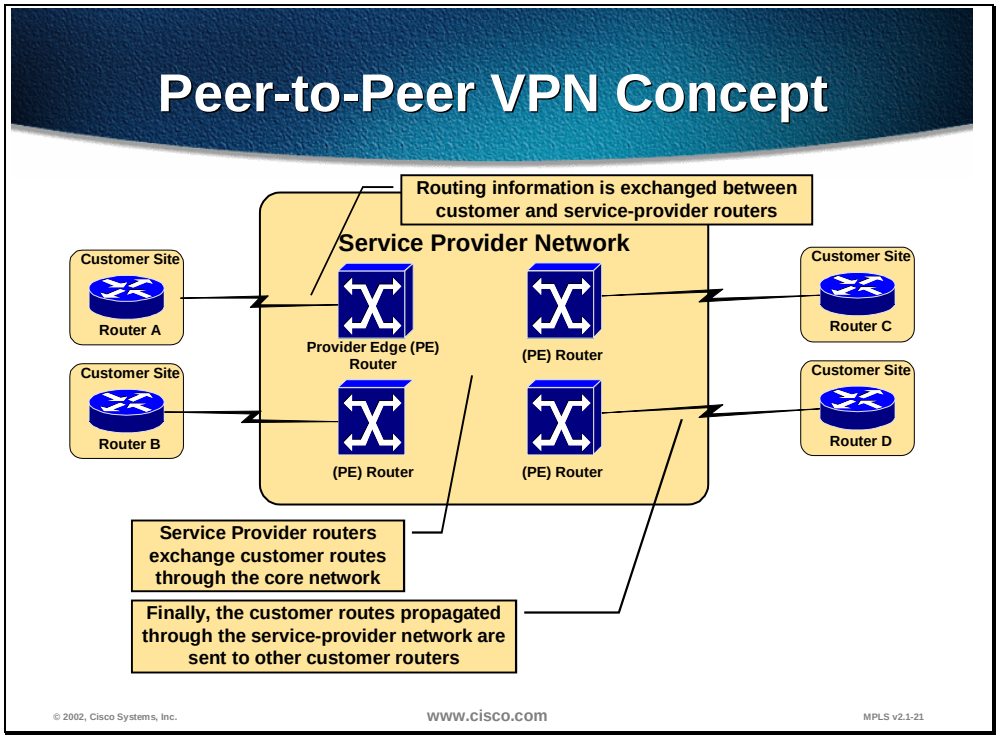

Overlay VPN paradigm has a number of drawbacks, most significant of them being the need for the customer to establish point-to-point links or virtual circuits between sites. The formula to calculate how many point-to-point links or virtual circuits you need in the worst case is ((n)(n-1))/2, where n is the number of sites you need to connect. For example, if you need to have full–mesh connectivity between 4 sites, you will need a total of 6 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 where the Service Provider actively participates in the customer routing, accepting customer routes, transporting them across the Service Provider backbone and finally propagating them to other customer sites.

Practice

Q1) What is the major benefit of peer-to-peer VPN as compared to overlay VPN?

A) P2P VPNs are more or less equal to overlay VPNs.

B) P2P VPNs are faster than overlay VPNs.

C) P2P VPNs are more scalable than overlay VPNs.

D) P2P VPNs guarantee optimum routing between sites without the need for full-mesh of VCs.

Peer-to-Peer VPN Implemented with IP Packet Filters

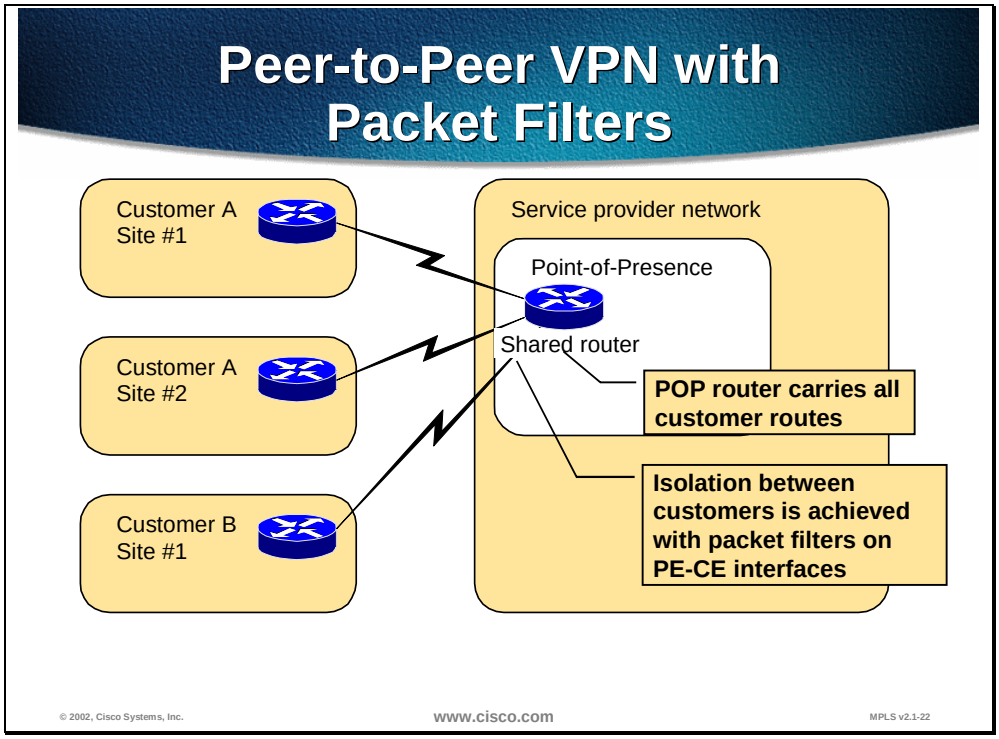

The first peer-to-peer VPN solutions appeared several years ago. Architectures similar to the Internet were used to build them and special provisions had to be taken in account to transform the architecture, which was targeted toward public backbones (Internet) into a solution where the customers would be totally isolated and able to exchange their corporate data securely.

The more common peer-to-peer VPN implementation uses packet filters on the PE-routers to isolate the customers. 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.

Peer-to-Peer VPN Implemented with Controlled Route Distribution

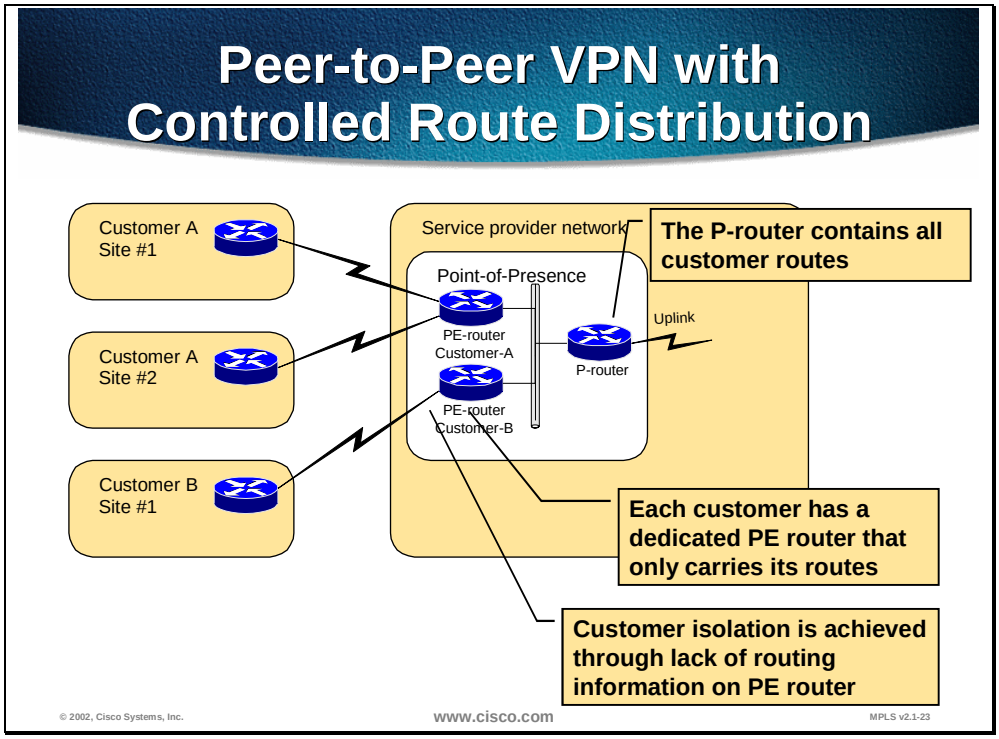

Maintaining packet filters is a mundane and error-prone task. Some Service Providers thus implemented more innovative solutions based on controlled route distribution. In this approach, the core Service Provider routers (the P-routers) would contain all customer routes and the PE-routers would only contain routes of a single customer, requiring a dedicated PE-router per customer per Point-of-Presence (POP). The customer isolation is achieved solely through lack of routing information on the PE-router. Using route filtering between the P-router and the PE-routers, the PE-router for Customer A will only learn routes belonging to Customer A, and the PE-router for Customer B will only learn routes belonging to Customer B. Border Gateway Protocol (BGP) with BGP communities is usually used inside the Provider backbone since it offers the most versatile route filtering tools.

Note Default routes used anywhere in the customer or Service Provider network break isolation between the customers and have to be avoided.

Practice

Q1) Where is Border Gateway Protocol (BGP) with BGP communities usually used?

A) Inside the customer network.

B) Inside the Provider backbone.

C) Outside the Provider backbone.

D) Between the Provider backbone and the customer network.

Benefits and Drawbacks of VPN Implementation Options

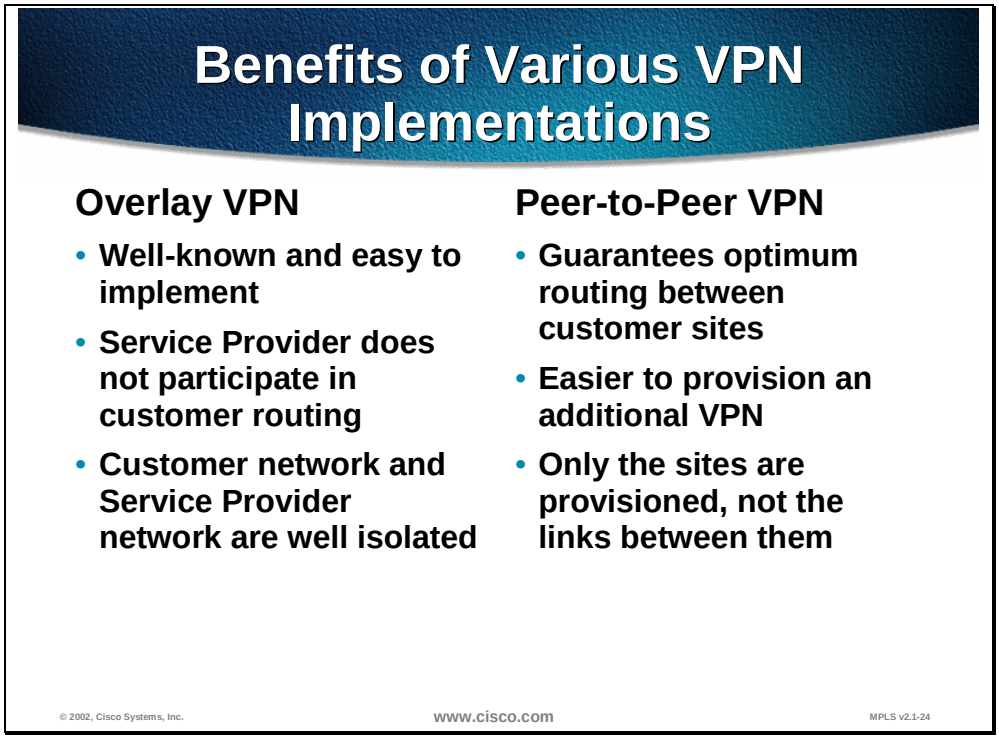

Each VPN paradigm has a number of benefits:

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

On the other hand, the peer-to-peer VPN give you:

- Optimum routing between customer sites without any special design or configuration effort
- Easy provisioning of additional VPNs or customer sites, as the Service Provider only needs to provision individual sites, not the links between individual customer sites.

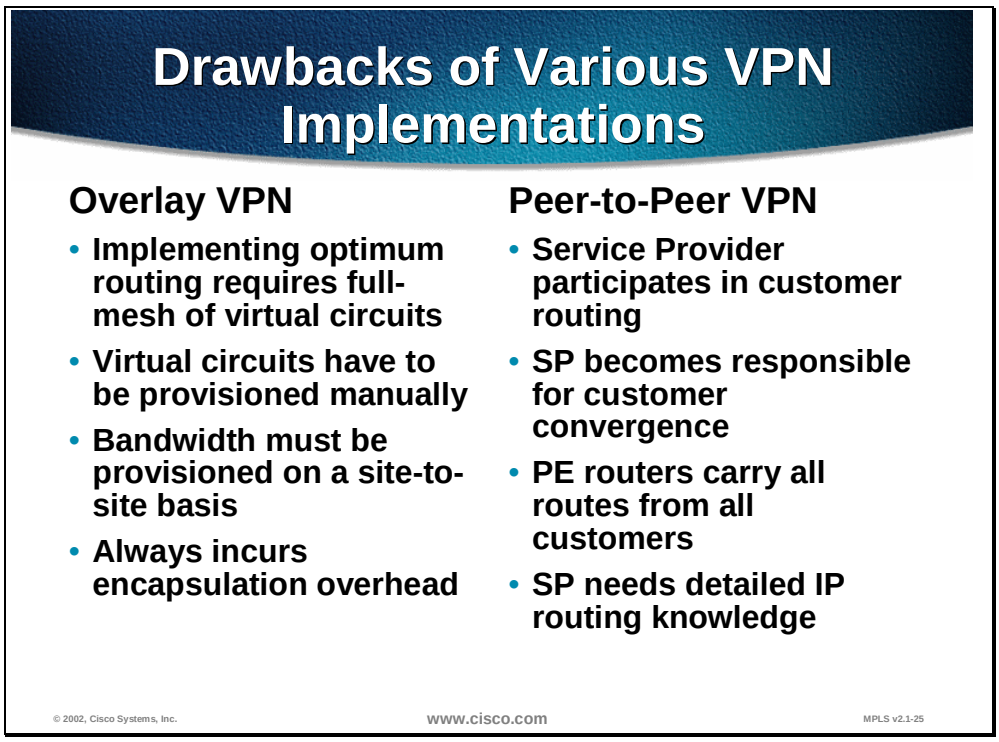

Each VPN paradigm also has a number of drawbacks:

- Overlay VPNs require a full mesh of virtual circuit between customer sites to provide optimum inter-site routing
- All the virtual circuits between customer sites in an overlay VPN 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) also incur high encapsulation overhead (ranging from 20 to 80 bytes per transported datagram).

The major drawbacks of peer-to-peer VPN arise from the Service Provider's involvement in customer routing:

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

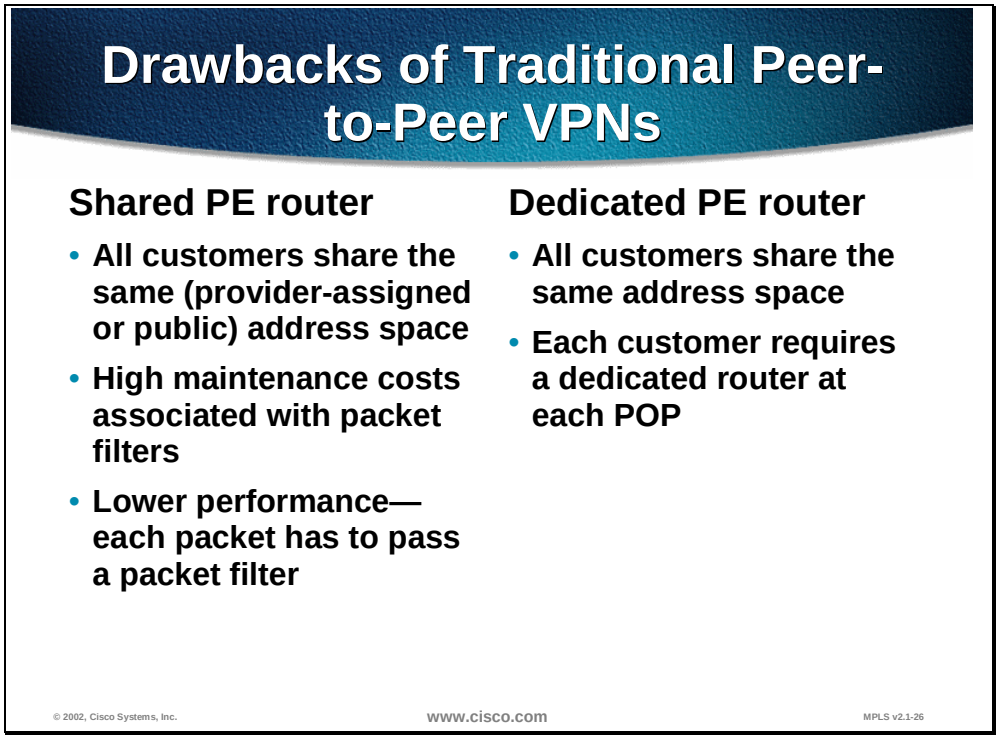

The pre-MPLS VPN implementations of peer-to-peer VPNs all shared a common drawback – the customers have to share the same address space, either using public IP addresses in their private networks or relying on service provider-assigned IP addresses. In both cases, connecting a new customer to a peer-to-peer VPN service usually requires IP renumbering inside the customer network – an operation, which most customers are reluctant to perform.

The peer-to-peer VPNs based on packet filters also incur high operational costs associated with packet filter maintenance as well as performance degradation due to heavy usage of packet filters.

The peer-to-peer VPNs implemented with per-customer PE-routers are easier to maintain and can give you optimum routing performance, but are usually more expensive since every customer requires a dedicated router in every POP. This approach is thus usually used in scenarios where the Service Provider only provides service to a small number of large customers.

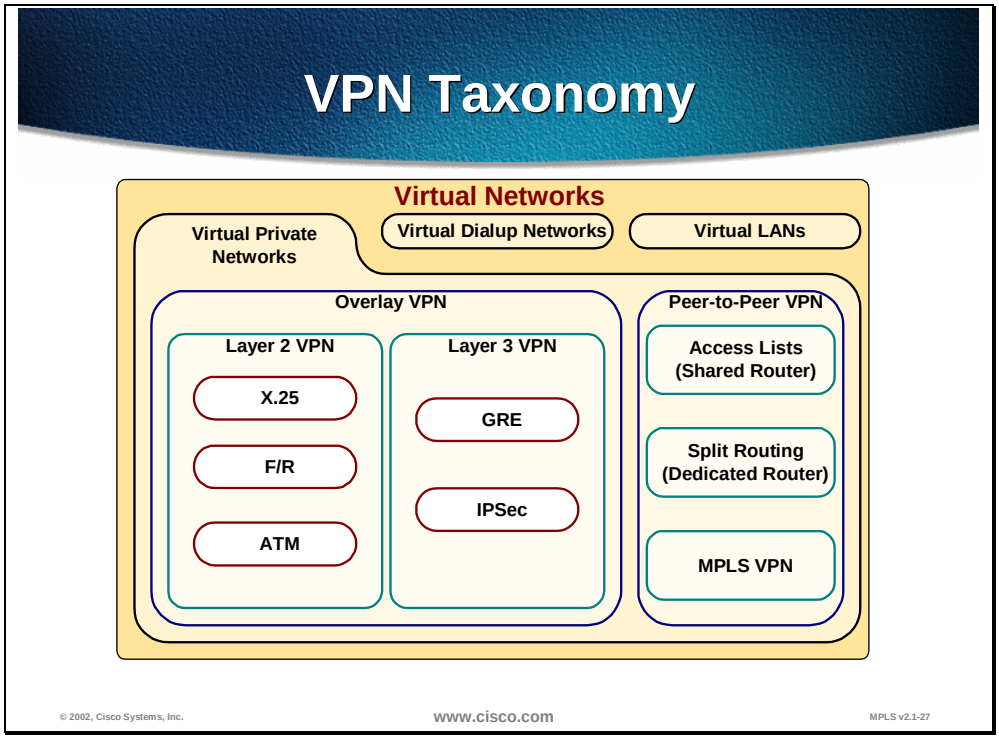

Practice

Q1) What is the drawback of all traditional peer-to-peer VPN implementations?

A) Operational costs are high

B) Maintenance is not as easy as it is with overlay VPN

C) The customers cannot use overlapping IP address-spaces

Summary

After completing this lesson, you should be able to perform the following tasks:

- Describe the differences between overlay and peer-to-peer VPN
- Describe the benefits and drawbacks of each VPN implementation option
- List major technologies supporting overlay VPNs
- Describe traditional peer-to-peer VPN implementation options

Next Steps

After completing this lesson, go to:

- Major VPN Topologies

Lesson Review

Instructions

Answer the following questions:

1. What is an overlay VPN?
2. Which routing protocol runs between the customer and the service provider in an overlay VPN?
3. Which routers are routing protocol neighbors of a CE-router in overlay VPN?
4. List three IP-based overlay VPN technologies.
5. What is the major benefit of peer-to-peer VPN as compared to overlay VPN?
6. List two traditional peer-to-peer VPN implementations.
7. What is the drawback of all traditional peer-to-peer VPN implementations?

Major VPN Topologies

Overview

This lesson defines major VPN topologies used today and gives usage guidelines for each category.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Identify major VPN topologies
- Describe the implications of using overlay VPN or peer-to-peer VPN approach with each topology
- List sample usage scenarios for each topology

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- Hub-and-Spoke Overlay VPN Topology
- Partial-Mesh or Full-Mesh Overlay VPN Topology
- Simple Extranet Topology
- Central-Services Extranet
- Managed Network VPN Topology
- Summary

Hub-and-Spoke Overlay VPN Topology

VPN Categorizations

There are three major VPN categorizations:

- Topology categorization, which only applies to overlay VPNs
- Business categorization, which categorizes VPNs based on the business needs they fulfill
- Connectivity categorization, which classifies VPNs based on their connectivity requirements.

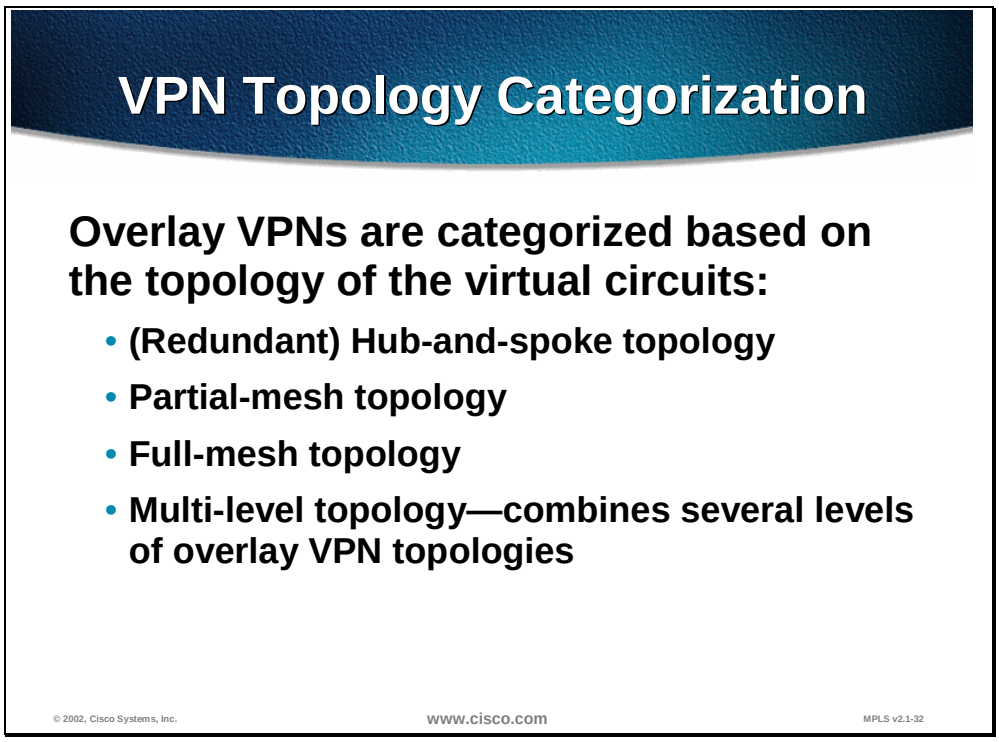

The oldest VPN categorization was based on the topology of point-to-point links in an overlay VPN implementation:

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

Large networks usually deploy a layered combination of these technologies, for example:

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

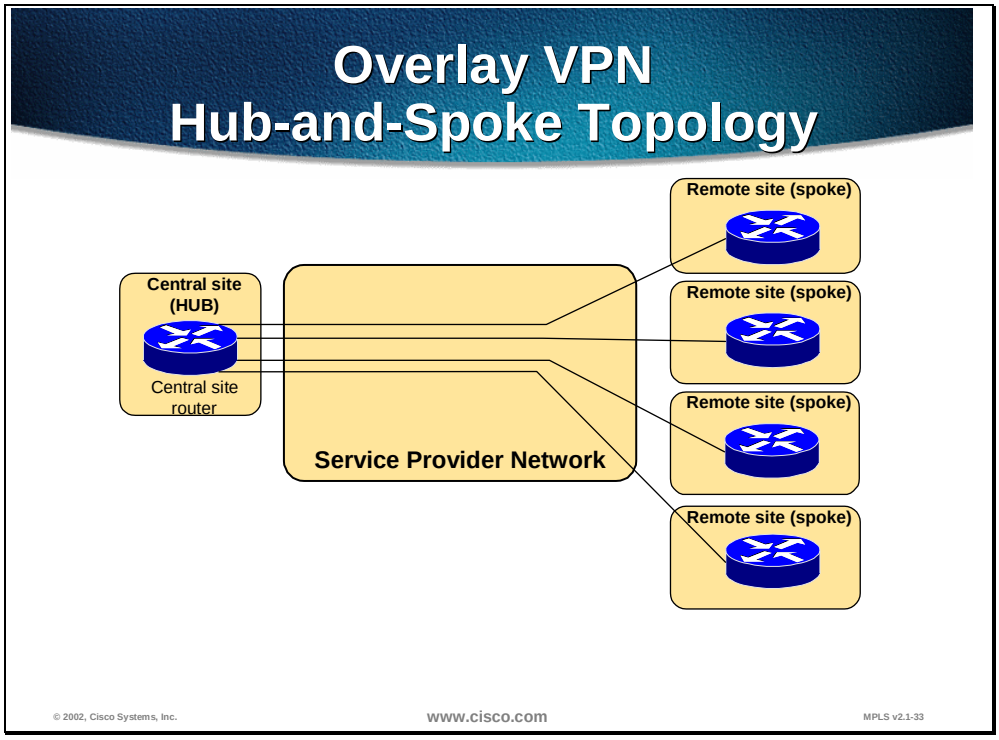

The hub-and-spoke topology is the simplest overlay VPN topology – all remote sites are linked with a single virtual circuit to a central CE-router. The routing is also extremely simple – static routing or distance-vector protocol like RIP are more than adequate. If you are using dynamic routing protocol like RIP, split-horizon must be disabled at the hub router, or you must use point-to-point sub-interfaces at the hub router to overcome the split-horizon problem.

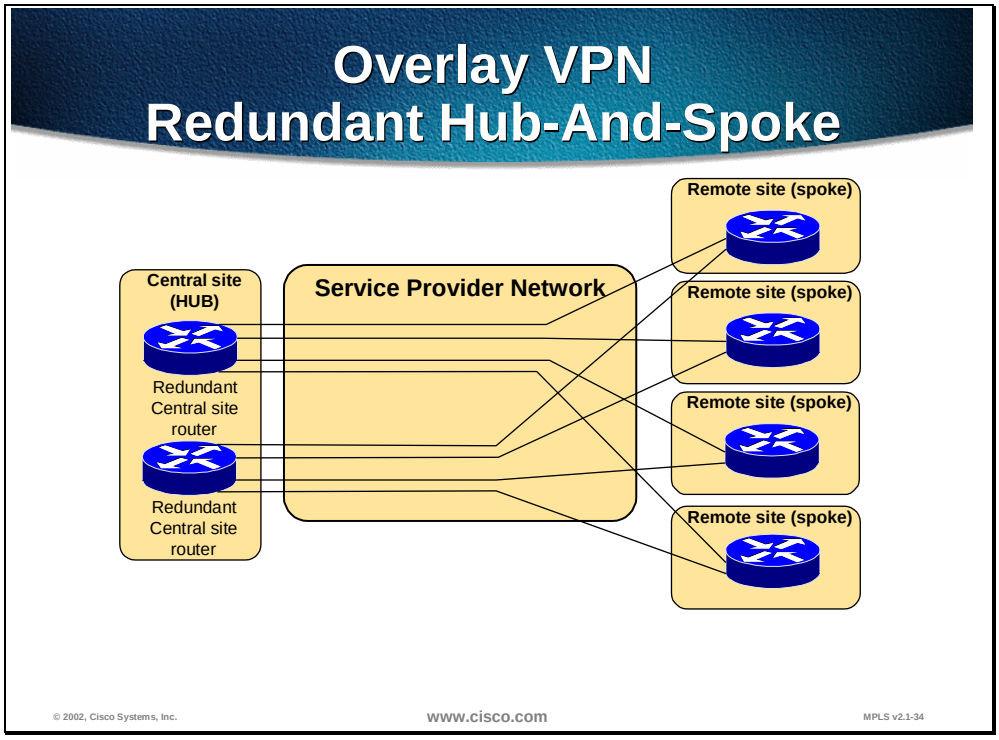

A typical redundant hub-and-spoke topology introduces central site redundancy (more complex topologies might also introduce router redundancy at spokes).

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

Practice

Q1) Which overlay VPN topology is the simplest?

A) Hub-and-spoke topology

B) Partial-mesh topology

C) Full-mesh topology

Partial-Mesh or Full-Mesh Overlay VPN Topology

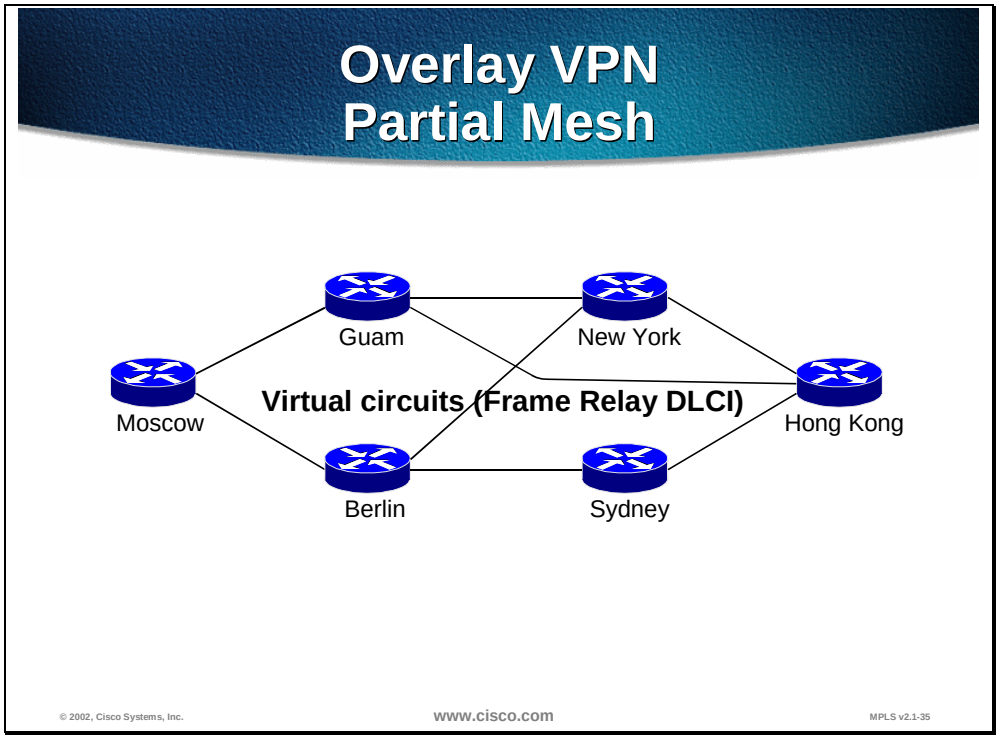

Partial mesh is used in environments where the cost or complexity factors prevent a full-mesh between customer sites. The virtual circuits in a partial mesh can be established based on a wide range of criteria:

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

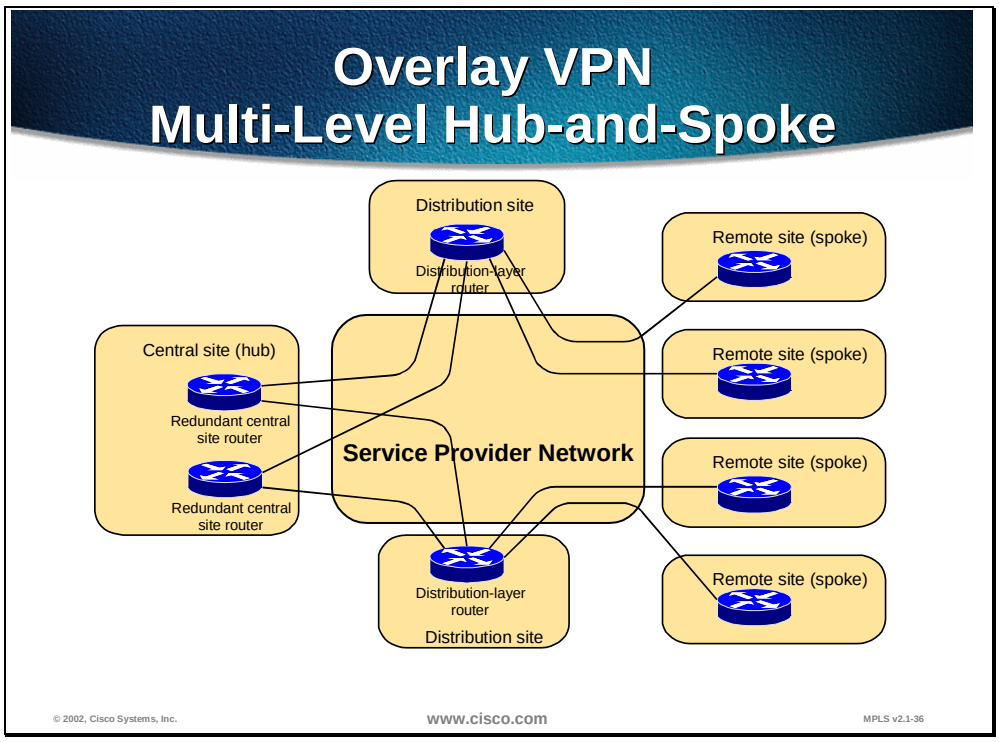

Various overlay VPN topologies are usually combined in a large network. For example, in the diagram above, a redundant hub-and-spoke topology is used in network core and a non-redundant hub-and-spoke is used between distribution sites and remote sites. This topology would be commonly used in environments where all traffic flows between the central site and remote sites and there is little (or no) traffic exchanged directly between the remote sites.

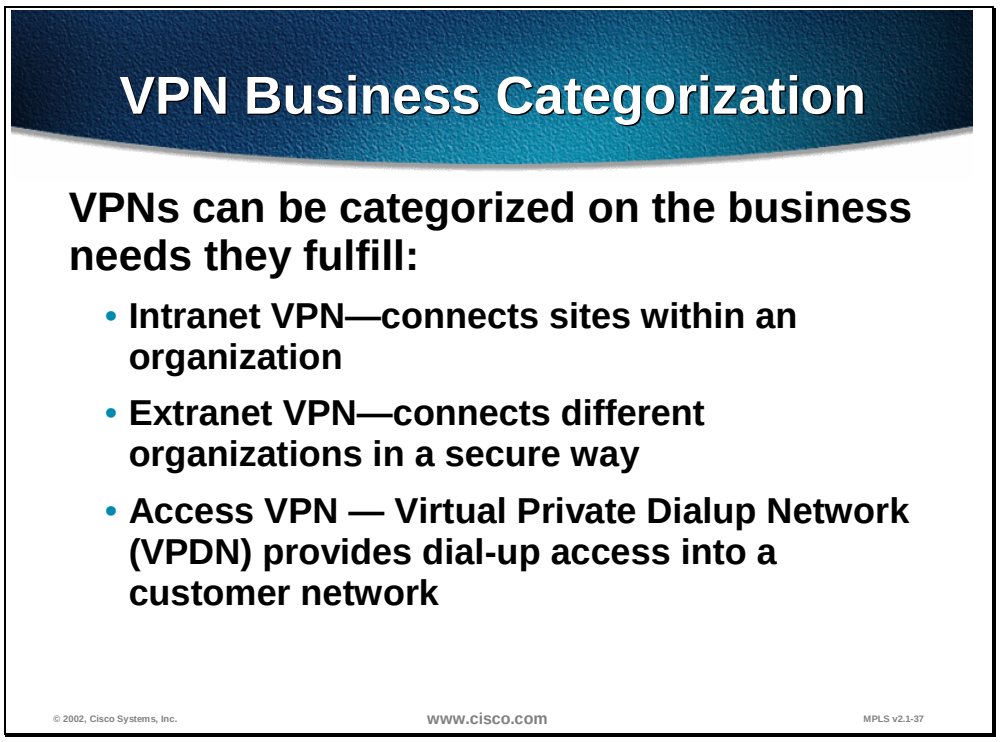

Another very popular VPN categorization classifies VPNs based on the business needs they fulfill:

- Intranet VPNs connect sites within an organization. Security mechanisms are usually not deployed in an Intranet, as all sites belong to the same organization.
- Extranet VPN connects different organizations. Extranets implementations usually rely on security mechanisms to ensure protection of individual organizations participating in the Extranet. The security mechanisms are usually the responsibility of individual participation organizations.
- Access VPN - Virtual Private Dialup Networks that provide dial-up access into a customer network.

Practice

Q1) What is the major driving force for customers to prefer partial mesh over full mesh topology?

A) Partial mesh is easier to configure.

B) Partial mesh is optimized for traffic patterns.

C) Connectivity costs usually dictate use of partial mesh.

D) Full mesh requires more networking equipment.

Simple Extranet Topology

The following two diagrams compare overlay VPN implementation of an Extranet with a peer-to-peer one. Similar comparisons could be made for Intranets as well.

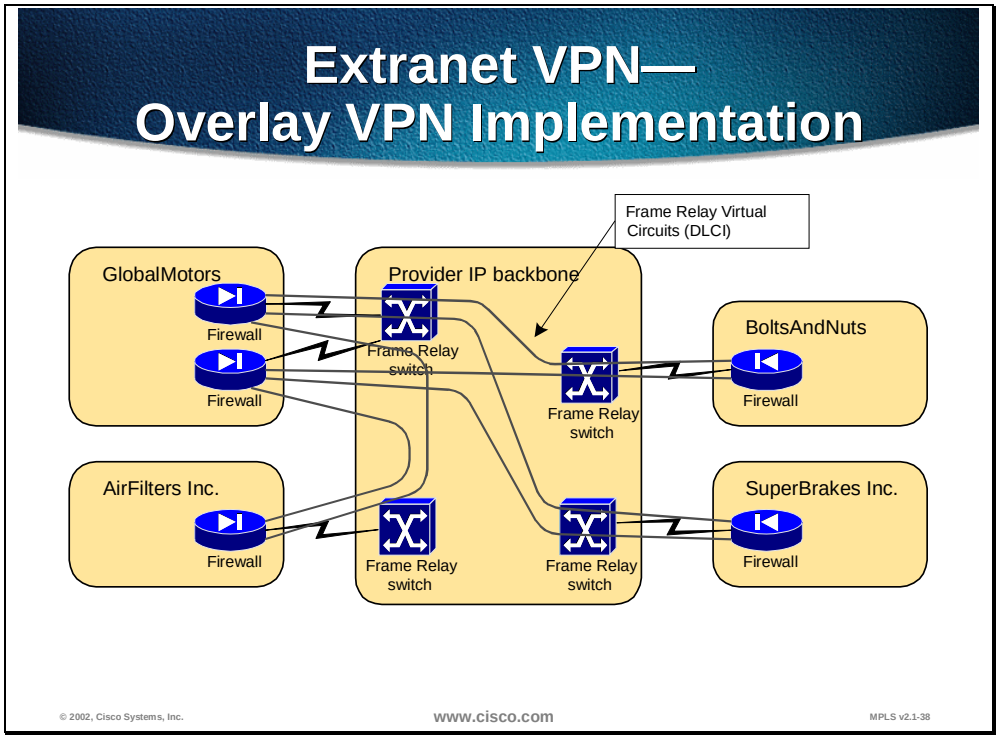

In an overlay implementation of an Extranet, organizations are linked with dedicated virtual circuits. Traffic between two organizations can only flow if:

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

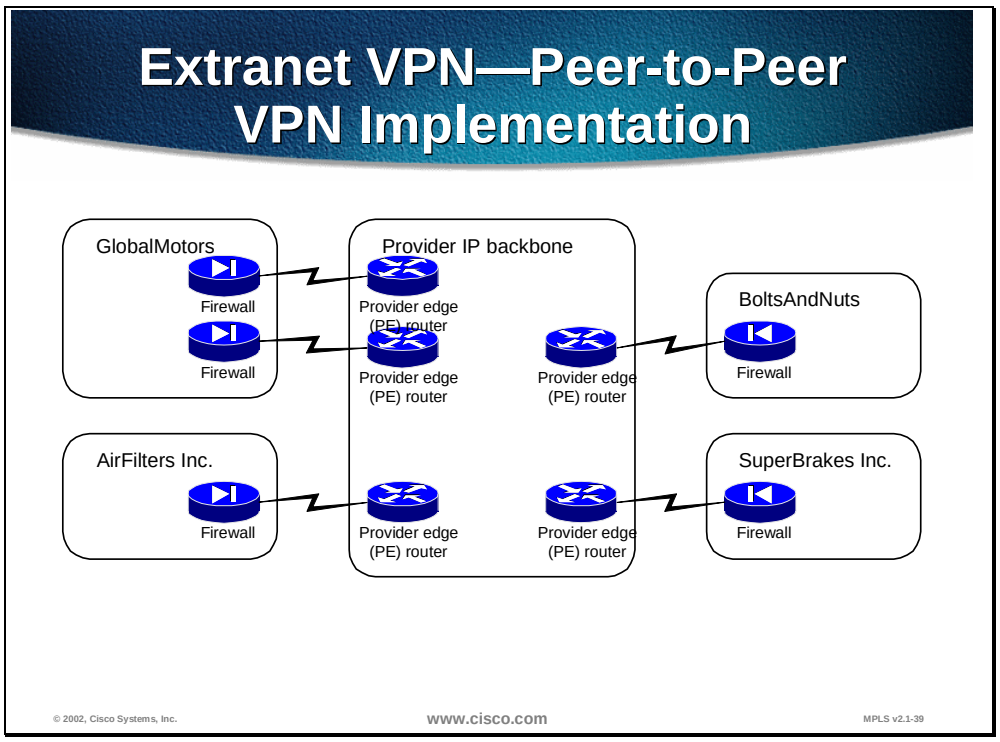

Peer-to-peer VPN implementation of an Extranet VPN is very simple compared to an overlay VPN implementation – all sites are connected to the Service Provider network and the optimum routing between sites is enabled by default.

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.

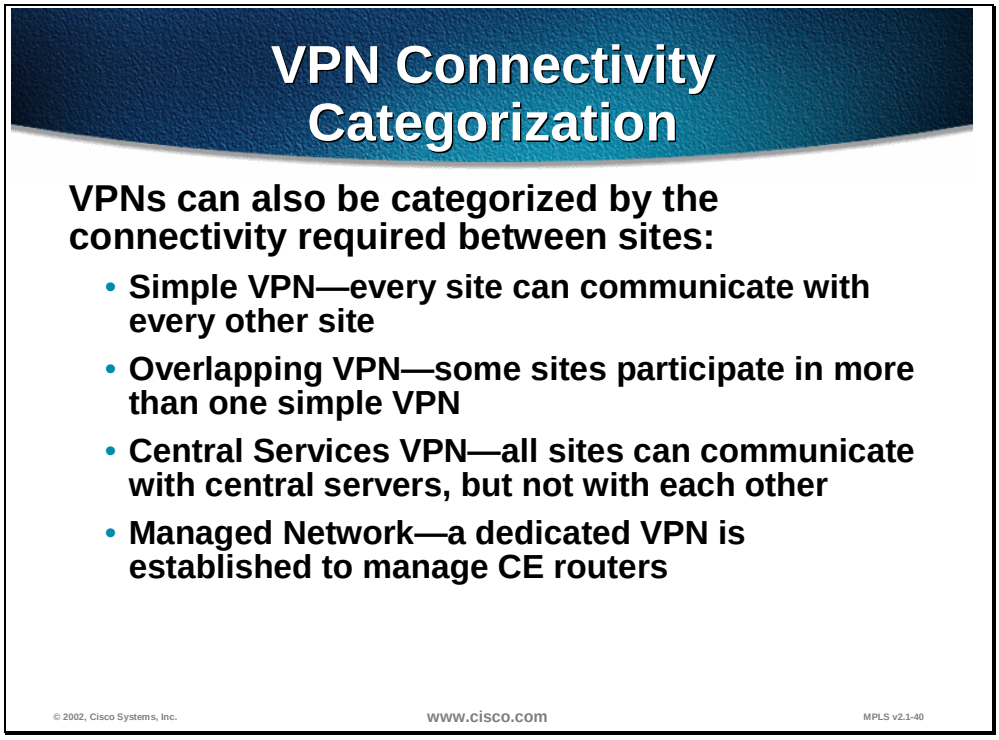

The virtual private networks discussed so far were usually very simple in connectivity terms:

- In most cases, full connectivity between sites was required (in overlay Intranet VPN implementations, this usually means that some customer sites act as transit sites)
- In the overlay implementation of the Extranet VPN, the connectivity was limited to sites that had direct virtual circuits established between them.

There are, however, a number of advanced VPN topologies with more complex connectivity requirements:

- **Overlapping VPNs**, where a site participates in more than one VPN
- **Central Services VPN**, where the sites are split in two classes – **server** sites that can communicate with all other sites and **client** sites that can only communicate with the servers, but not with other clients.
- **Network Management** VPN, which is used to manage CE devices in scenarios where the Service Provider owns and manages CE devices.

Practice

Q1) What is the difference between a simple VPN and a Central Services VPN?

A) Central services VPN is used to interconnect multiple simple VPNs.

B) Simple VPN allows any-to-any while Central services VPN allows client-to-server communication.

C) Central services VPN requires all packets to go through the central site.

D) Simple VPNs do not provide the connection to the Internet.

Central-Services Extranet

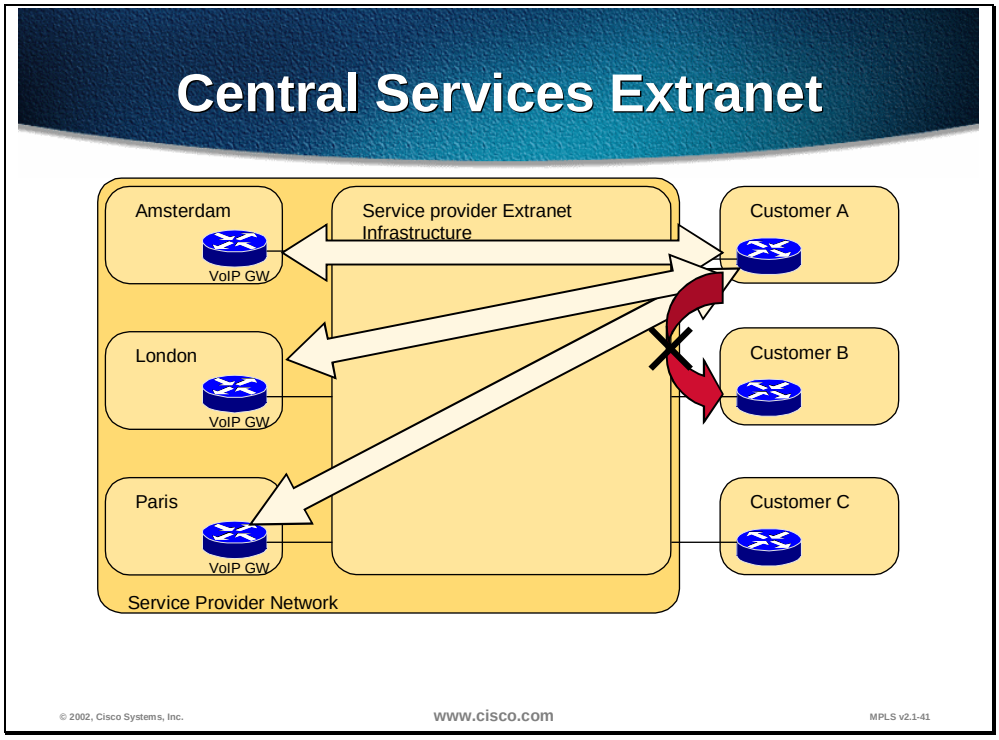

This diagram shows a sample Central Services extranet implementing international Voice-over-IP service. Every customer of this service can access voice gateways in various countries, but cannot access other customers using the same service.

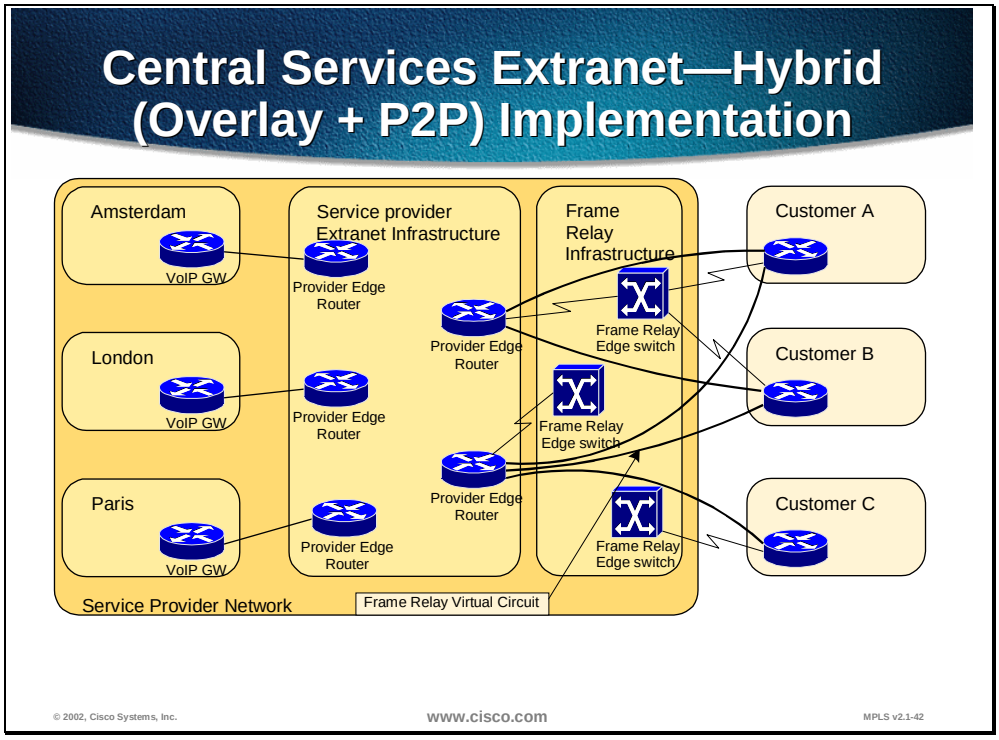

The network diagram shown above describes an interesting scenario where peer-to-peer VPN and overlay VPN implementation can be used to provide end-to-end service to the customer.

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

Managed Network VPN Topology

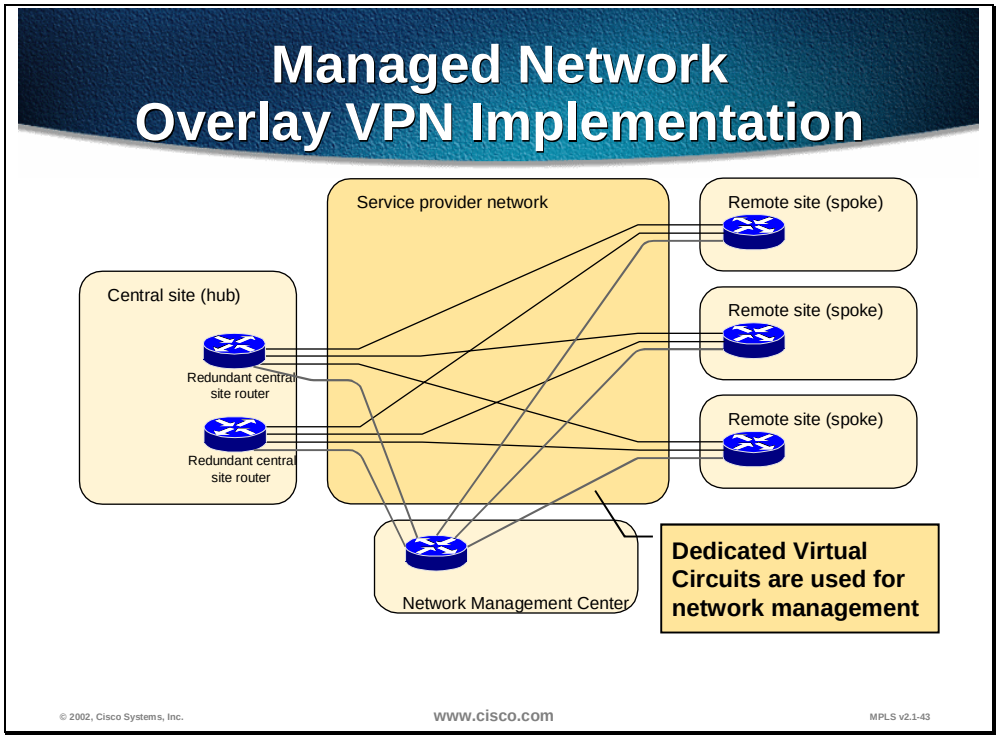

Network management VPN is traditionally implemented in combination with overlay VPN services. Dedicated virtual circuits are deployed between any managed CE-router and the central network management router (NMS-router) to which the Network Management Station (NMS) is connected.

This network management VPN implementation is sometimes called **rainbow** implementation, as 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

After completing this lesson, you should be able to perform the following tasks:

- Identify major VPN topologies
- Describe the implications of using overlay VPN or peer-to-peer VPN approach with each topology
- List sample usage scenarios for each topology

Next Steps

After completing this lesson, go to:

- MPLS VPN Architecture

Lesson Review

Instructions

Answer the following questions:

1. What are the major Overlay VPN topologies?
2. Why would the customers prefer partial mesh over full mesh topology?
3. What is the difference between an Intranet and an Extranet?
4. What is the difference between a simple VPN and a Central Services VPN?
5. What are the connectivity requirements of a Central Services VPN?

MPLS VPN Architecture

Overview

This lesson compares MPLS VPN with other peer-to-peer VPN implementations and describes the major benefits of MPLS VPN.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the difference between traditional peer-to-peer models and MPLS VPN
- List the benefits of MPLS VPN
- Describe major architectural blocks of MPLS VPN
- Explain the need for route distinguisher and route target

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- Peer-to-Peer VPN Review
- Overall MPLS VPN Architecture
- Virtual Routing Tables and Route Distinguishers
- Support for Complex VPN Topologies—Route Targets
- Impact of Complex VPN Topologies on Route Distinguishers
- Benefits of MPLS VPN Versus Other Peer-to-Peer VPN Technologies
- Summary

Peer-to-Peer VPN Review

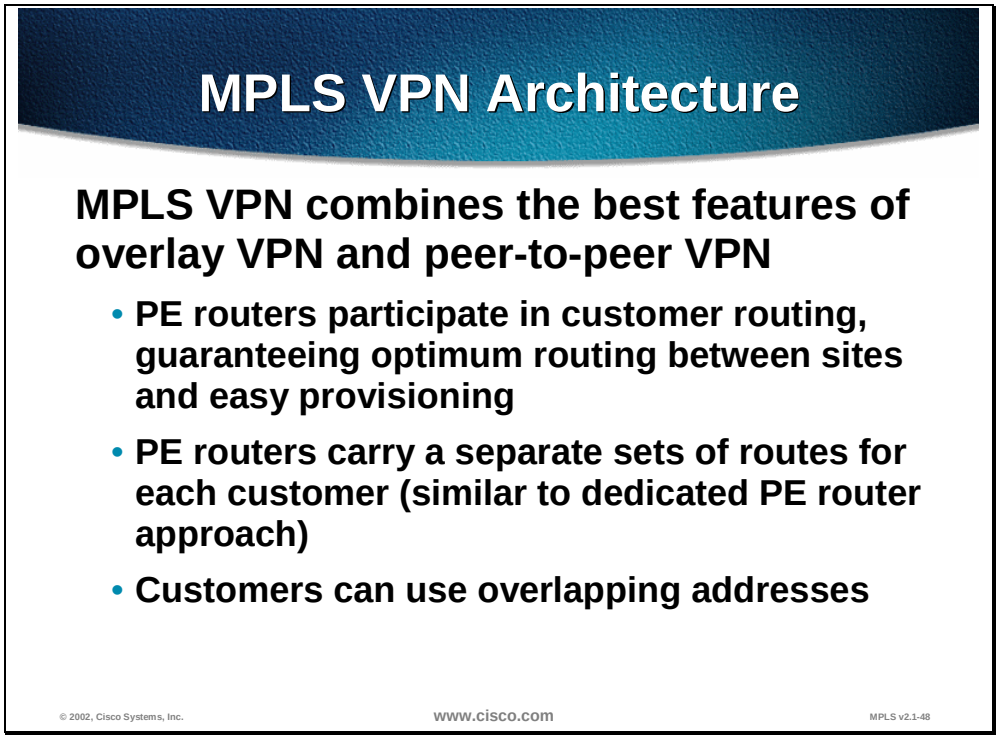

The MPLS VPN architecture provides the Service Providers with a peer-to-peer VPN architecture that combines the best features of overlay VPN (support for overlapping customer address spaces) with the best features of peer-to-peer VPNs:

- PE routers participate in customer routing, guaranteeing optimum routing between customer sites
- PE routers carry separate set of routes for each customer, resulting in perfect isolation between the customers.

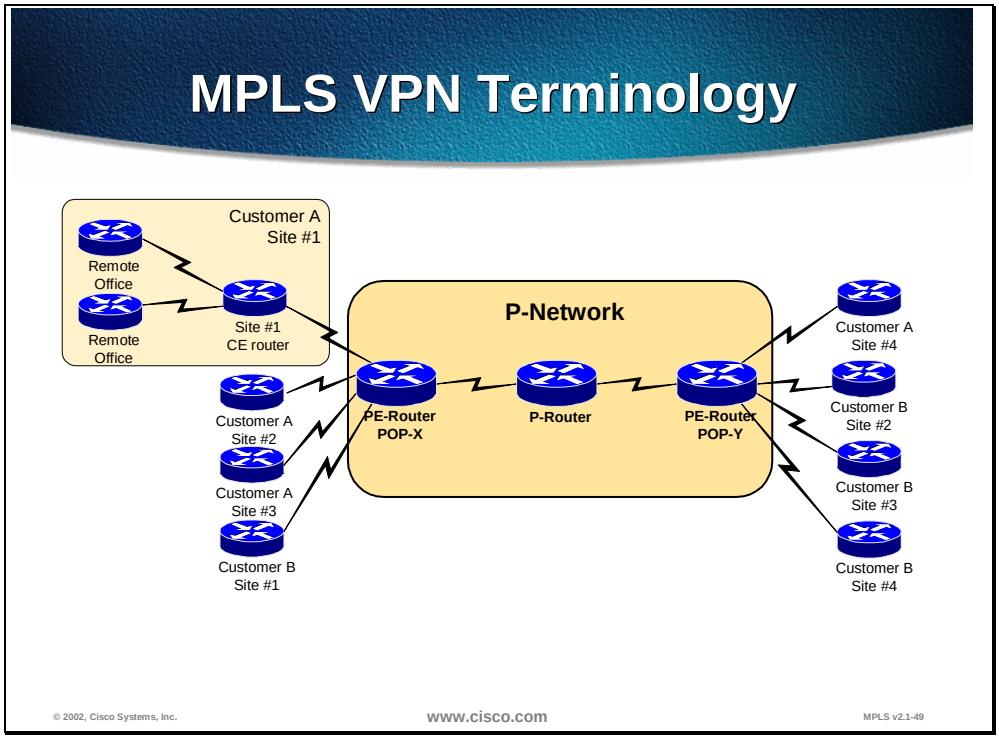

The MPLS VPN terminology divides the overall network into customer controlled part (**C-network**) and provider controlled part (**P-network**). Contiguous portions of C-network are called **sites** and are linked with the P-network via **CE-routers**. The CE-routers are connected to the **PE-routers**,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

Overall MPLS VPN Architecture

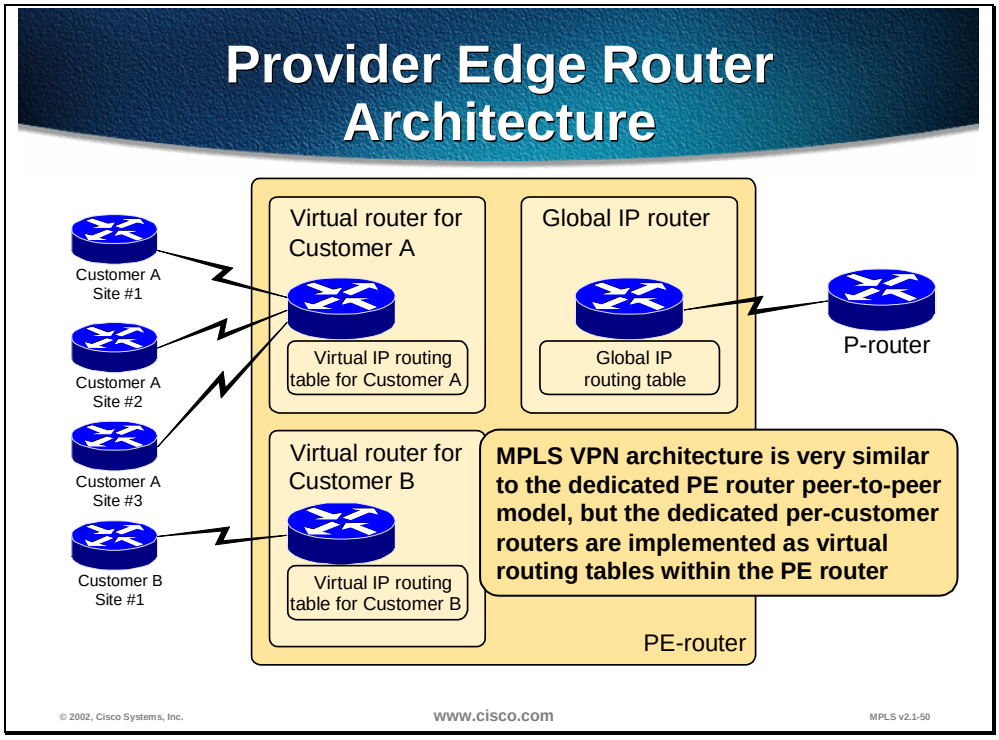

The architecture of a PE-router in MPLS VPN is very similar to the architecture of a Point-of-Presence (POP) in the dedicated PE-router peer-to-peer model, the only difference being that the whole architecture is condensed into one physical device. Each customer is assigned an independent routing table (virtual routing table) that corresponds to the dedicated PE-router in traditional peer-to-peer model. Routing across the provider backbone is performed by another routing process that uses global IP routing table, corresponding to the intra-POP P-router in traditional peer-to-peer model.

Note IOS implements isolation between customers via **virtual routing and forwarding tables (VRFs)**. The whole PE-router is still configured and managed as a single device, not as a set of virtual routers.

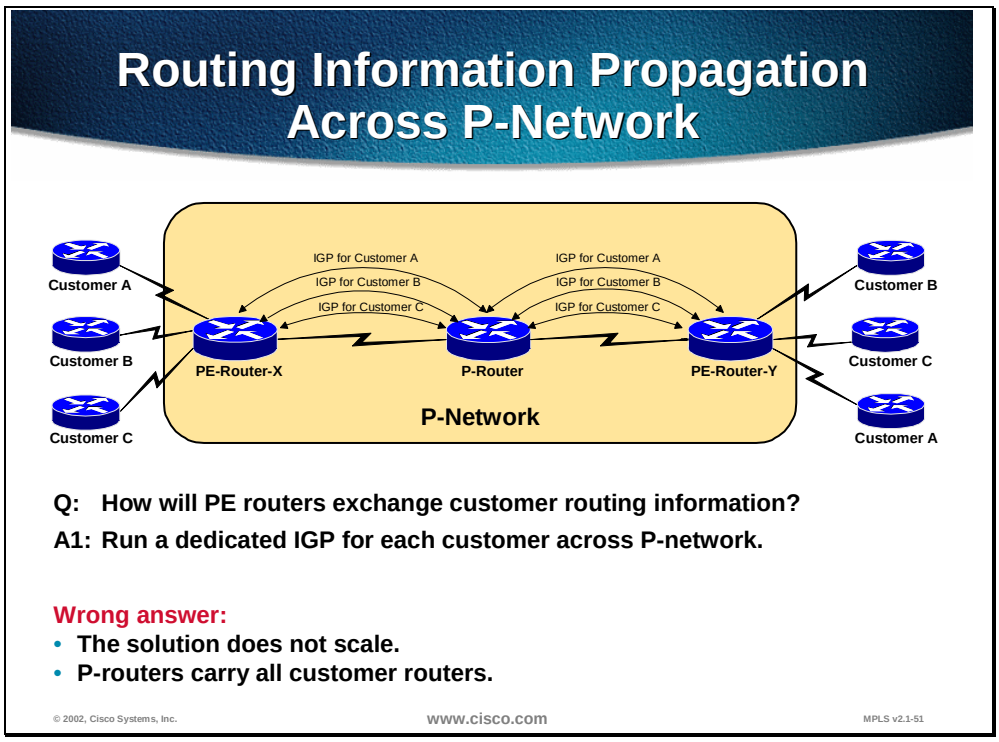

While the virtual routing tables provide the isolation between customers, the data from these routing tables still needs to be exchanged between PE-routers to enable data transfer between sites attached to different PE-routers. We therefore need a routing protocol that will transport all customer routes across the Provider network while maintaining the independency of individual customer address spaces.

An obvious solution, implemented by various VPN vendors, is to run a separate routing protocol for each customer. The PE-routers could be connected via point-to-point tunnels (and the per-customer routing protocols would run between PE-routers) or the P-routers could participate in the customer routing.

This solution, although very simple to implement (and even used by some customers), is not appropriate in Service Provider environments, as it simply does not scale:

- The PE-routers have to run a large number of routing protocols
- The P-routers have to carry all customer routes.

Routing Information Propagation Across P-Network (Cont.)

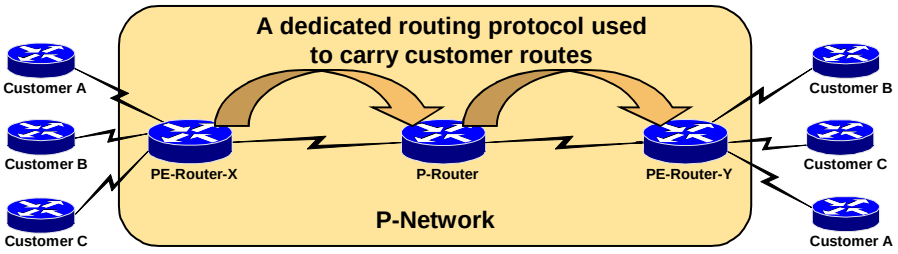

Q: How will PE routers exchange customer routing information?

A2: 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 routers.

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

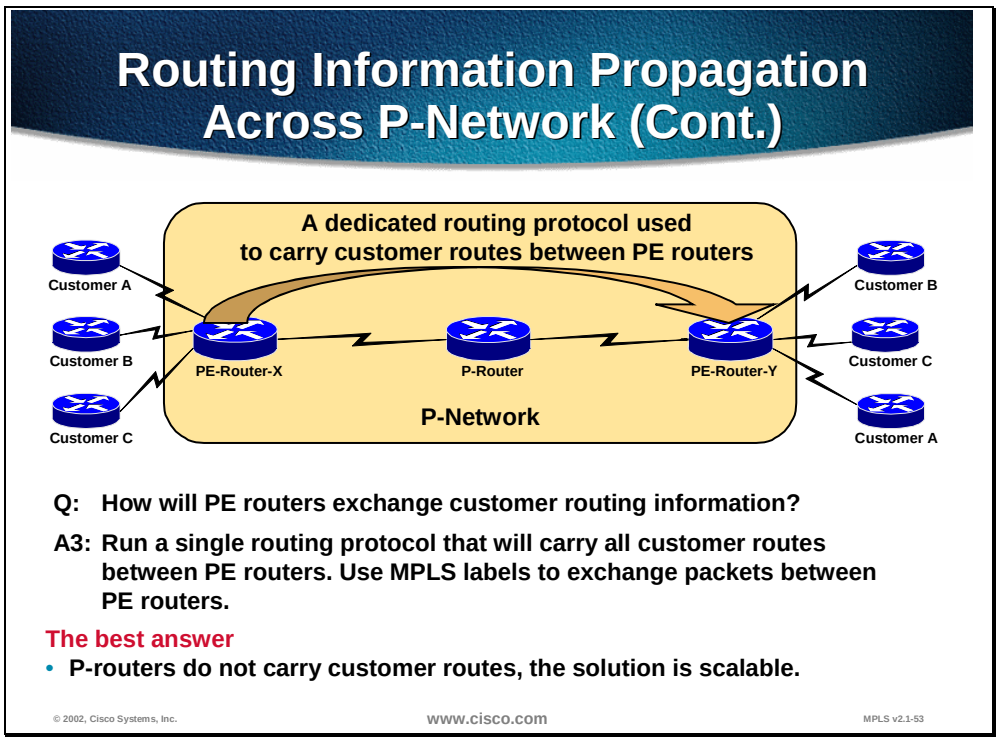

The best solution to customer route propagation is hence to run a single routing protocol between PE-routers that will exchange all customer routes without the involvement of the P-routers. This solution is scalable:

- The number of routing protocols running between PE-routers does not increase with increasing number of customers
- The P-routers do not carry customer routes.

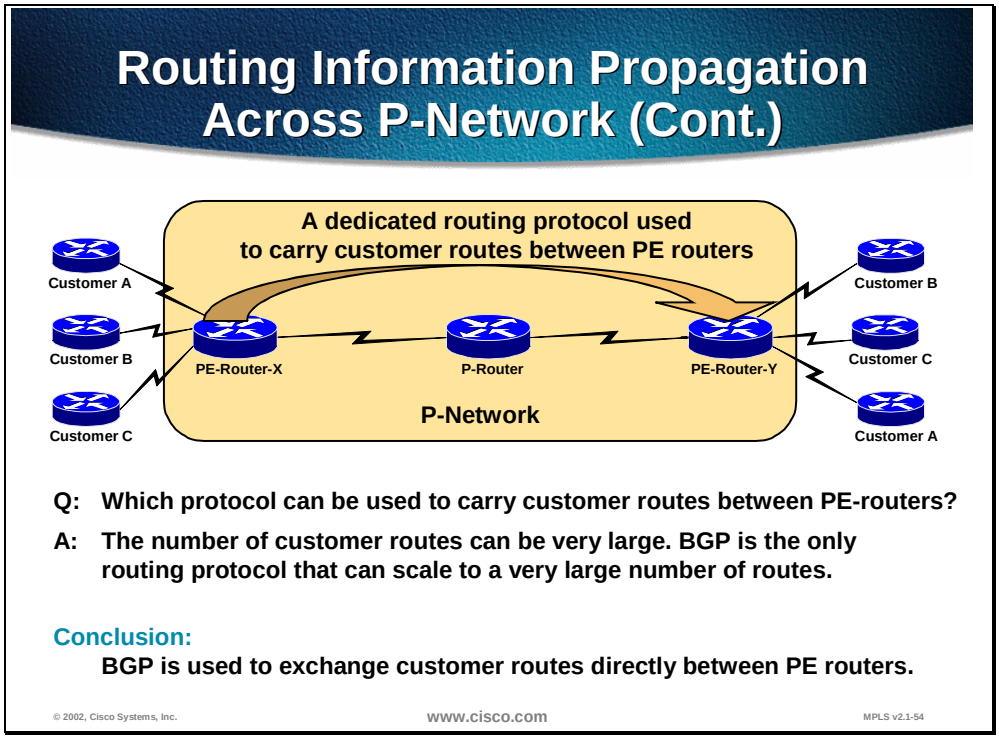

The next design decision to be made is the choice of the routing protocol running between PE-routers. As the total number of customer routes is expected to be very large, the only well known protocol with the required scalability is Border Gateway Protocol (BGP).

Conclusion: BGP is used in MPLS VPN architecture to transport customer routes directly between PE-routers

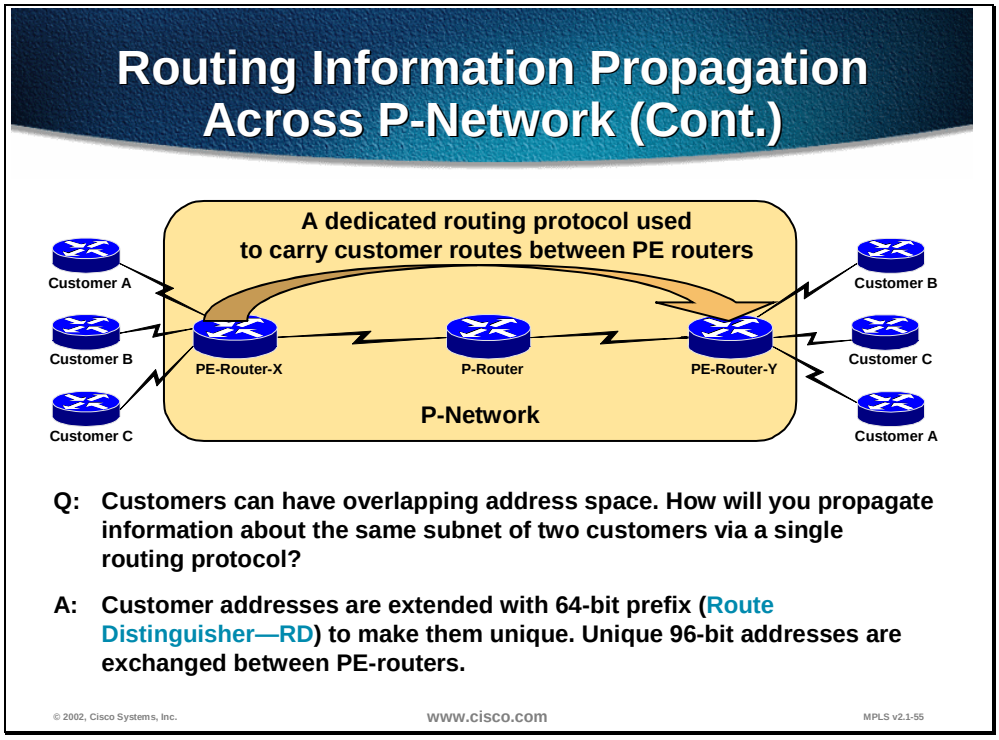

MPLS VPN architecture provides an important differentiator against traditional peer-to-peer VPN solutions – the support of overlapping customer address spaces.

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 will make them unique even if they were previously overlapping. A 64-bit prefix, called **route distinguisher (RD)**, is used in MPLS VPN to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE-routers.

BGP is the dedicated routing protocol used between the PE routers. But BGP had to be enhanced before it could be used to carry the 96-bit routes. The enhanced version of BGP is called Multi-Protocol Border Gateway Protocol (MP-BGP). The 96-bit routes are treated by MP-BGP as addresses of a new address family, the VPNv4 address family.

The PE routers will use traditional IP version 4 (IPv4) addresses when it is exchanging routes with the CE routers. But then they are internally converted into VPNv4 routes by prefixing them with the RD. The VPNv4 routes are the propagated to the other PE router which can remove the RD before propagating it to the CE router.

Practice

Q1) How are customer routes exchanged across the P-network?

A) Running a single BGP that can exchange all customer routes across the Provider network.

B) Running a separate IBGP for each customer.

C) Multiple IGPs (one per customer) are used to exchange customer routes across the P-network.

D) Multi-protocol BGP (MP-BGP) is used to exchange customer routes across the P-network.

Virtual Routing Tables and Route Distinguishers

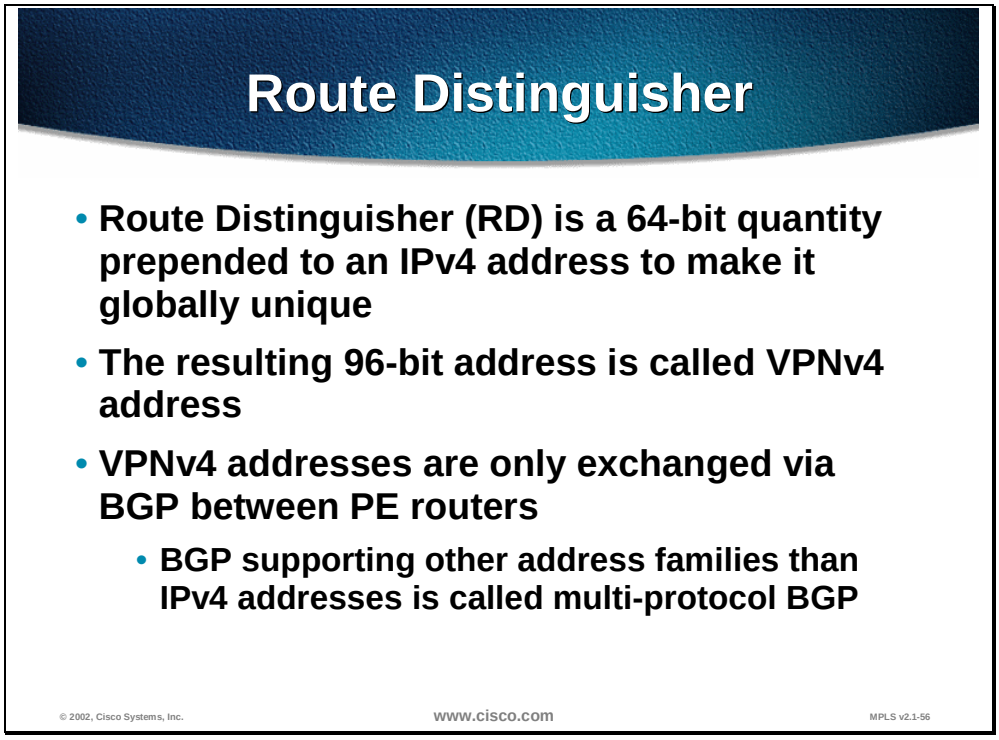

Route Distinguisher (RD) is a 64-bit prefix that is only used to transform non-unique 32-bit customer IPv4 addresses into unique 96-bit VPNv4 addresses (also called VPN_IPv4 addresses).

The RD is an arbitrary number. It is configurable and should be selected according to a user defined allocation scheme. But from a technical perspective, the RD number has no meaning and is used only to distinguish two identical IPv4 routes received from two different customers as two different VPNv4 routes.

The VPNv4 addresses are only exchanged between PE-routers; they are never used between CE-routers and CE-routers. BGP between PE-routers must therefore support exchange of traditional IPv4 prefixes as well as exchange of VPNv4 prefixes. The BGP session between PE-routers is consequently called **multi-protocol** BGP (MP-BGP) session.

Note	A MP-BGP session between two PE routers inside a single autonomous system is a MP-IBGP session. This session obeys the same rules as all other IBGP sessions with respect to the BGP attributes next-hop and AS-path etc.

Note	In the case where a VPN spans more than one AS, then the VPNv4 routes must be exchanged across the AS boundary using a MP-EBGP session.

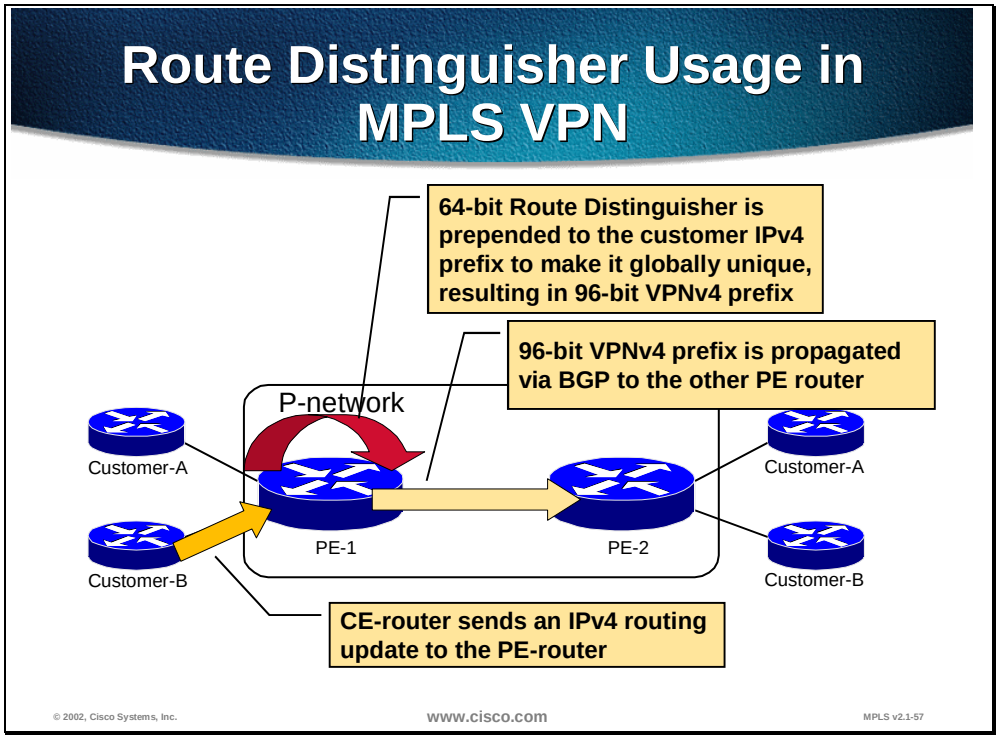

The customer route propagation across MPLS VPN network is performed in the following steps:

Step 1 CE-router sends an IPv4 routing update to the PE-router

Step 2 PE-router prepends 64-bit route distinguisher to the IPv4 routing update, resulting in globally unique 96-bit VPNv4 prefix

Step 3 The VPNv4 prefix is propagated via Multi-Protocol Internal BGP (MP-IBGP) session to other PE-routers

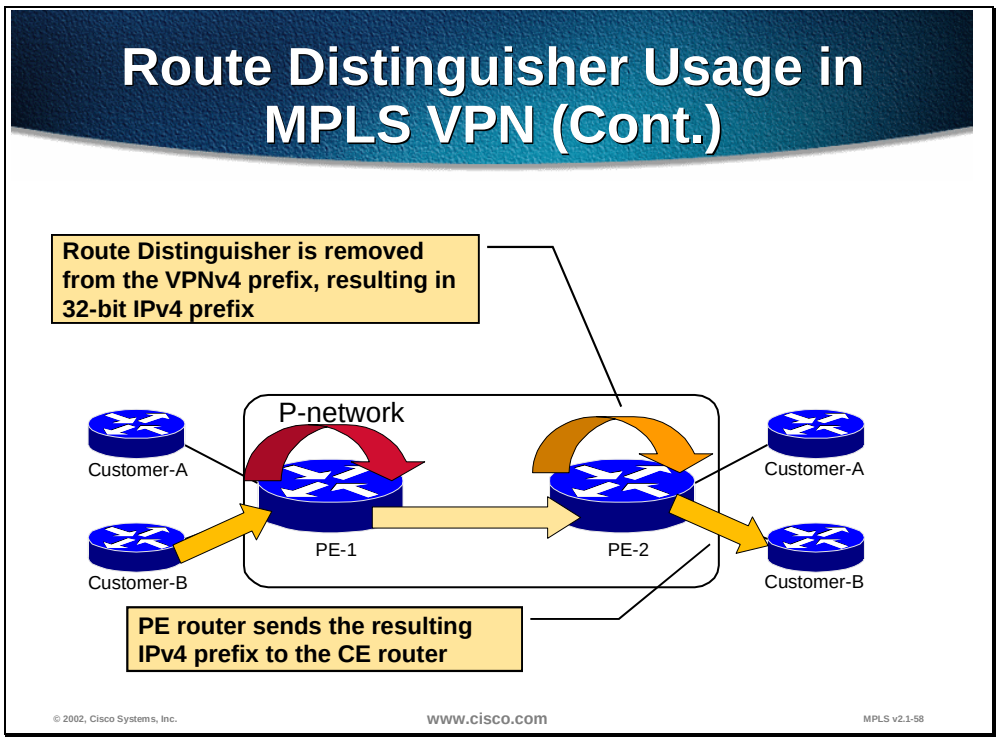

Step 4 The receiving PE-routers strip the route distinguisher from the VPNv4 prefix, resulting in IPv4 prefix

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

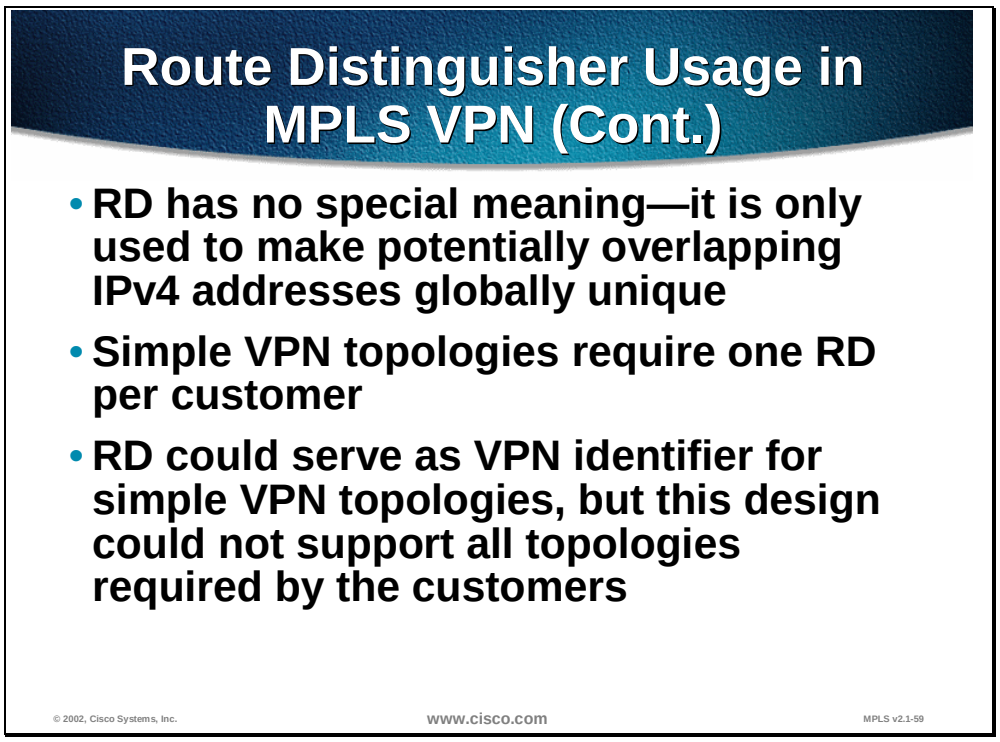

The route distinguisher has no special meaning or role in MPLS VPN architecture – its only function is to make overlapping IPv4 addresses globally unique.

Note As there has to be a unique one-to-one mapping between the route distinguishers and virtual routing and forwarding tables, the route distinguisher could be viewed as the VRF identifier in Cisco's implementation of MPLS VPN.

The route distinguisher is configured at the PE router as part of the setup of a VPN site. It is not configured on the customer equipment, and is not visible to the customer.

Simple VPN topologies only require one route distinguisher per customer, raising the possibility that RD could serve as VPN identifier. This design, however, would not allow implementation of more complex VPN topologies, like when a customer site belongs to multiple VPNs.

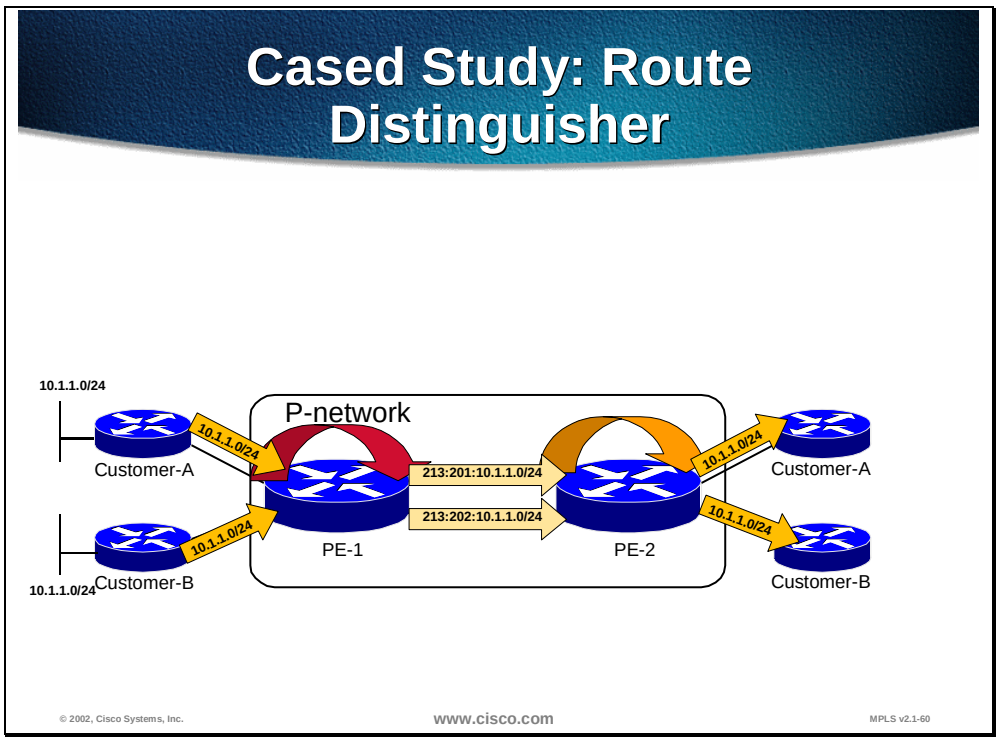

The figure illustrates a network where two different customers, Customer-A and Customer-B, are connected to an MPLS/VPN service provider network. Customer-A has two sites connected and Customer-B also has two sites connected.

Both customers are using class A network 10.0.0.0 private addresses. This results in an overlapping address space.

The MPLS/VPN architecture handles this overlapping by using route distinguishers. When PE-1 receives routing information from customer-A regarding subnet 10.1.1.0/24, PE-1 internally prefixes the RD 213:201 associated with the customer access link to the IPv4 subnet. Thus a VPNv4 route 213:201:10.1.1.0/24 is created.

PE-1 will use the RD 213:202 for all routes received on the access link to the Customer-B site. This results in two identical IPv4 routes being converted into two different and distinguishable VPNv4 routes. The VPNv4 routes are propagated using MP-BGP to PE-2.

When PE-2 receives the two different VPNv4 routes, it will use a MP-BGP attribute not explained here to determine that one of them is intended for Customer-A and the other is intended for Customer-B. Before the route to Customer-A is propagated to the Customer-A CE router, the RD is removed, thus recreating the original IPv4 route. PE-2 does the same with the Customer-B route. Now the two routes are identical again, but that is no problem because they are propagated on two different access links to two different CE routers.

Practice

Q1) What is a route distinguisher?

A) It is a new BGP attribute used to determine to which VPN a route belongs.

B) It is a process that identifies the VRF to which the address belongs.

C) It is a 64-bit value assigned to VRFs to distinguish between instances of routing protocols.

D) It is a 64-bit prefix prepended to customer IPv4 address to make it globally unique.

Support for Complex VPN Topologies—Route Targets

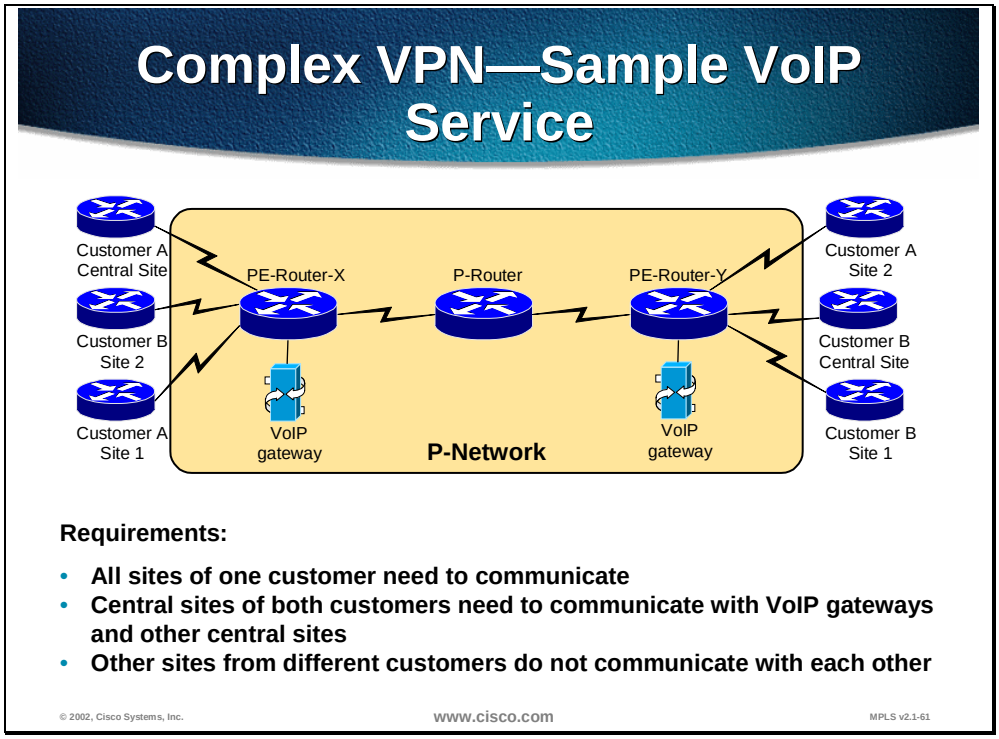

To illustrate the need for more versatile VPN indicator than the route distinguisher, consider the Voice-over-IP service illustrated in the figure above. The connectivity requirements of this service are as follows:

- All sites of a single customer need to communicate
- Central sites of different customers subscribed to VoIP service need to communicate with the VoIP gateways (to originate and receive calls toward public voice network) as well as with other central sites to exchange inter-company voice calls.

Note Additional security measures would have to be put in place at central sites to make sure that the central sites only exchange VoIP calls with other central sites, otherwise the corporate network of a customer could be compromised by another customer using VoIP service.

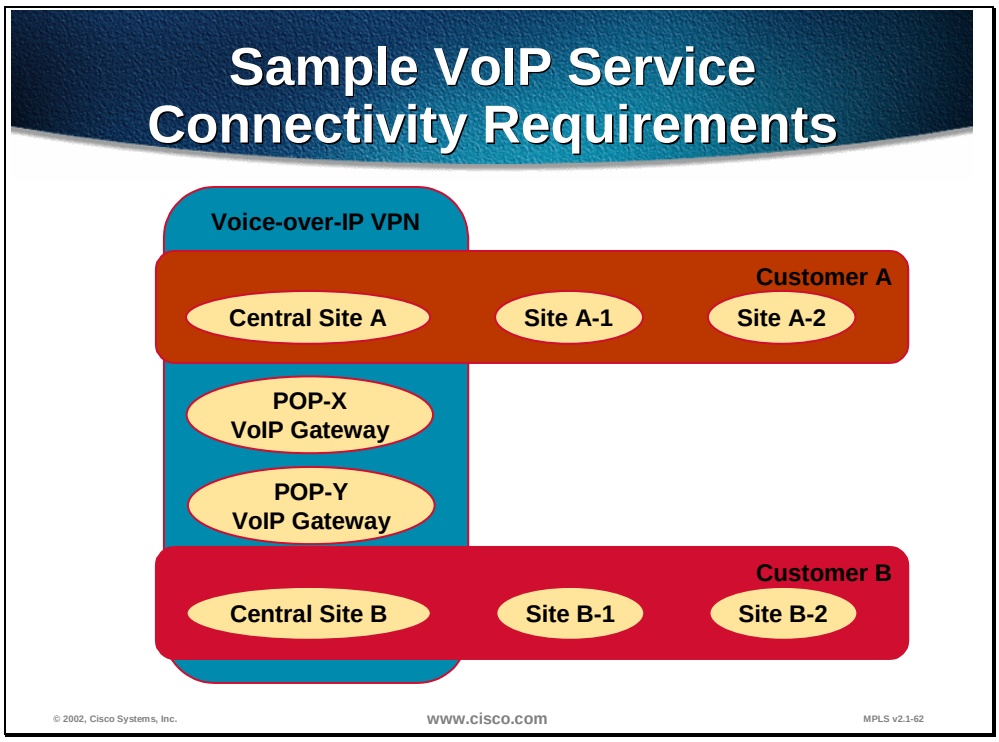

The connectivity requirements of the VoIP service are illustrated in the diagram above. There are three VPNs needed to implement the desired connectivity – two customer VPNs and a shared Voice-over-IP VPN. Central customer sites participate in the customer VPN as well as in the Voice-over-IP VPN.

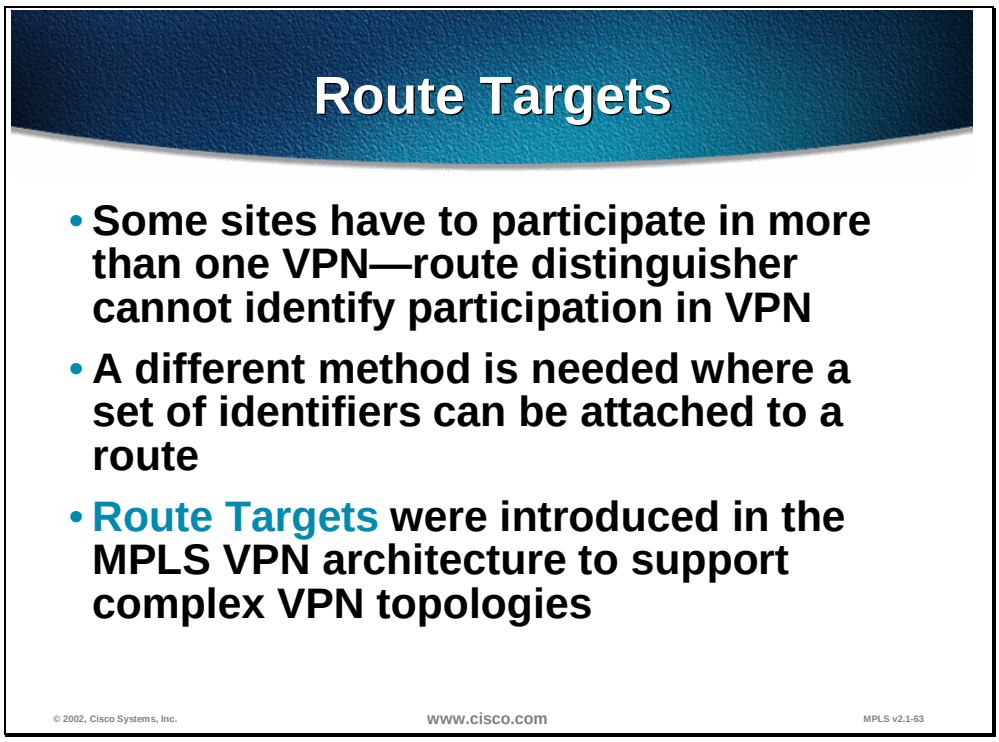

The route distinguisher (which is a single entity prepended to an IPv4 route) cannot indicate that a site participates in more than one VPN. A different method is needed where a set of VPN identifiers could be attached to a route to indicate its membership in several VPNs.

The **route targets** were introduced in the MPLS VPN architecture to support this requirement.

The generic nature of the MPLS/VPN architecture requires the Route Target to be used to indicate VPN membership in all topologies. Also simple topologies require the use of Route Targets.

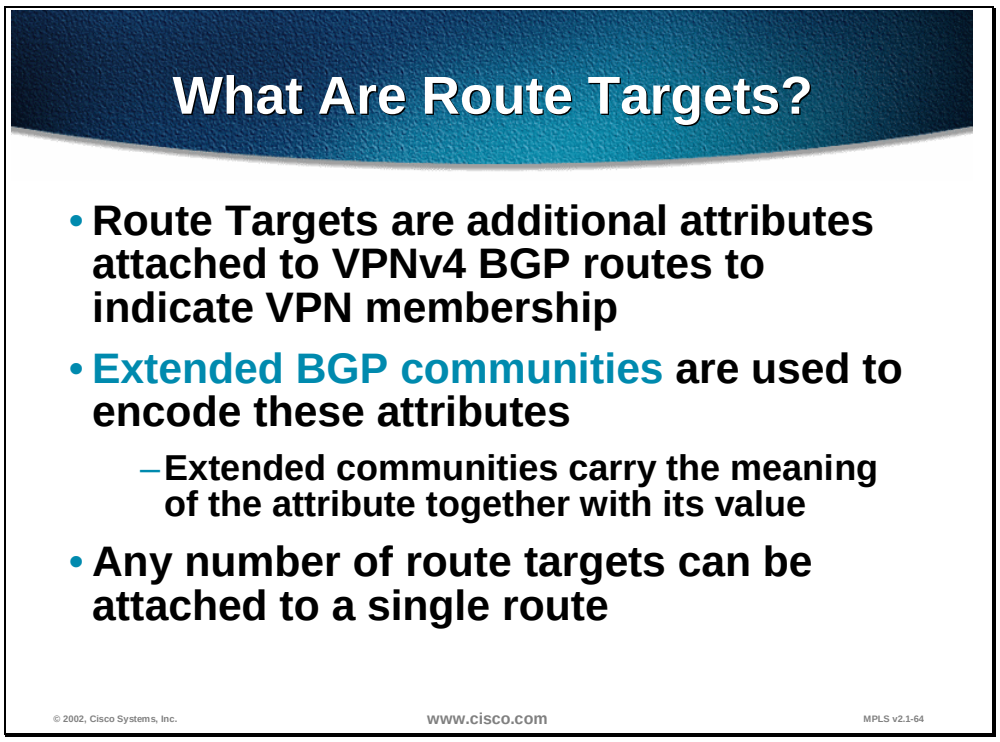

Route targets (RT) are extended BGP communities that are attached to a VPNv4 BGP route to indicate its VPN membership. As with standard BGP communities, a set of extended communities can be attached to a single BGP route, satisfying the requirements of complex VPN topologies. A route can carry one RT which indicates membership in one VPN at the same time as it carries an additional RT which indicates membership in a completely different VPN.

Extended BGP communities are 64-bit values. The semantics of the extended BGP community is encoded in the high-order 16 bits of the value, making them useful for a number of different applications. For example, the value of high-order 16 bits of extended BGP community is two (2) for MPLS VPN Route Targets.

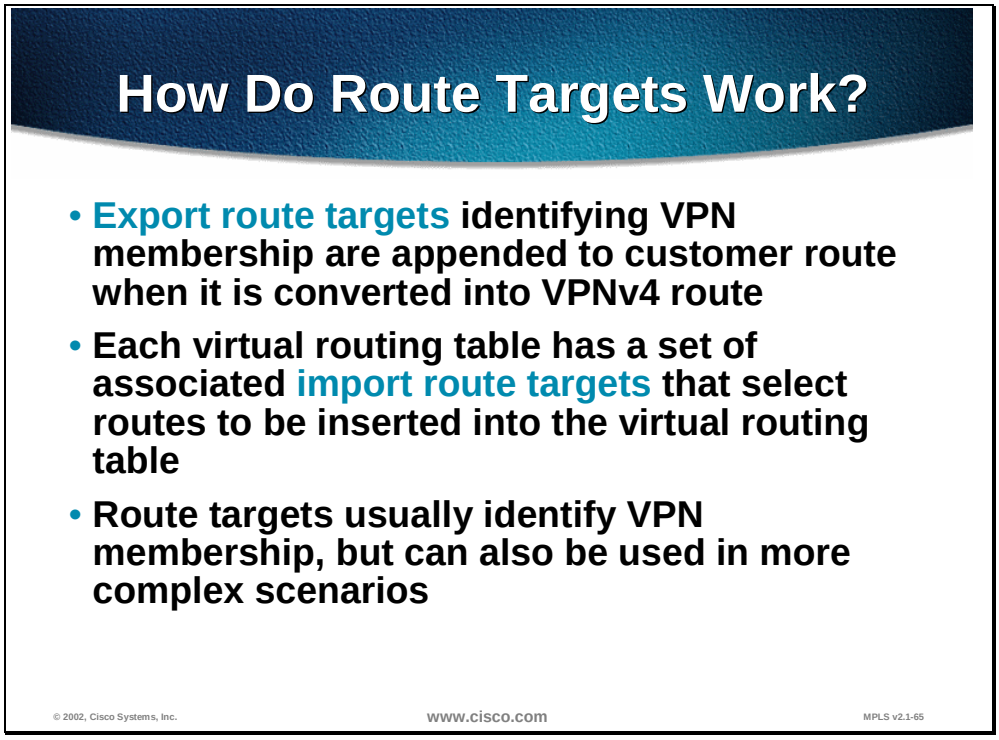

MPLS VPN route targets are attached to a customer route at the moment when it's converted from IPv4 route to a VPNv4 route by the PE-router. The route targets attached to the route are called **export route target** and are configured separately for each virtual routing table in a PE-router. The export route targets 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 done based on **import route targets**. Each virtual routing table in a PE-router can have a number of import route targets configured, identifying the set of VPNs from which this virtual routing table is accepting routes.

Note Please refer to **MPLS VPN Implementation on Cisco IOS** chapter for more details on import and export route targets.

In overlapping VPN topologies, the route targets are used to identify VPN membership. Advanced VPN topologies (for example, central services VPN) use route targets in more complex scenarios – please refer to **MPLS VPN Topologies** chapter of **MPLS VPN Solutions** lesson for more details.

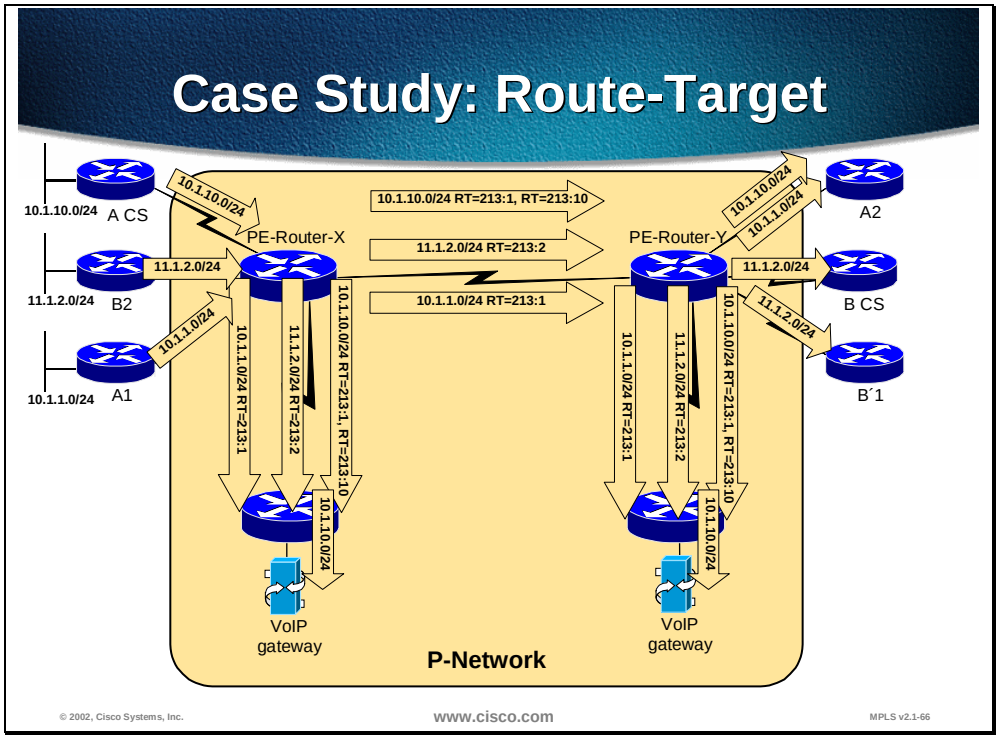

The figure illustrates how two customers, A and B, have several sites connected to an MPLS/VPN service provider network. Customer A has three sites: the central site A CS and two other sites, A1 and A2. Customer B also has the corresponding three sites: B CS, B1 and B2. The service provider is also providing a value added service of Voice Over IP, (VoIP).

The connectivity requirements say that all the customer A sites should communicate. At the same time all customer B sites should communicate. But customer A sites should not be able to communicate with customer B sites. In addition to these basic VPN requirements the site A CS should be able to communicate with the VoIP gateways.

This is an example of overlapping VPNs.

Route Targets are assigned to satisfy the requirements. RT 213:1 is used for the customer A VPN, RT 213:2 is used for customer B and RT 213:10 is used for the VoIP gateways VPN.

All routes received from any customer A site is assigned the RT 213:1. The PE routers will use the RT value of incoming VPNv4 routes to determine if the routes should be imported and propagated to other customer A sites.

All routes received from any customer B site is assigned the RT 213:2. The PE routers use this value to select routes to be imported to customer B sites.

The routes that are received by PE-Router-X from A CS is in addition to the RT 213:1 also assigned the RT 213:10. That means that the route 10.1.10.0/24 from A CS is converted into a VPNv4 route which is assigned two different RTs. When the route is received by PE-Router-Y, it finds the RT 213:1, which is enough for it to be imported and propagated to site A2. When the route is received by the two PE routers for VoIP, they find the RT 213:10, which is enough for it to be propagated to the VoIP sites.

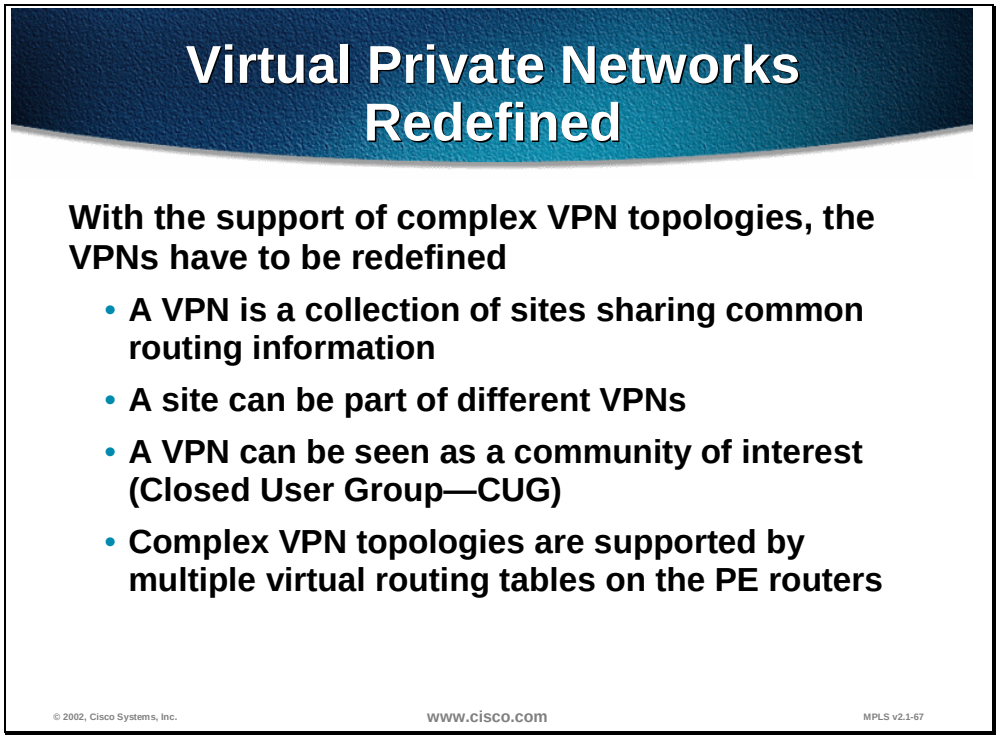

With the introduction of complex VPN topologies, the definition of a Virtual Private Network needs to be changed – a VPN is simply **a collection of sites sharing common routing information**. In traditional switched WAN terms (for example, in X.25 terminology), such a concept would be called **closed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

A site can be part of different VPNs, resulting in differing routing requirements for sites that belong to different sets of VPNs. These routing requirements have to be supported with multiple virtual routing tables on the PE-routers.

Practice

Q1) Why were the route targets introduced in MPLS VPN architecture?

A) The route target is equivalent to the route distinguisher. The two terms are used interchangable.

B) To support complex VPN topologies.

C) To identify the VPN by assigning the route target value as the MPLS label.

Impact of Complex VPN Topologies on Route Distinguishers

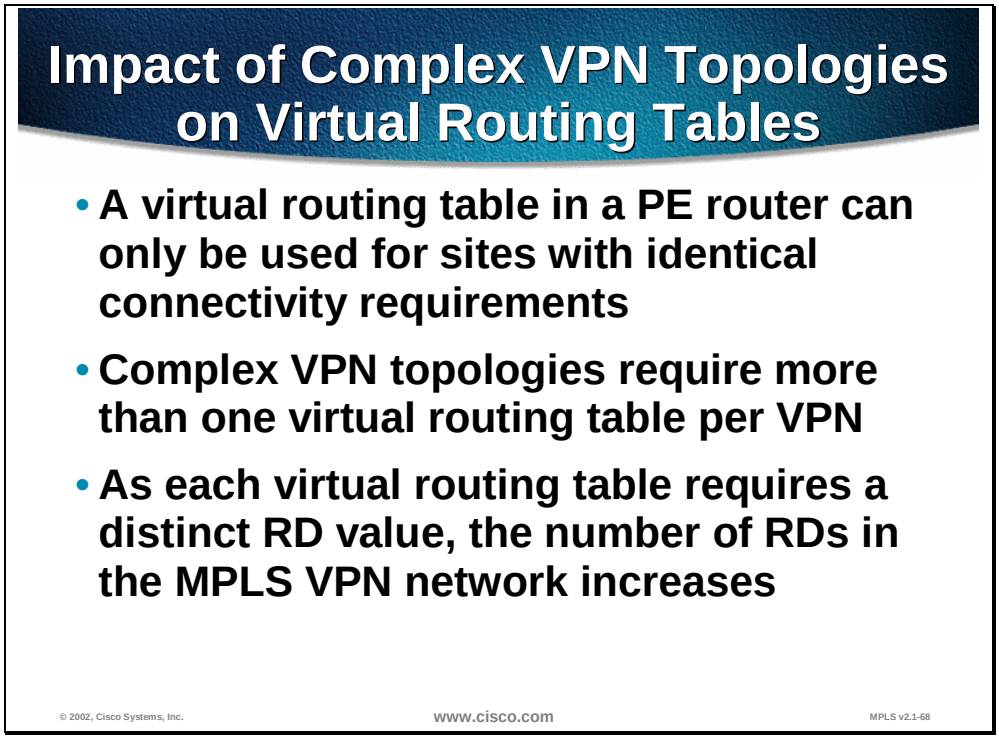

A single virtual routing table can only be used for sites with identical connectivity requirements. Complex VPN topologies therefore require more than one virtual routing table per VPN.

Note If you would associate sites with different requirements with the same virtual routing table, some of them might be able to access destinations that should not be accessible to them otherwise.

As each virtual routing table requires a distinctive route distinguisher value, the number of route distinguisher in MPLS VPN network increases with the introduction of overlapping VPNs. Moreover, the simple association between route distinguisher and VPN that was true for simple VPNs is also gone.

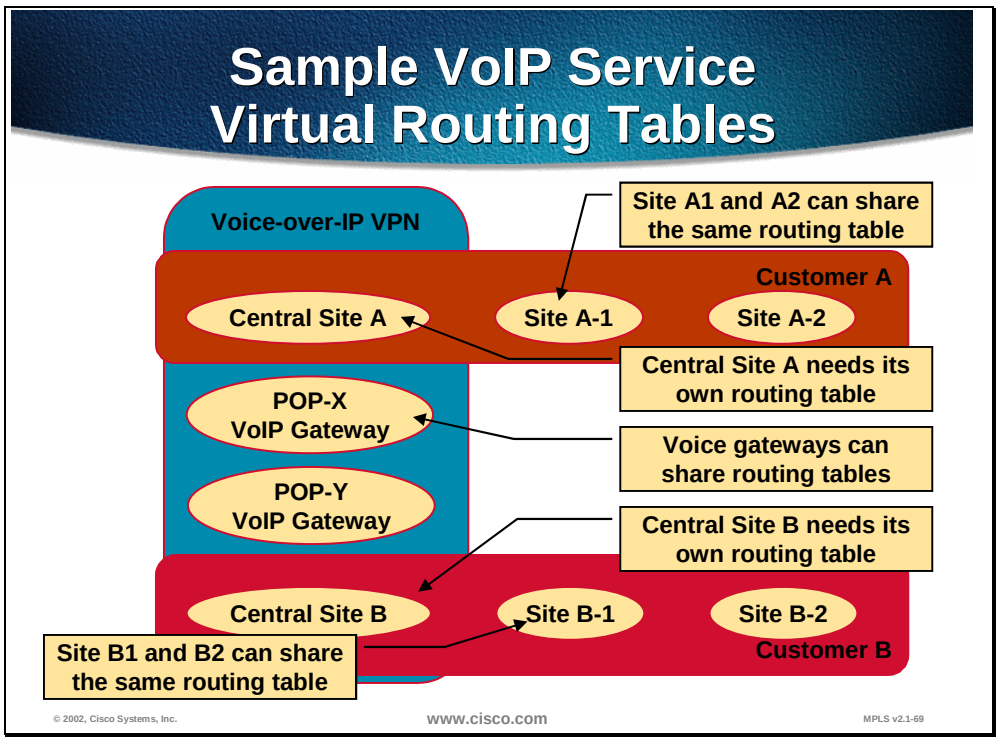

To illustrate the requirements for multiple virtual routing tables, consider the sample VoIP service with 3 VPNs (Customer A VPN, Customer B VPN, and the Voice-over-IP VPN). The following five virtual routing tables are needed to implement this service:

- All sites of customer A (apart from the central site) can share the same virtual routing table, as they only belong in a single VPN
- The same is true for all sites of Customer B (apart from the central site)
- The VoIP gateways are only participating in 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.

Practice

Q1) What is the impact of complex VPN topologies on virtual routing tables in the PE routers?

A) There is no impact. MPLS/VPNs were designed to be scalable.

B) Complex VPN topologies might require more than one VRF per VPN.

C) Customers may need to use BGP with the service provider's PE router.

D) Customers may need to use OSPF with the service provider's PE router.

Benefits of MPLS VPN versus other Peer-to-Peer VPN Technologies

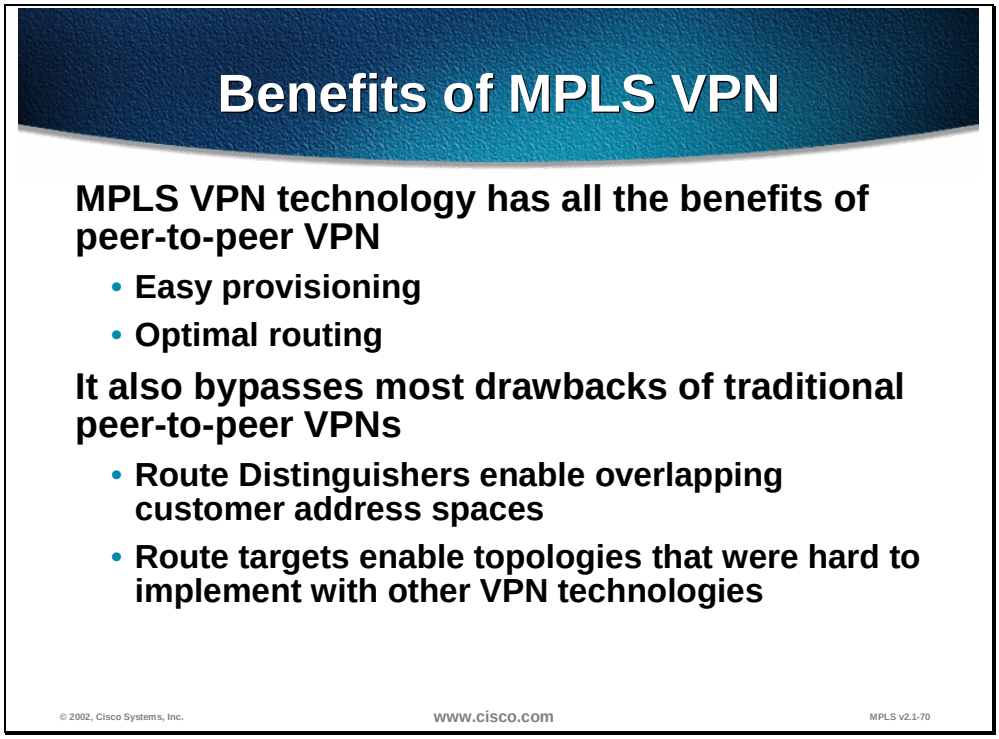

MPLS VPN architecture combines the benefits of peer-to-peer VPN paradigm with the benefits of the overlay VPN paradigm while avoiding most of the drawbacks of both of them:

- Like all peer-to-peer VPNs, MPLS VPN is easier to provision and provides automatic optimum routing between customer sites
- Like the overlay VPNs, MPLS VPN allow overlapping customer address space through the use of *route distinguishers*, 64-bit quantities that make overlapping customer addresses globally unique when prepended to them.

Another building block of MPLS VPN architecture, *route targets*, allow you to build complex VPN topologies that far surpass anything that can be built with peer-to-peer VPNs.

Summary

After completing this lesson, you should be able to perform the following tasks:

- Describe the difference between traditional peer-to-peer models and MPLS VPN
- List the benefits of MPLS VPN
- Describe major architectural blocks of MPLS VPN
- Explain the need for route distinguishers and route targets

Next Steps

After completing this lesson, go to:

- MPLS VPN Routing Model

Lesson Review

Instructions

Answer the following questions:

1. How does MPLS VPN support overlapping customer address spaces?
2. How are customer routes exchanged across the P-network?
3. What is a route distinguisher?
4. Why is the RD not usable as VPN identifier?
5. What is a route target?
6. Why were the route targets introduced in MPLS VPN architecture?
7. How are route targets used to build virtual routing tables in the PE routers?
8. What is the impact of complex VPN topologies on virtual routing tables in the PE routers?

MPLS VPN Routing Model

Overview

This lesson will help you understand how the VPN routing information is exchanged in an MPLS VPN network from both an end user and provider point of view.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the routing model of MPLS VPN
- Describe the MPLS VPN routing model from customer and provider perspectives
- Identify the routing requirements of CE-routers, PE-routers and P-routers

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- MPLS VPN Routing Model
- End-to-End Routing Information Flow
- Routing Requirements of CE-routers, PE-routers and P-routers
- Summary

MPLS VPN Routing Model

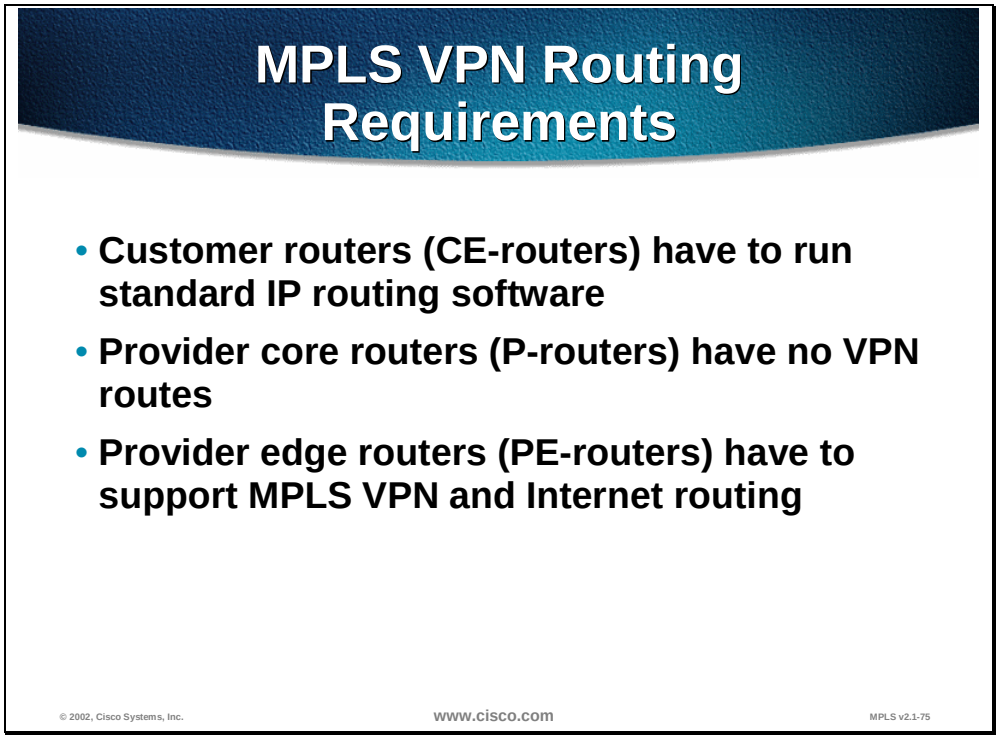

The designers of MPLS VPN technology were faced with the following routing requirements:

- The customer routers should not be MPLS VPN-aware. They should run standard IP routing software
- The provider core routers (P-routers) must not carry VPN routes to make the MPLS VPN solution scalable
- The provider edge routers (PE-routers) must support MPLS VPN services and traditional Internet services.

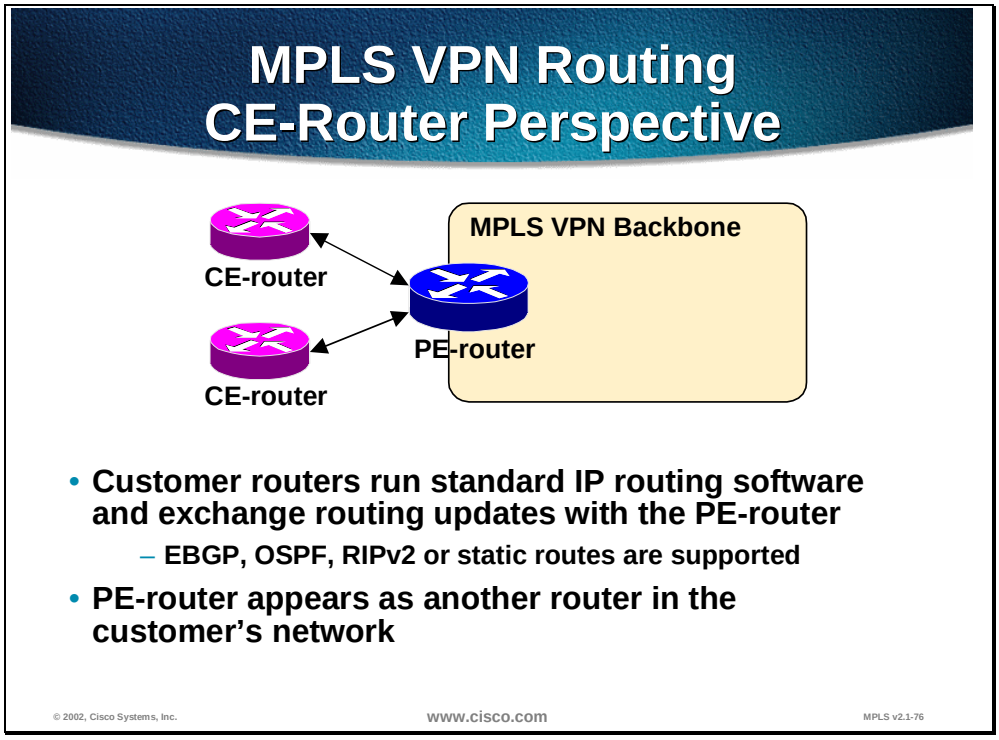

The MPLS VPN backbone should look like a standard corporate backbone to the CE-routers. The CE-routers run standard IP routing software and exchange routing updates with the PE-routers that appear as to them as normal routers in customer's network.

Note In Cisco IOS 12.1, the choice of routing protocols that can be run between CE-router and PE-router is limited to static routes, RIP version 2, OSPF and external BGP.

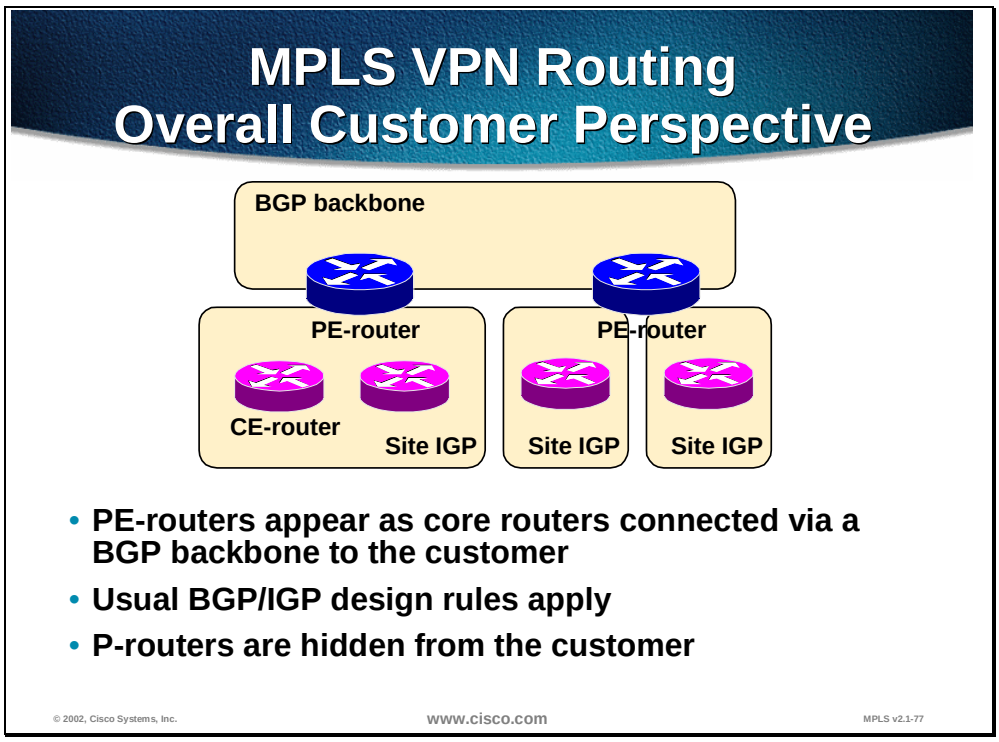

From the customer's network designer, the MPLS VPN backbone looks like intra-company BGP backbone with PE-routers performing the route redistribution between individual sites and the core backbone. The standard design rules that are used for enterprise BGP backbones can be applied to the design of the customer's network.

The P-routers are hidden from the customer's view; the internal topology of the BGP backbone is therefore totally transparent to the customer.

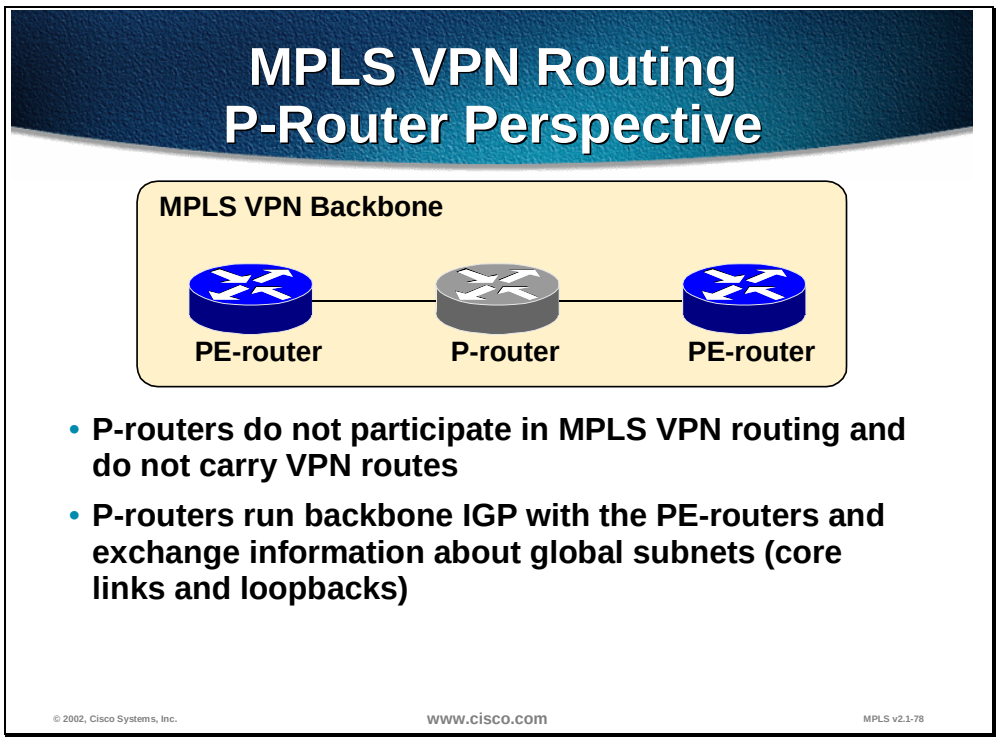

From the P-router perspective, the MPLS VPN backbone looks even simpler – the P-routers do not participate in MPLS VPN routing and do not carry VPN routes. They only run backbone IGP with other P-routers and with PE-routers and exchange information about core subnets. BGP deployment on P-routers is not needed for proper MPLS VPN operation; it might be needed, however, to support traditional Internet connectivity that was not yet migrated to MPLS.

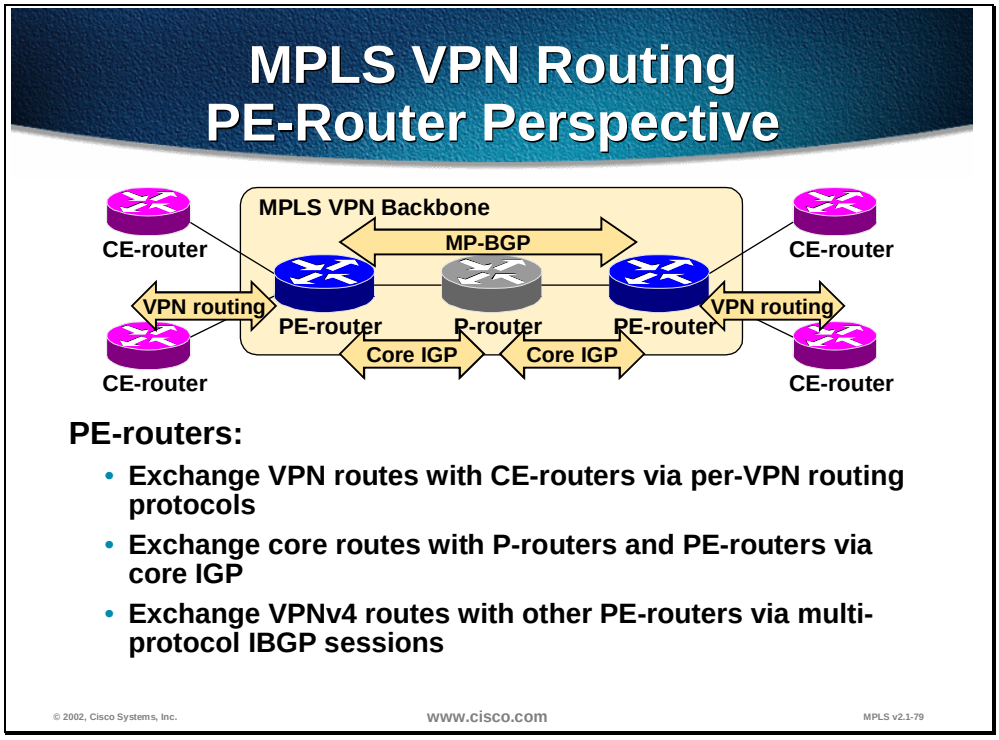

The PE-routers are the only routers in the MPLS VPN architecture that see all routing aspects of the MPLS VPN:

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

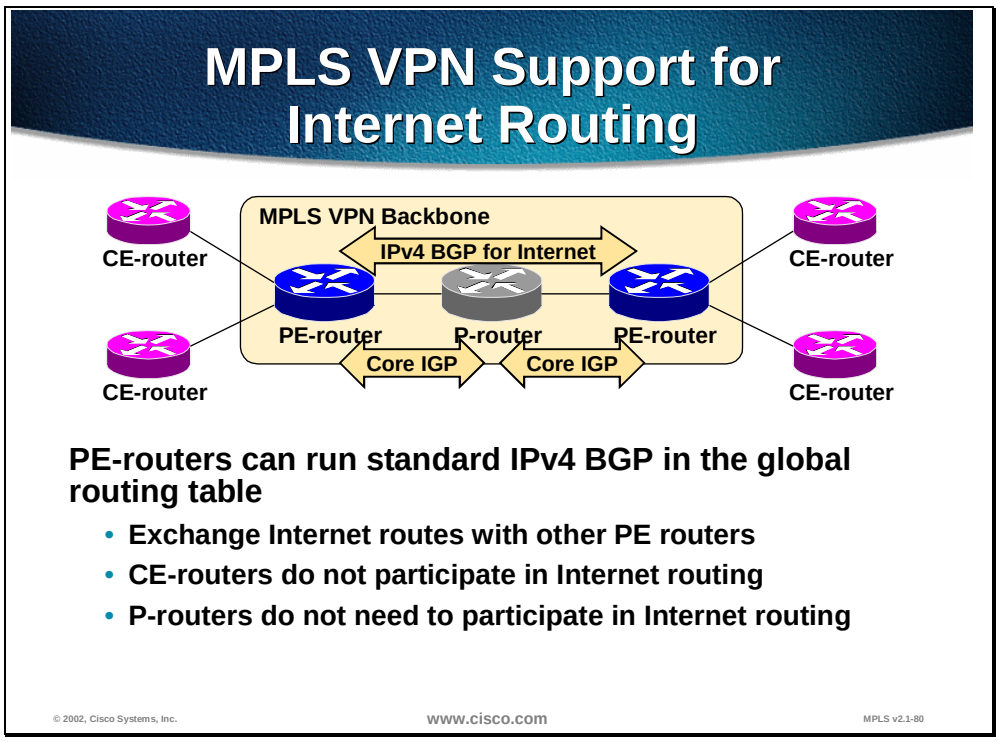

The routing requirements for PE-routers also extend to supporting Internet connectivity - PE-routers have to exchange Internet routes with other PE-routers. The CE-routers cannot participate in Internet routing if the Internet routing is performed in global address space. The P-routers could participate in Internet routing, however, you should disable Internet routing on the P-routers to make your network core more stable (please see the design guidelines in **Core MPLS Technology** module for more details).

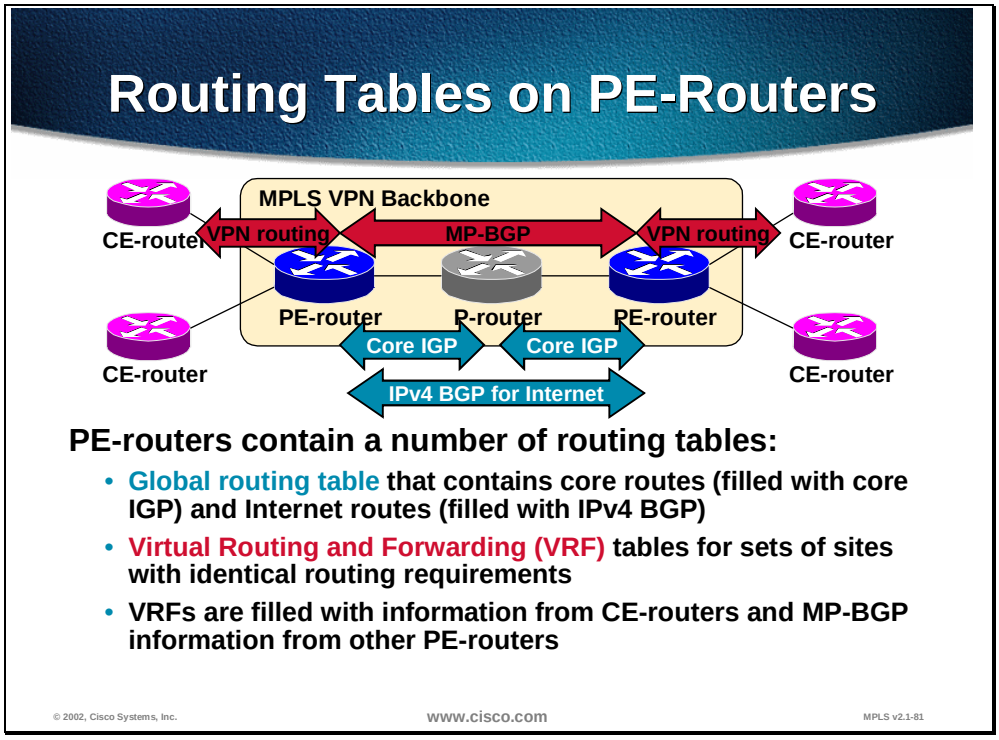

The PE-routers support various routing requirements imposed on them by using a number of IP routing tables:

- The global IP routing table (the IP routing table that is always present in an IOS-based router even if it's not running MPLS VPN) contains all core routes (inserted by core IGP) and the Internet routes (inserted from global IPv4 BGP table)
- The Virtual Routing and Forwarding (VRF) tables contain sets of routes for sites with identical routing requirements. The VRFs are filled with intra-VPN IGP information exchanged with the CE-routers and with VPNv4 routes received through MP-BGP sessions from the other PE-routers

Practice

Q1) What is the P-router perception of end-to-end MPLS VPN routing?

A) P-router takes part in the customers routing and exchange routes with the CE-routers.

B) P-router is fully MPLS VPN aware. It has a virtual routing table for each VPN where it stores customer routes.

C) P-router is not MPLS VPN aware. It only sees global subnets in the MPLS VPN backbone, not the customer routes.

End-to-End Routing Information Flow

The following slides give you an overview of end-to-end routing information flow in an MPLS VPN network.

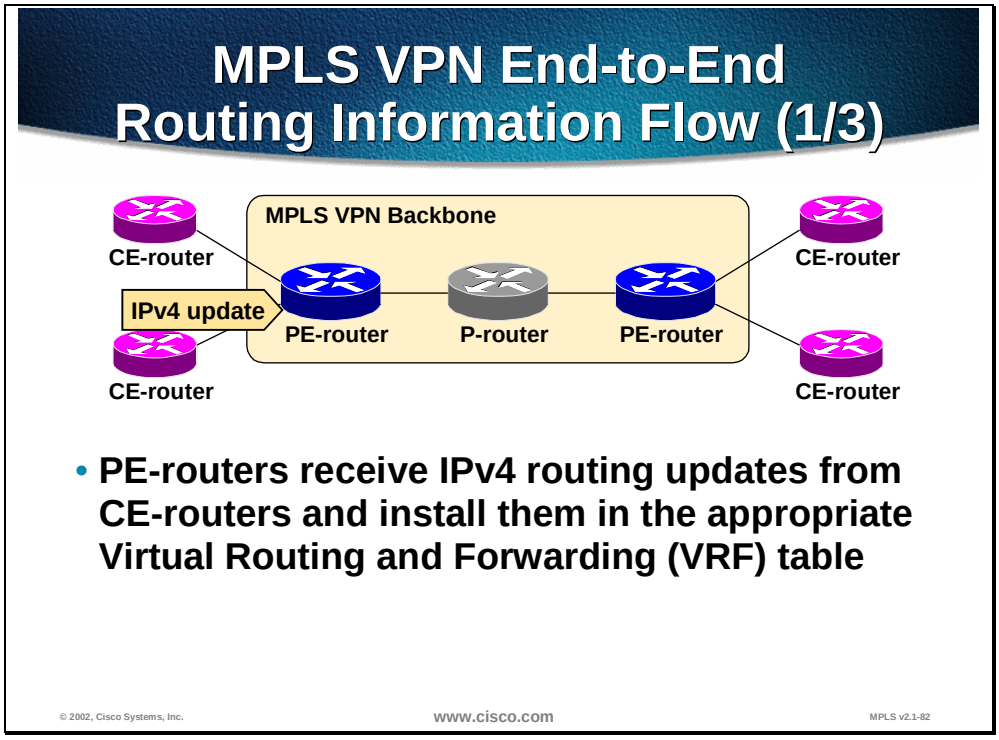

The PE-routers receive IPv4 routing updates from the CE-routers and install them in appropriate Virtual Routing and Forwarding table.

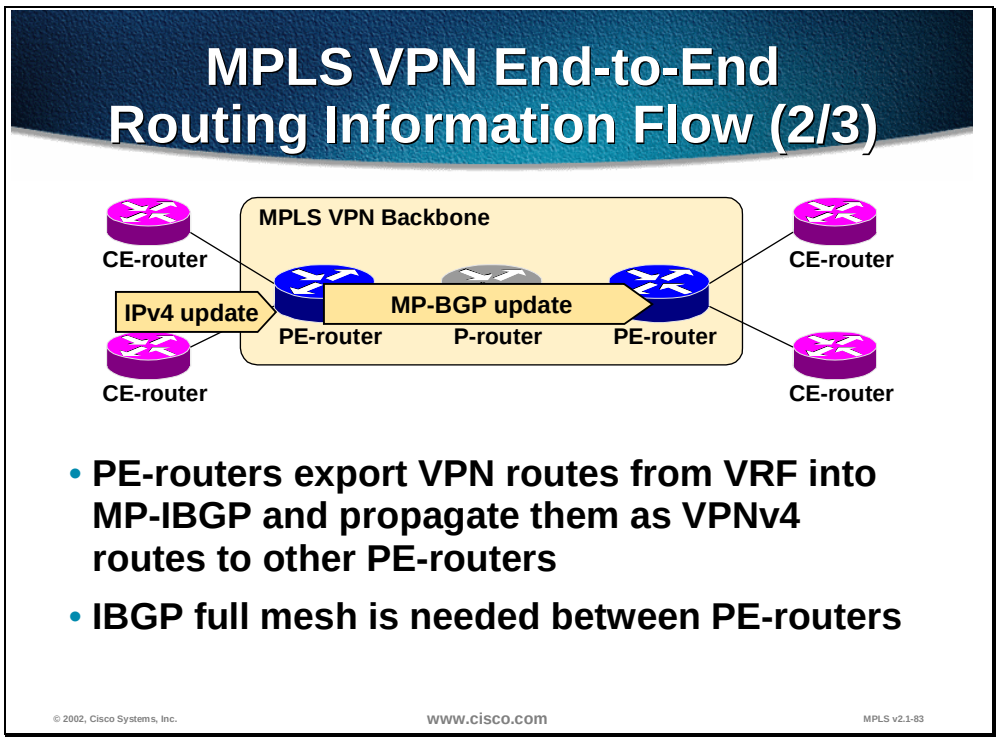

The customer routes from VRFs tables are exported as VPNv4 routes into MP-BGP and propagated to other PE-routers.

The MP-BGP sessions between the PE-routers are therefore IBGP sessions and are subject to the IBGP split horizon rules. Full mesh of MP-IBGP sessions is thus required between PE-routers or you could use route reflectors to reduce the full mesh IBGP requirement.

The MPLS/VPN architecture also support MP-EBGP sessions. This topic is covered in the lesson **MPLS VPN Spanning more than One AS.**

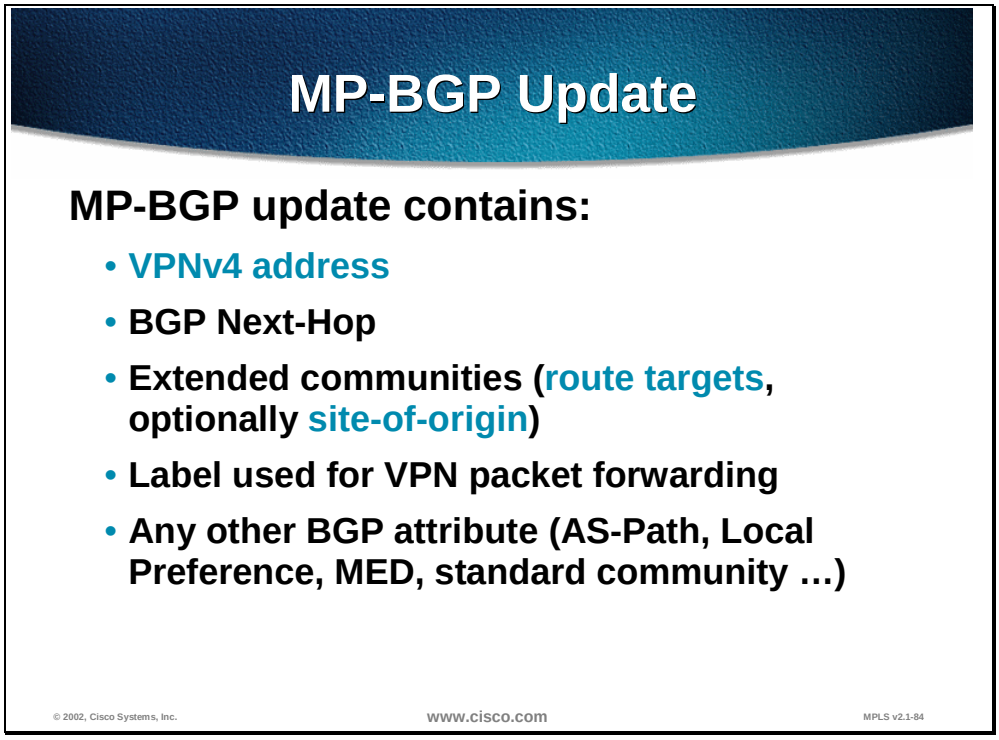

Multi-protocol BGP update exchange between PE-routers contains:

- VPNv4 address
- The BGP next-hop IP address
- Extended BGP communities (route targets are required, site of origin is optional)
- Label used for VPN packet forwarding (the “MPLS VPN Packet Forwarding” lesson later in this lesson explains how the label is used)
- Mandatory BGP attributes (for example, AS-path)

Optionally, the MP-BGP update can contain any other BGP attribute, for example, local preference, MED or standard BGP community.

The BGP next-hop is, normally, the IP address of the PE router. This information is used by the receiving PE router to make a lookup in the LIB and find a label switched path to the remote PE. This will be the next-hop label.

The extended community attribute RT is used by the receiving PE to determine into which virtual routing table(s) to import the route.

Site of origin (SOO) is an optional extended community attribute which is used to recognize from which site a route was received in order to block propagation of the route back to the same site. This is only applicable when a site has multiple access links to the provider network, and there is no other means of detecting a potential routing loop.

The VPN label will be used in combination with the next-hop label to form a label stack. The next-hop label will be the top-most label which directs packets across the MPLS/VPN backbone to the remote PE. The VPN label is assigned by the remote PE and will be used by it to forward the packets to the correct site.

All other BGP attributes are available to implement various routing policies and preferences.

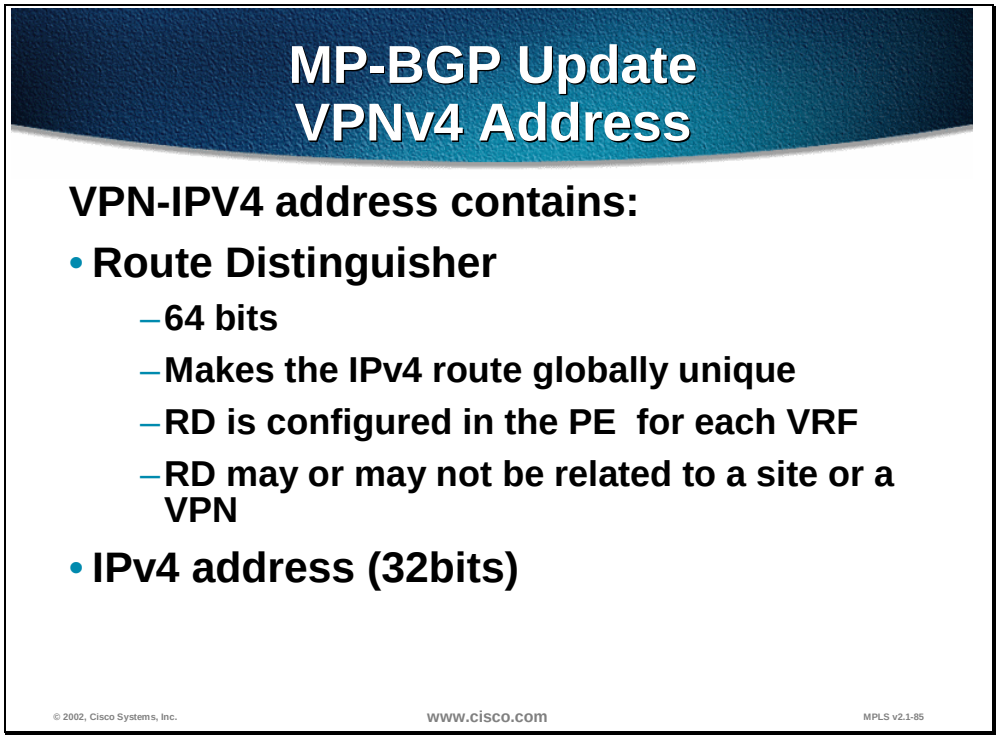

The VPNv4 address propagated in the MP-BGP update is composed of a 64-bit route distinguisher and the 32-bit customer IPv4 address. The route distinguisher is configured in the virtual routing and forwarding table on the PE-router.

In simple VPN topologies, where all sites in a VPN have identical routing requirements, the route distinguisher may be related to a VPN.

In other complex VPN topologies, every site may require a dedicated VRF based on the connectivity requirements. In this case, the RD may be related to a particular site rather than to a particular VPN.

In general, however, there is no clear correspondence between route distinguisher and either customer VPN or customer site.

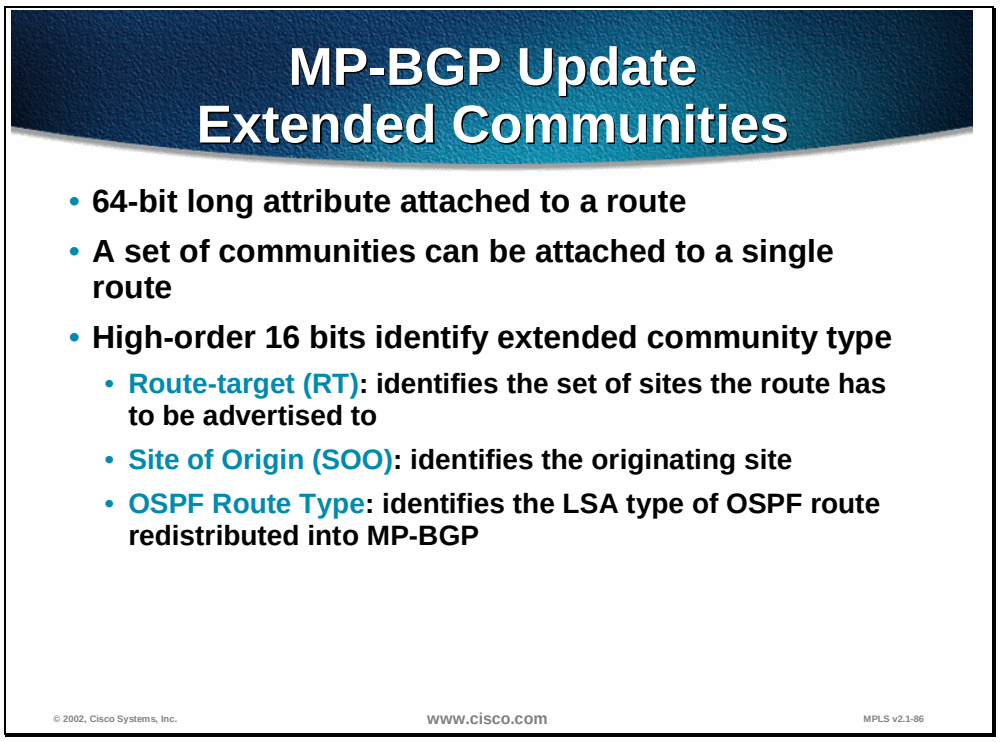

Extended BGP communities (at least route targets) are always attached to the VPNv4 routes in MP-BGP updates. These communities are 64-bit long attributes, where the high-order 16 bits identify the community meaning and the network administrator defines the low-order 48 bits.

So far, three extended community types have been defined:

- *Route target*, which is used to indicate VPN membership of a customer route. Route targets are used to facilitate transfer of customer routes between virtual routing and forwarding tables.
- *Site of origin (SOO)*, which identifies the customer site originating the route. Site of origin is used to prevent routing loops in network designs with multihomed sites and there are no other means of detecting routing loops. An SOO value may optionally be assigned to all routes received on any of the links from a multihomed site. The SOO value can then be checked before the any route is propagated to the site. Routes received from a specific site (regardless on which link) is not sent back to the same site on any of the links.
- *OSPF route type*, which identifies the LSA type of an OSPF route converted into MP-BGP VPNv4 route.

The following values are used in the high-order 16 bits of the extended BGP community to indicate community type:

Community type	Value in high-order 16 bits
Route target	0x0002
Site of origin	0x0003
OSPF route type	0x8000

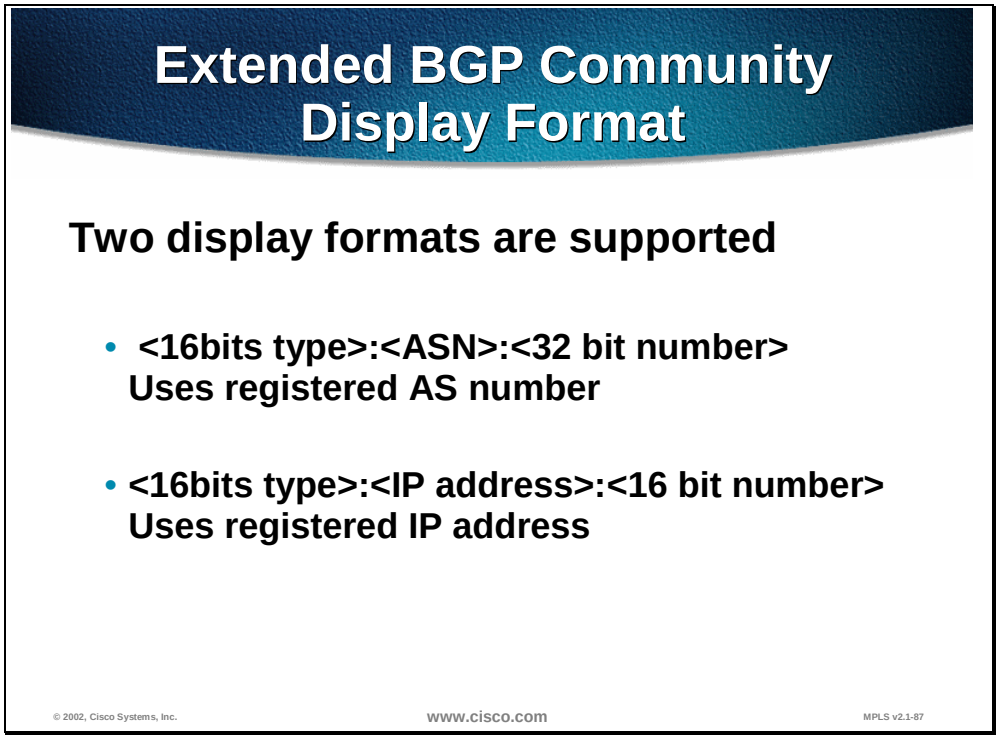

The low-order 48 bits of the extended BGP community can be displayed in two different formats:

- Higher-order 16 bits are the public AS number of the Service Provider defining the community, lower-order 32 bits are defined by the network administrator. This is the recommended format
- Higher-order 32 bits are a public IP address belonging to the Service Provider defining the community; the network administrator defines lower-order 16 bits

The display format is encoded in one of the high-order 16 bits of the extended community to ensure consistent formatting across all routers participating in an MPLS VPN network.

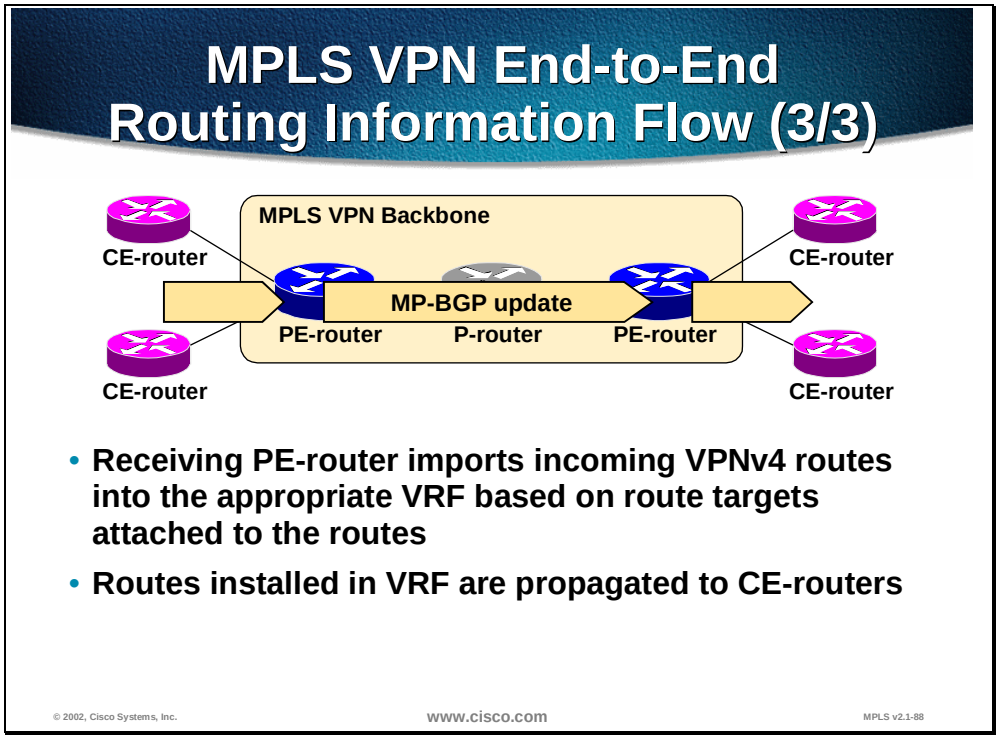

The PE-routers receiving MP-BGP updates will import the incoming VPNv4 routes into their VRFs based on route targets attached to the incoming VPNv4 routes and import route targets configured in the VRFs. The VPNv4 routes installed in VRFs are converted to IPv4 routes then propagated to the CE-routers.

Practice

Q1) How is the VPN routing information exchanged between the PE-routers?

A) PE-routers are not VPN aware and need not to exchange VPN routing information.

B) PE-routers exchange VPN routing information with MP-BGP.

C) PE-routers exchange VPN routing information with IPv4 routing updates.

Routing Requirements of CE-routers, PE-routers and P-routers

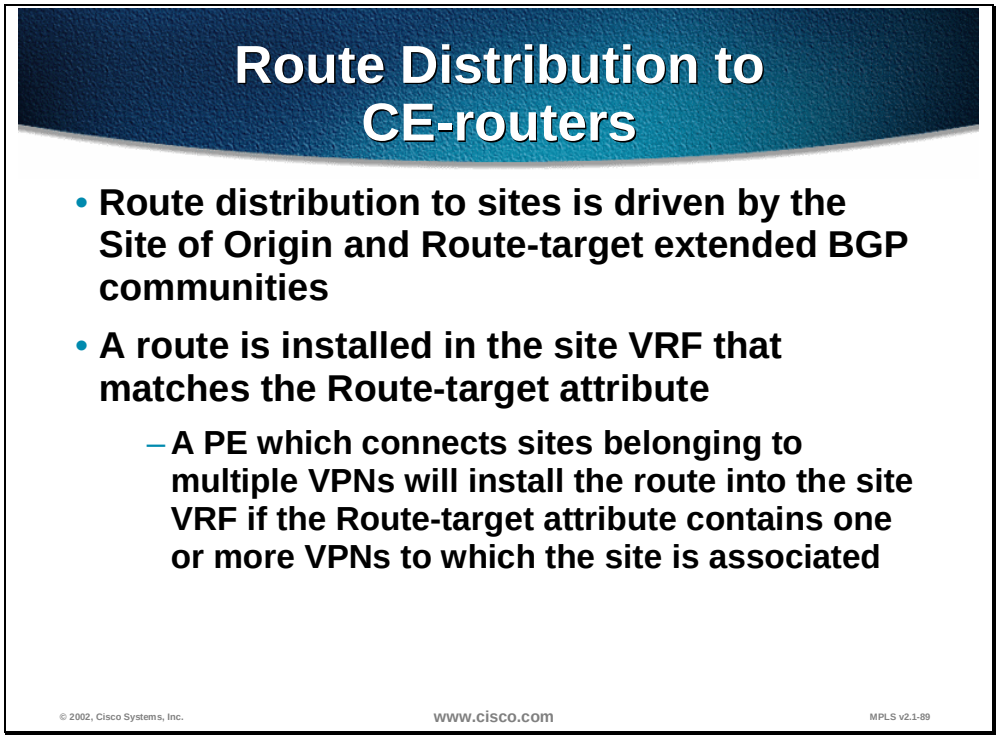

The route targets attached to a route and the import route targets configured in the VRF drive the import of VPNv4 routes into VRFs on the receiving PE-router – the incoming VPNv4 route is imported into the VRF only if at least one route target attached to the route matches at least one import route target configured in the VRF.

The optional site-of-origin 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 site-of-origin attached to the route is equal to the site-of-origin attribute associated with the CE-router. The site-of-origin can thus be used to prevent routing loops in MPLS VPN networks with multihomed sites where the routing protocol used between PE and CE routers does not itself prevent the loop.

Practice

Q1) Which BGP attributes control the VPN route distribution toward CE-routers?

A) MED

B) Local preference

C) Route targets

D) Site-of-origin

Summary

After completing this lesson, you should be able to perform the following tasks:

- Describe the routing model of MPLS VPN
- Describe the MPLS VPN routing model from customer and provider perspective
- Identify the routing requirements of CE-routers, PE-routers and P-routers

Next Steps

After completing this lesson, go to:

- MPLS VPN Packet Forwarding

Lesson Review

Instructions

Answer the following questions:

1. What is the impact of MPLS VPN on CE-routers?
2. What is the customer's perception of end-to-end MPLS VPN routing?
3. What is the P-router perception of end-to-end MPLS VPN routing?
4. How many routing tables does a PE-router have?
5. How many routing tables reside on a P-router?
6. Which routing protocols fill the global routing table of a PE-router?
7. Which routing protocols fill the Virtual Routing table of a PE-router?
8. How is the Internet routing supported by MPLS VPN architecture?
9. How is the VPN routing information exchanged between the PE-routers?
10. Which attributes are always present in a MP-BGP update?
11. Which attributes can be optionally present in a MP-BGP update?
12. Which BGP attributes drive the import of VPNv4 route into a VRF?
13. Which BGP attributes control the VPN route distribution toward CE-routers?

MPLS VPN Packet Forwarding

Overview

The lesson documents the MPLS VPN forwarding mechanisms.

Importance

This lesson is the foundation lesson for the MPLS VPN Curriculum.

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the MPLS VPN forwarding mechanisms
- Describe the VPN and backbone label propagation
- Explain the need for end-to-end LSP between PE routers
- Explain the implications of BGP next-hop on MPLS VPN forwarding

Learner Skills and Knowledge

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

- Cisco Certified Network Professional (CCNP) level of knowledge or equivalent level of IP routing and Cisco IOS knowledge;
- Core MPLS knowledge
- Advanced BGP knowledge,

Optional knowledge:

- ATM knowledge,
- OSPF or IS-IS knowledge
- MPLS Traffic Engineering and associated prerequisites
- MPLS Quality of Service and associated prerequisites

Mandatory prerequisite modules:

- MPLS Core Services
- BGP Curriculum

Optional prerequisite modules:

- MPLS Quality of Service
- MPLS Traffic Engineering
- ATM curriculum
- OSPF or IS-IS curriculum

Outline

This lesson includes these lessons:

- Overview
- MPLS VPN Forwarding
- VPN Label Propagation
- Impact of BGP Next-Hop Processing on MPLS VPN Forwarding
- Impact of IGP Route Summarization on MPLS VPN Forwarding
- Summary

MPLS VPN Forwarding

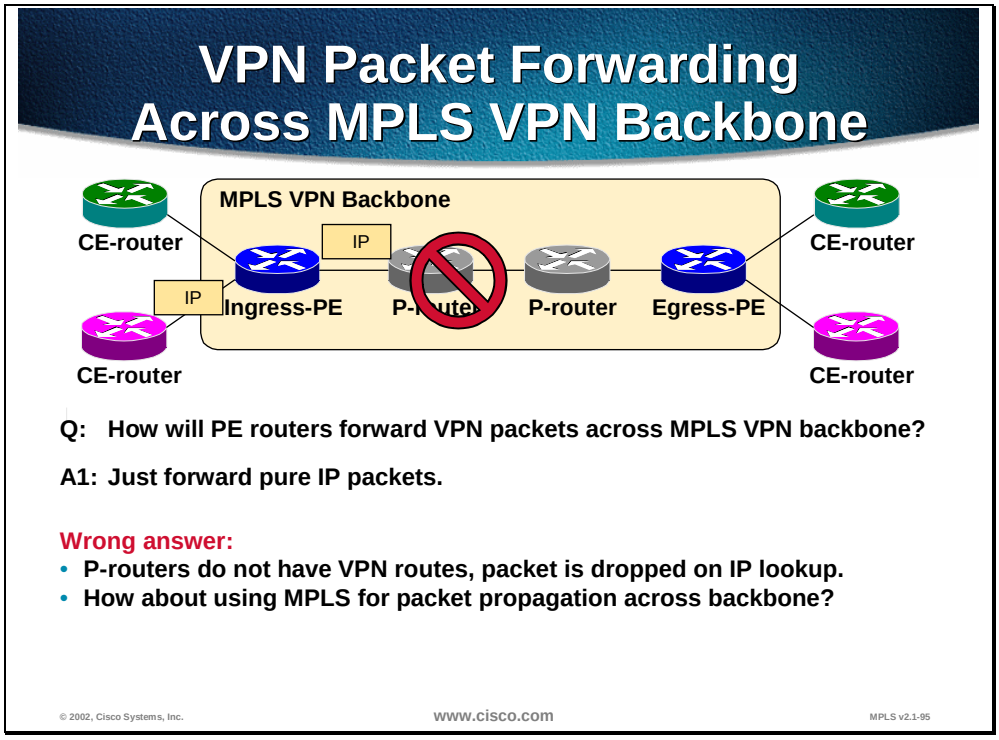

With the customer routes being propagated across MPLS VPN backbone, all the routers are ready to start forwarding customer data. The customer traffic between CE-routers and PE-routers is always sent as pure IP packets, satisfying the requirement that the CE-routers run standard IP software and are not MPLS VPN-aware.

In a very simplistic approach to packet forwarding across MPLS VPN backbone, the PE-routers might just forward IP packets received from the customer routers toward other PE-routers. This approach would clearly fail, as the P-routers have no knowledge of the customer routes and therefore cannot forward customer IP-packets. A better approach would be to use MPLS Label Switched Path (LSP) between PE-routers and a label to determine the proper LSP to use.

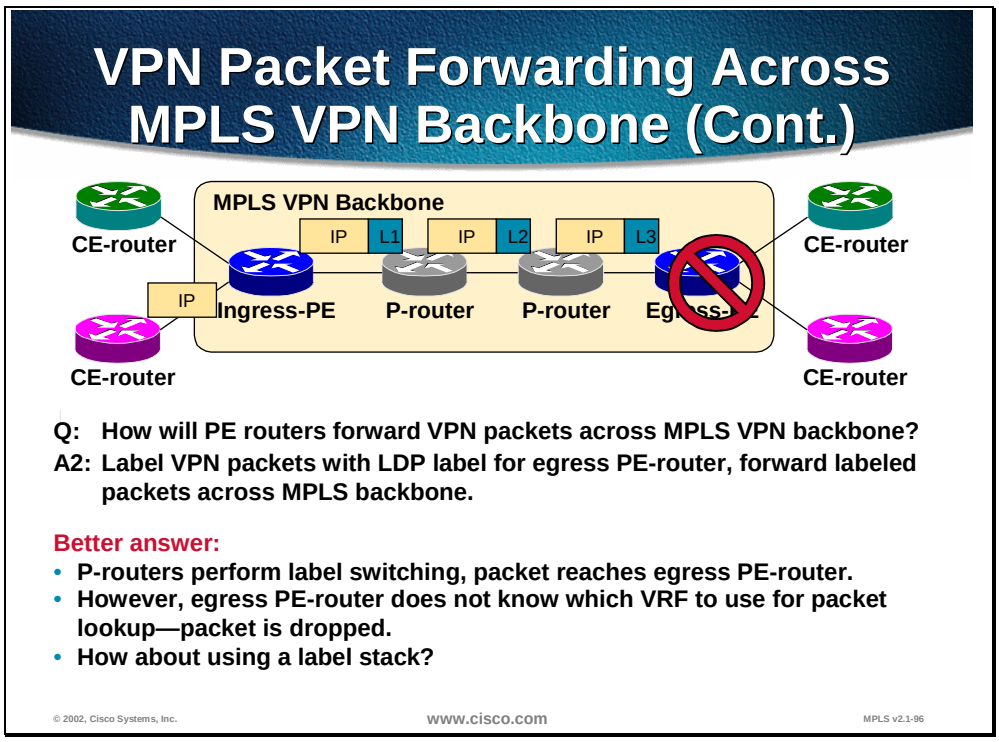

An MPLS-oriented approach to MPLS VPN packet forwarding across the MPLS VPN backbone would be to label the customer packet with the LDP-assigned label for egress PE-router. The core routers would consequently never see the customer IP packet, just a labeled packet targeted toward egress PE-router. They would perform simple label switching operations, finally delivering the customer packet to the egress PE-router. Unfortunately, the customer IP packet contains 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.

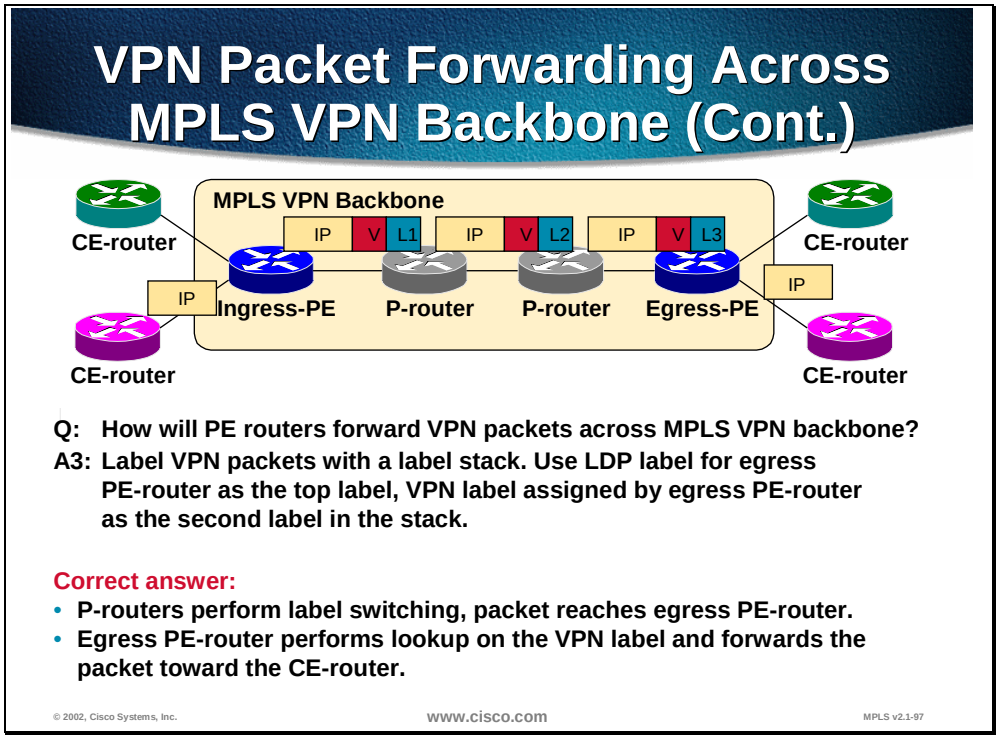

MPLS label stack can be used to indicate to the egress PE-router what to do with the VPN packet. When using the label stack, the ingress PE-router labels incoming IP packet with two labels. The top label in the stack is the LDP label for the egress PE-router that will guarantee that the packet will traverse 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 in the stack could point directly toward an outgoing interface, in which case the egress PE-router only performs label lookup on the VPN packet. The second label could also point to a VRF, in which case the egress PE-router performs a label lookup first to find the target VRF and then performs an IP lookup within the VRF.

Both methods are used in Cisco IOS. 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 **null** interface and routes for directly connected VPN interfaces.

Two-level MPLS label stack satisfies all MPLS VPN forwarding requirements:

- P-routers perform label switching on the LDP-assigned label toward the egress PE-router
- Egress PE-router performs label switching on the second label (that 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.

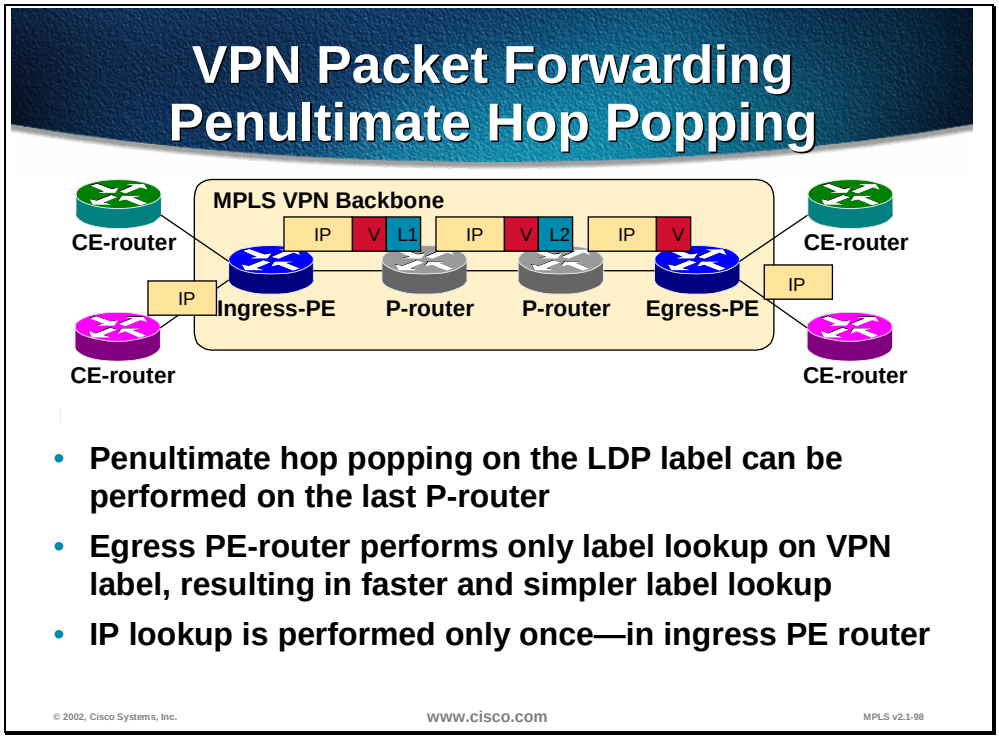

Penultimate hop popping (removal of top label in the stack on hop prior to the egress router) can be performed in frame-based MPLS networks. In these networks, the last P-router in the label switched path pops the LDP label (as previously requested by the egress PE-router through LDP) and the PE-router 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 Forwarding Information Base (FIB) is therefore performed only once – in the ingress PE-router; even without the penultimate hop popping.

Note Please refer to MPLS Technology chapter for more information on penultimate hop popping.

Practice

Q1) How are VPN packets propagated across MPLS VPN backbone?

A) They are propagated with a VPN label instead of the TDP/LDP label.

B) They are propagated with VPNv4 destination and source addresses in the header.

C) They are propagated as standard IPv4 packets.

D) They are propagated across MPLS VPN backbone with two labels in the MPLS label stack.

VPN Label Propagation

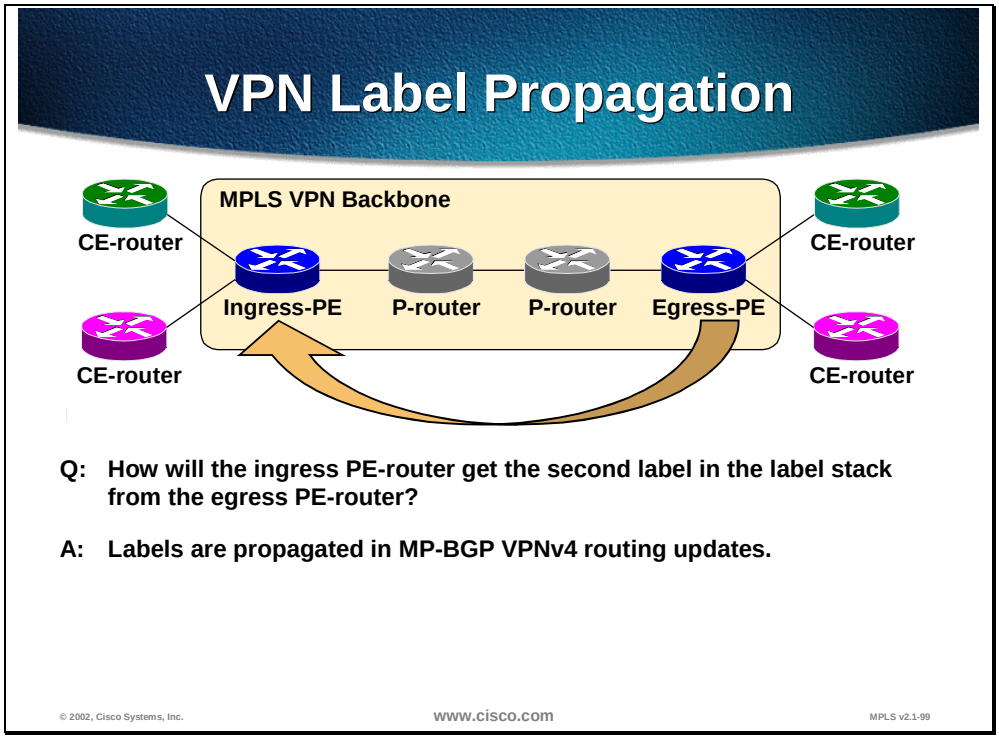

The MPLS label stack, with the second label being assigned by the egress PE-router, is mandatory for proper MPLS VPN operation. These labels have to be propagated between PE-routers to enable proper packet forwarding and MP-BGP was chosen as the propagation mechanism. Every MP-BGP update thus carries a label assigned by the egress PE-router together with the 96-bit VPNv4 prefix.

The following slides illustrate the VPN label propagation between PE-routers.

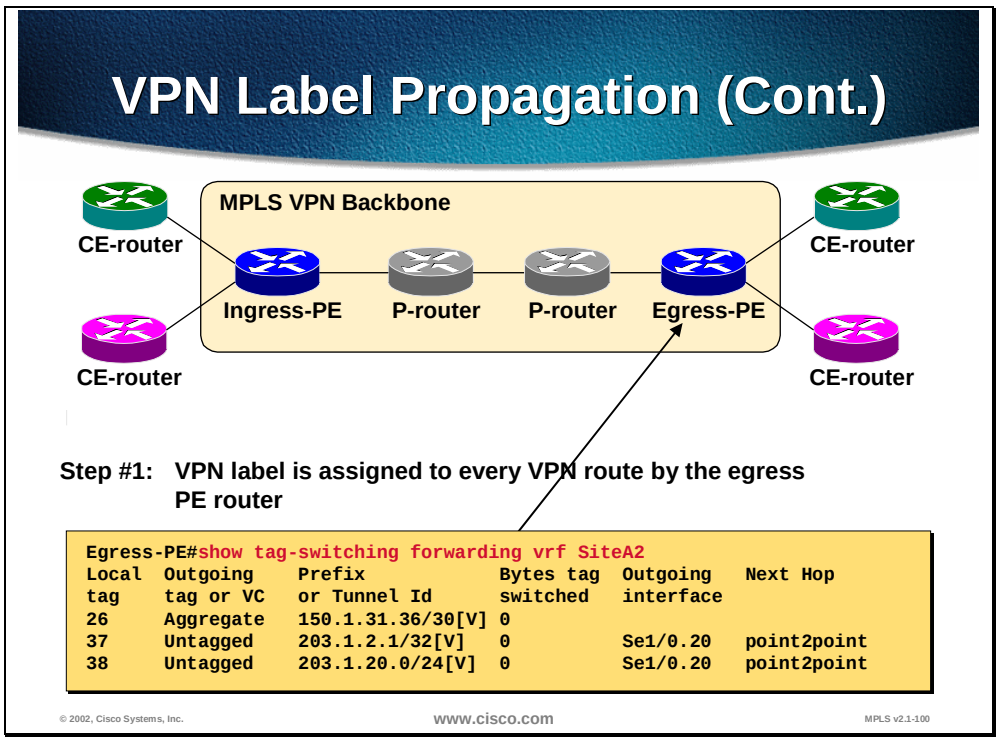

Step 1 Egress PE-routers assign a label to every VPN route received from attached CE-routers and to every summary route summarized inside the PE-router. This label is then used as the second label in the MPLS label stack by the ingress PE-routers when labeling VPN packets.

The VPN labels assigned locally by the PE-router can be inspected with the **show tag-switching forwarding vrf** command.

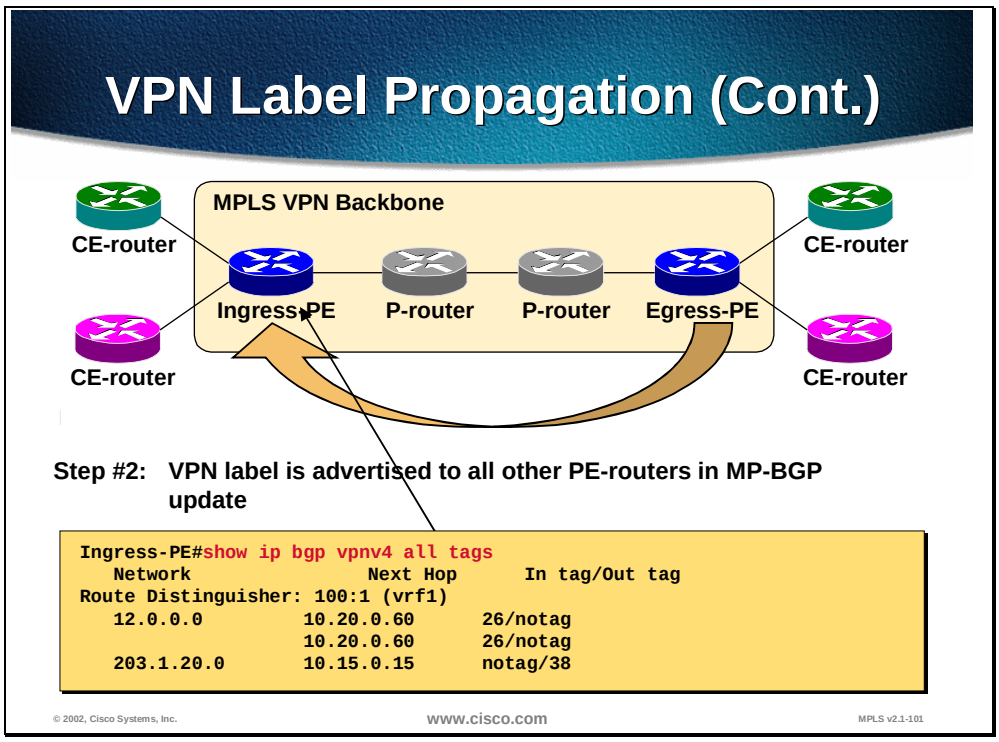

Step 2 VPN labels assigned by the egress PE-routers are advertised to all other PE-routers together with VPNv4 prefix in MP-BGP updates.

These 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 203.1.20.0 is 38 and was assigned by another PE-router (Egress-PE in the previous slide).

Note Like many other IOS **show** commands, the **show ip bgp vpnv4 tags** command uses the old terminology – labels are still called tags.

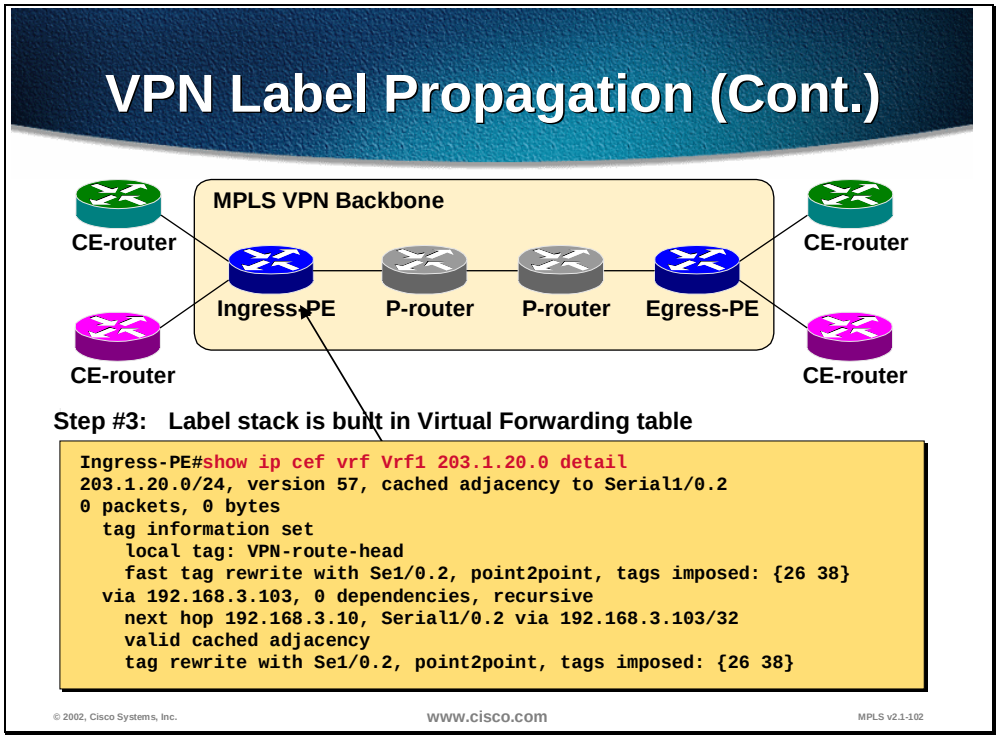

Step 3 The ingress PE-router has two labels associated with a remote VPN route – a label for BGP next-hop assigned by the next-hop P-router via LDP (and taken from local Label Information Base – LIB) as well as the label assigned by remote PE-router and propagated via MP-BGP update. Both labels are combined in a label stack and installed in the virtual forwarding (VRF) table.

The label stack in the virtual forwarding table can be inspected with 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 TDP/LDP label toward the egress PE-router and the second label is the VPN label advertised by the egress PE-router.

Practice

Q1) Which router assigns the VPN label?

A) Every PE- or P-router in the path generates its own local label.

B) The advertising CE-router also assigns the label.

C) The egress PE-router assigns the VPN label.

D) The ingress PE-router assigns the VPN label.

Impact of BGP Next-Hop Processing on MPLS VPN Forwarding

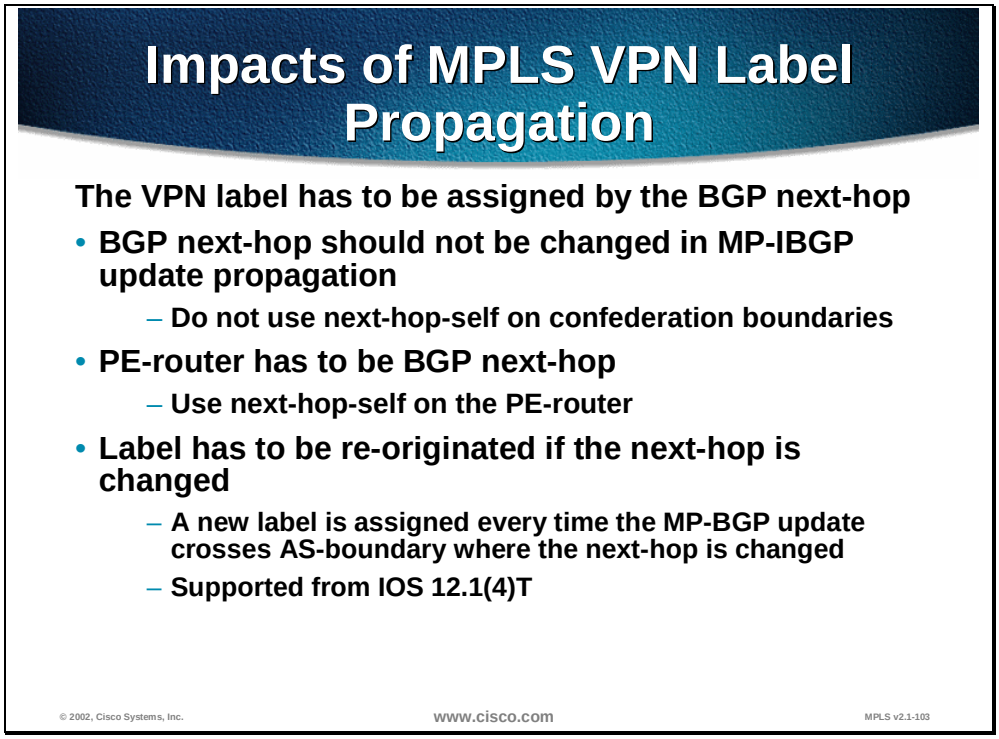

MPLS VPN packet forwarding works optimally if and only if the router specified as the BGP next-hop in incoming BGP update is the same router as the one that has assigned the second label in the label stack. There are three scenarios that can cause the BGP next hop to be different from the IP address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

- If the customer route is received from the CE-router via external BGP session, the next-hop of the VPNv4 route is still the IP address of the CE-router (BGP next hop of an outgoing IBGP update is always identical to the BGP next hop of the incoming EBGP update). In older software, it used to be required 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 autonomous system. It can change, however, if you use **next-hop-self** on inter-AS boundary inside a BGP confederation or if you use inbound route-map on a PE-route to change next-hop (a strongly discouraged practice). To prevent this, make sure that you never change BGP next-hop with a **route-map** or **next-hop-self** inside an autonomous system.
- The BGP next hop is always changed on an external BGP session. If the MPLS VPN network spans multiple public autonomous systems (not just autonomous systems within a BGP confederation), special provisions are

made automatically by IOS in the AS boundary routers to re-originate the VPN label at the same time as the BGP next hop is changed. This functionality is called LSP stitching and is supported from IOS releases 12.1(5)T and 12.2.

In software which does not support LSP stitching, changing of the next-hop is fatal for the networks operation. In more recent software, where LSP stitching is supported, changing of the next-hop may introduce suboptimal routing.

The stitching forces the packet to be forwarded along the same path as the routing update has been forwarded. In the case where route-reflectors are used, the optimal transit path across the P network may not be passing via the route-reflector. But if the route-reflector does next-hop-self, the LSP stitching will force packets to take a suboptimal path.

BGP confederations may suffer from the same drawback when next-hop-self is used on intra-confederation AS boundaries, but the IGP is spanning the entire AS and may find more optimum paths between ingress PE and egress PE.

Practice

Q1) What is the impact of changing BGP next-hop on MP-BGP update?

A) MPLS VPN connectivity is broken unless the MPLS VPN label is re-originated.

B) There is no impact. MP-BGP update automatically re-originate the MPLS VPN label.

C) There is no impact. Every VPN packet is labeled with appropriate BGP next-hop.

Impact of IGP Route Summarization on MPLS VPN Forwarding

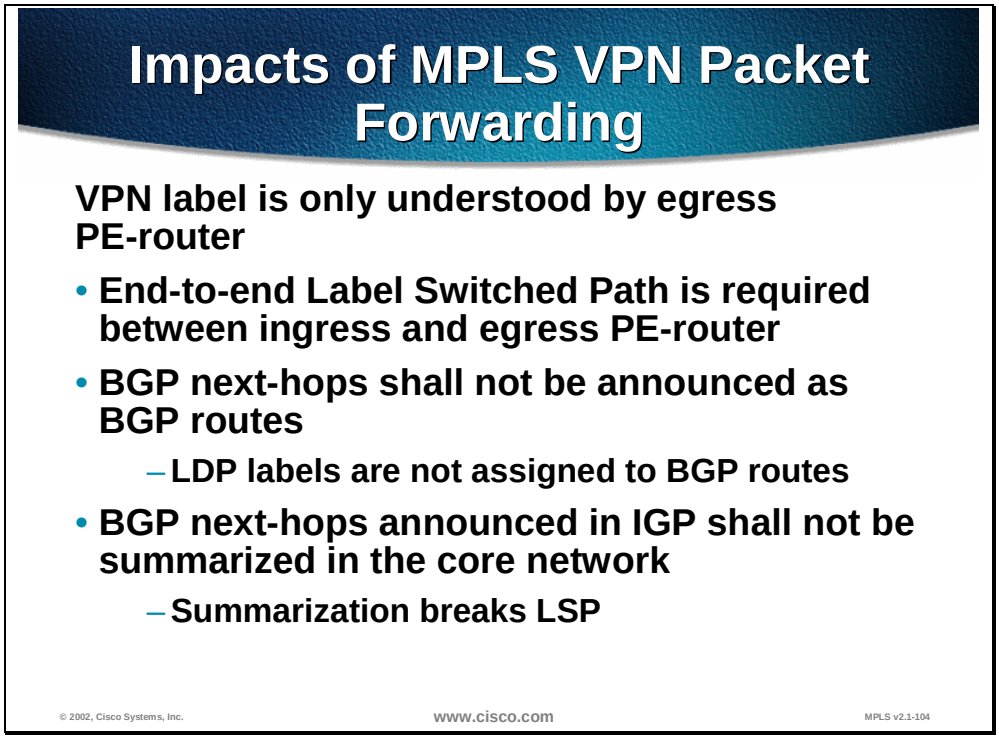

The second requirement for successful propagation of MPLS VPN packets across an MPLS backbone is an unbroken label switched path (LSP) between PE-routers. The second label in the stack is recognized only by the egress PE-router that has originated it and would not be understood by any other router, should it ever become exposed.

There are two scenarios that could cause the LSP between PE-routers to break:

- If the IP address of the PE-router is announced as a BGP route, it has no corresponding LDP label and the label stack could 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 will be disrupted at the summarization point, as illustrated on the next slide.

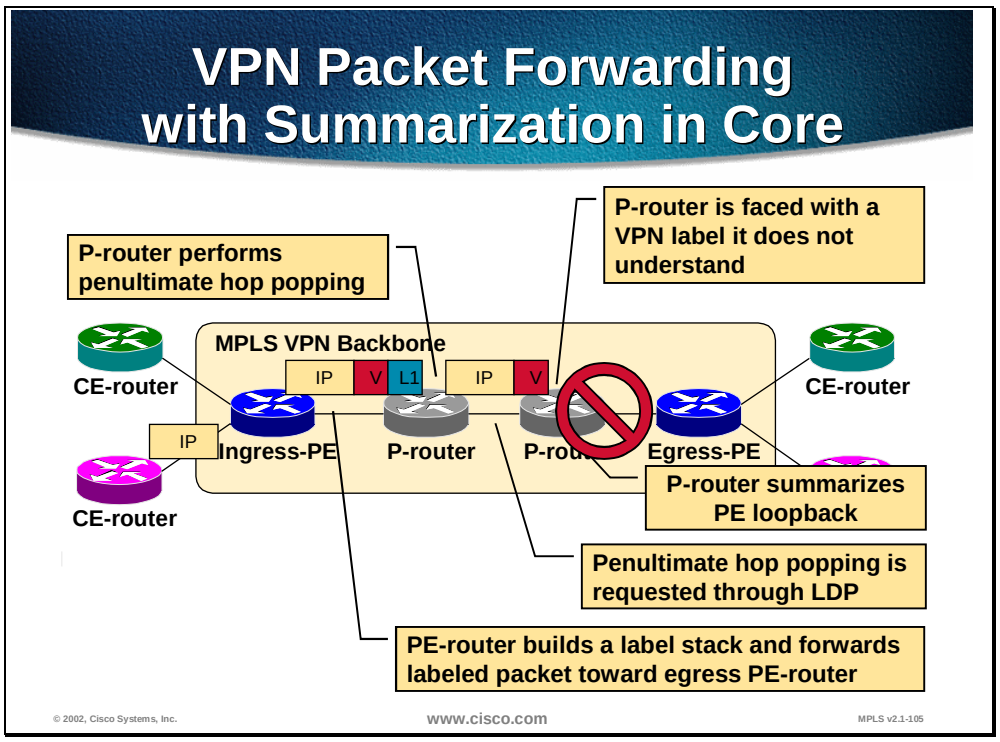

In the example above, the P-router summarizes the loopback address of the egress PE-router. LSP is broken at a summarization point, as the summarizing router needs to perform full IP lookup. In a frame-based MPLS network, the P-router would request penultimate hop popping for the summary route and the upstream P-router (or a PE-router) would remove the LDP label, exposing the VPN label to the P-router. As the VPN label was not assigned 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.

Practice

Q1) What is the impact of BGP next-hop summarization in the network core?

A) There is no impact.

B) Suboptimal routing may occur if BGP nexthops are summarized.

C) An additional routing lookup is needed.

D) MPLS VPN connectivity is broken if the BGP next-hops are summarized in the network core.

Summary

After completing this lesson, you should be able to perform the following tasks:

- Describe the MPLS VPN forwarding mechanisms
- Describe the VPN and backbone label propagation
- Explain the need for end-to-end LSP between PE routers
- Explain the implications of BGP next-hop on MPLS VPN forwarding

Next Steps

After completing this lesson, go to:

- MPLS VPN Spanning more than One AS

Lesson Review

Instructions

Answer the following questions:

1. How are VPN packets propagated across MPLS VPN backbone?
2. How can P-routers forward VPN packets if they don’t have VPN routes?
3. How is the VPN label propagated between PE-routers?
4. Which router assigns the VPN label?
5. How is the VPN label used on other PE-routers?
6. What is the impact of changing BGP next-hop on MP-BGP update?
7. How are MP-BGP updates propagated across AS boundary?
8. What is the impact of BGP next-hop summarization in the network core?

MPLS VPN Spanning more than One AS

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Explain the need for inter-AS VPNs
- Describe the propagation of VPN routes between two autonomous systems
- Explain how the label switched path in one AS is stitched together with an LSP in the other AS
- Describe the packet forwarding across AS boundaries
- Explain the benefits of using LSP stitching inside a BGP confederation

Benefits of Inter-AS VPNs

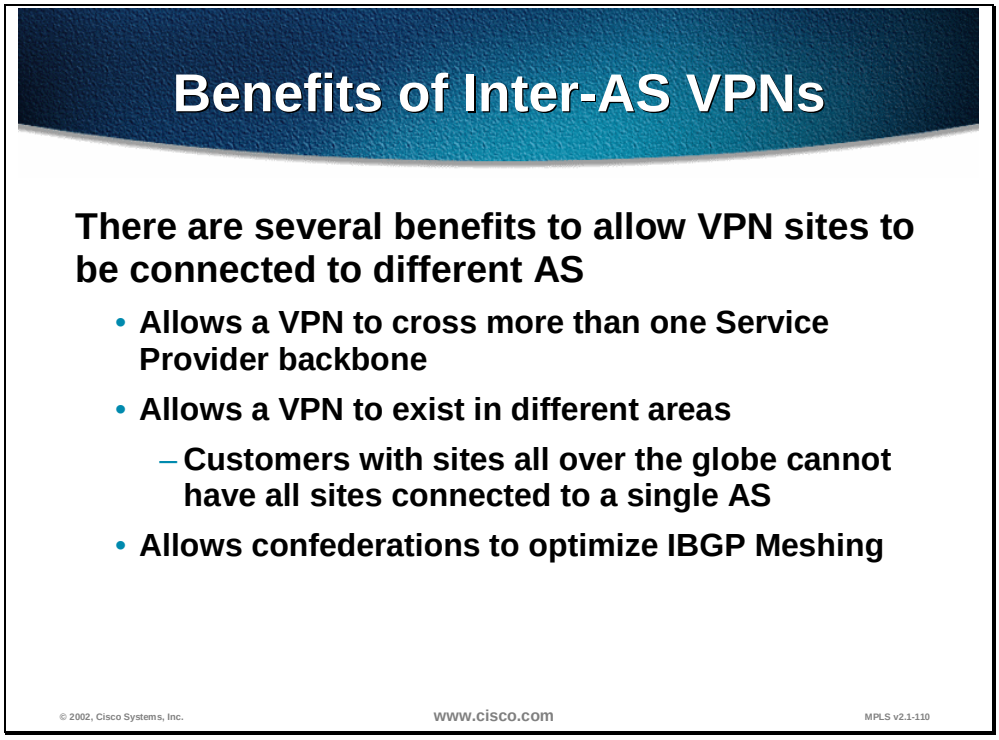

Allowing a VPN to span more than a single AS has several benefits:

- The VPN can cross more than one service provider backbone
- Customer sites can be spread over larger distances, potentially all over the globe, since they do not have to be connected to one and the same AS. A single AS does not scale as well as several cooperating AS.
- The provider AS can scale better by using BGP confederations to divide a large AS into smaller parts. Each part forms a subAS with its own independent IGP.

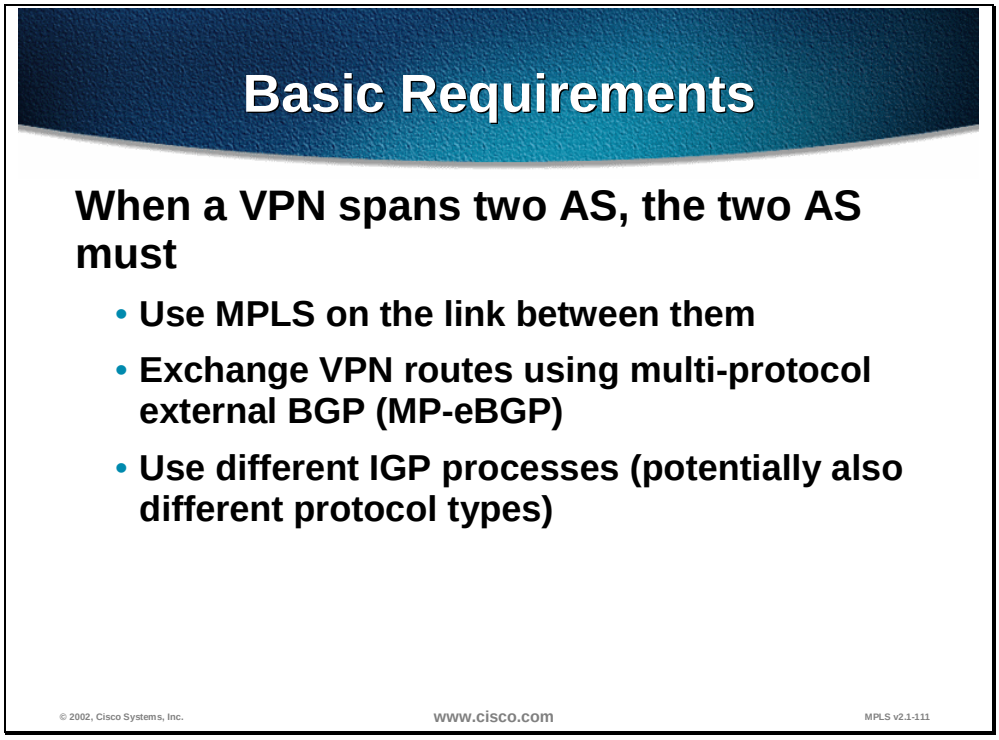

When a VPNS spans two (or more) AS, these AS must support the following basic features:

- The link between the two AS must support the exchange of MPLS packets. The customer's traffic must be MPLS encapsulated when it crosses all the provider AS and also the link between different AS.
- The two AS must exchange VPN routes between themselves using multi-protocol external BGP (MP-eBGP). MP-BGP is required to carry VPNv4 routes with label information, route-targets etc.
- Different AS should use different IGP processes. The routing interaction should be limited to MP-BGP. Two different providers may very well use different IGPs inside their AS.

LSP Stitching

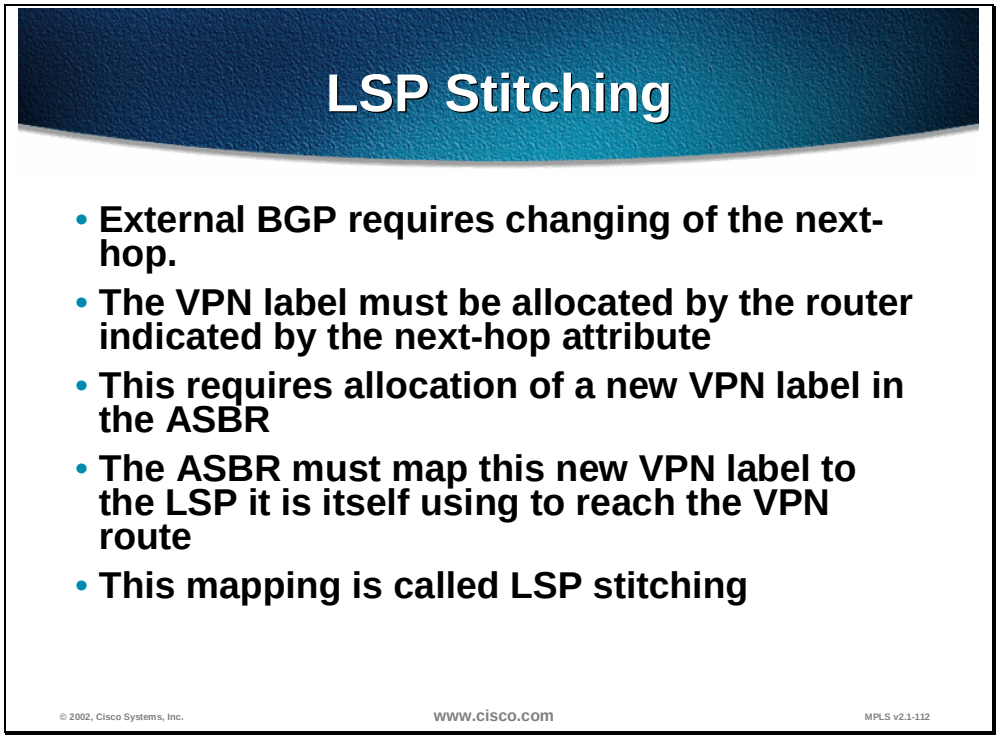

Label Switch Path (LSP) stitching means that two LSPs are connected (stitched) together. Packets arriving to a router on one of the LSPs are propagated into the other LSP.

In order to perform LSP stitching, the router, which performs the operation, must allocate a dedicated label for the stitching. This label must then be advertised to other routers. When these other routers are using the LSP associated with the dedicated label, the packets will reach the stitching router. The router recognizes the dedicated label and takes special action to forward the packets onto the next LSP.

LSP stitching is used in MPLS/VPN when the BGP next-hop attribute is changed. This is done by the ASBR on external BGP sessions and also when next-hop-self is configured in the **vpnv4 address family**.

The ASBR (or other router) that is changing the next-hop is allocating a dedicated label. The label is internally associated with the LSP to reach the final destination. This new label value is propagated in the MP-BGP update thus satisfying the requirement that the label value in the update must be assigned by the router which is the next-hop.

When the ASBR receives MPLS packets with the specially assigned label, it recognizes this and forwards the packets into the LSP that reaches the final destination.

Routing between ASes

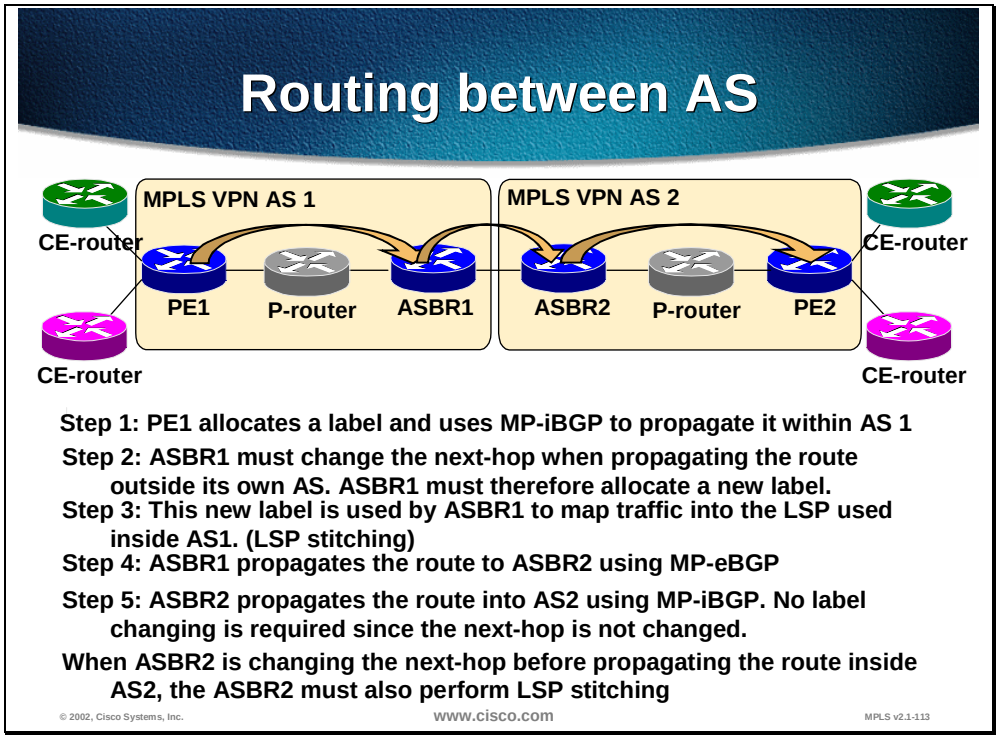

The figure illustrates how VPNs are spanning two AS, AS 1 and AS2.

PE1 will receive routing information from the CE routes. It creates VPNv4 routes, allocates local labels and forward the route on the MP-iBGP session to ASBR1. This is the standard MPLS/VPN procedures inside a single AS.

But ASBR1 is exchanging MP-eBGP with ASBR2. This means that ASBR1 must change the next-hop when the routes are propagated to ASBR2. ASBR1 therefore allocates a new label. This new label will be used by ASBR1 to map traffic into the LSP used inside AS1. ASBR1 is performing LSP stitching.

The VPNv4 route with the new label value is then propagated across the MP-eBGP session to ASBR2. ASBR2 does not have to change the next-hop when the VPNv4 route is propagated on the MP-iBGP session to PE2. ASBR2 therefore does not have to do LSP stitching in this direction.

If ASBR2 is configured to do next-hop-self on the MP-iBGP session to PE2, then ASBR2 must also perform LSP stitching.

Packet Forwarding between ASes

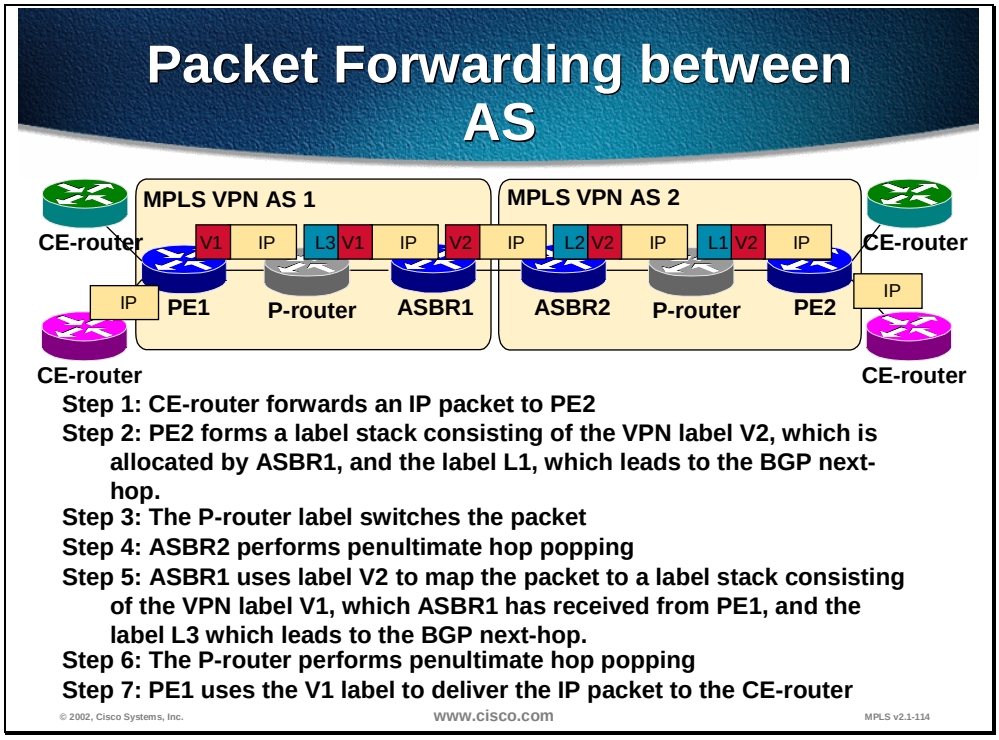

This figure illustrates how packets are propagated from the lower right CE router to the lower left CE router.

When PE2 receives the IP packet from the CE router, it makes a lookup in the virtual routing table for the customer site. It will find a label stack to use for forwarding. The top-most label, L1, is created by the LDP protocol and the local IGP inside AS 2. This reveals the requirement the PE2 must have IGP reachability to the next-hop. The second label is the VPN label, V2. It is the dedicated label assigned by ASBR1 to do the LSP stitching. PE2 forwards the packet with this label stack to the P-router.

When the P-router receives the MPLS packet, it looks only on the top-most label. The label instructs the P-router to swap top-most label to L2 and forward the packet to ASBR2. This operation is traditional label switching in an MPLS network.

ASBR2 receives the packet and performs penultimate-hop-popping. The top-most label is removed, thus exposing the VPN label. The MPLS packet is forwarded with a single label across the inter-AS link to ASBR1.

When ASBR1 receives the packet, it recognizes the label it has assigned for this LSP stitching. The label is associated with a label stack and an outgoing interface. ASBR1 can therefore remove the V2 label and impose the new label stack. This stack's top-most label, L3, indicates how the packet should be forwarded across AS 1 to PE1.

The VPN label in the new stack, V1, is assigned by PE1 and will be exposed as the P-router in AS 1 does penultimate-hop-popping.

PE1 now can use the label V1 to remove all labels and forward the packet as an IP packet to the CE-router.

Routing Inside a Confederation

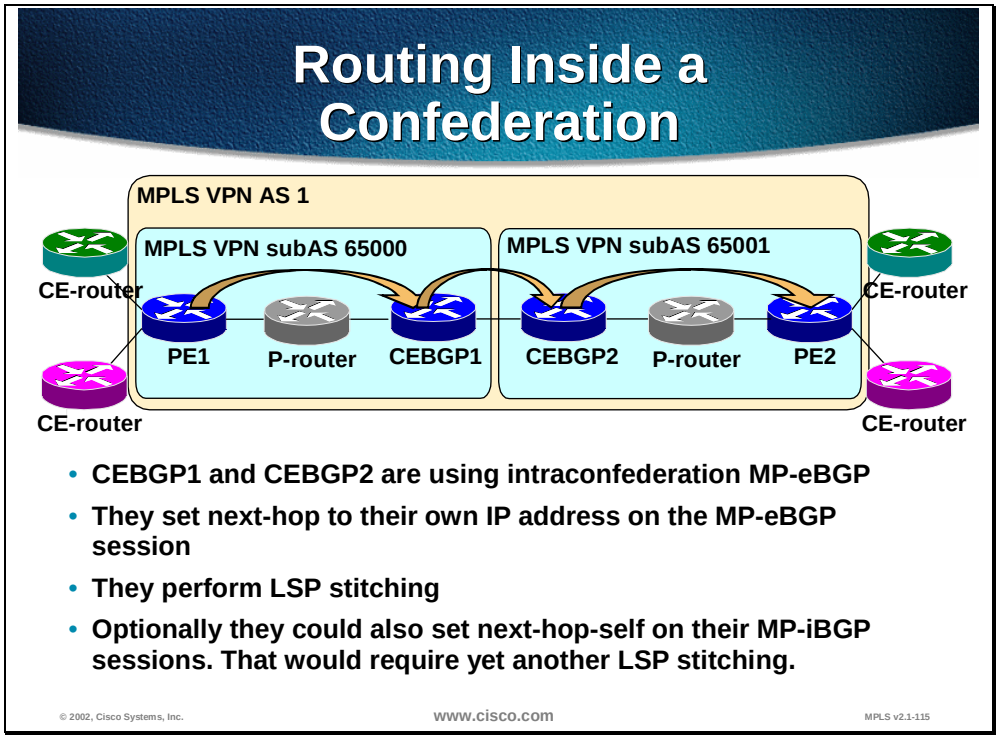

Large AS might want to use the BGP confederation tool for scaling. This scaling tool can be used both to reduce the IBGP full mesh, but also to reduce the complexity of the IGP by running independent IGP processes in each subAS.

When the latter is the case, the intra-confederation MP-eBGP session between CEBGP1 and CEBGP2 must change the next-hop. Since the two subAS are using different IGPs, the next-hop used inside subAS 65000 will not be a known address inside subAS 65001.

Changing the next-hop requires LSP stitching. Thus as CEBGP1 propagates the VPNv4 route to CEBGP2, the next-hop is changed and the VPN label value will be set to the dedicated label allocated for the stitching.

CEBGP2 does not have to change the next-hop as the VPNv4 route is propagated on the MP-iBGP session to PE2. But if CEBGP2 does change the next-hop, then CEBGP2 must also perform LSP stitching.

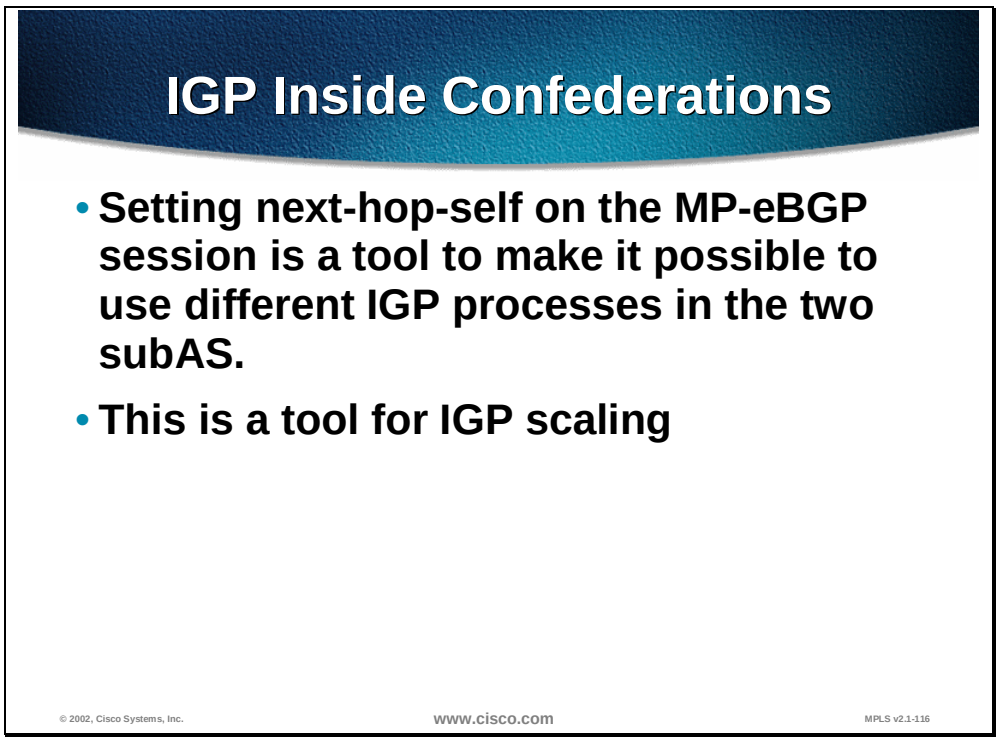

Setting next-hop-self on the MP-eBGP session is a tool to make it possible to use different IGP processes in the two subAS. It is therefore considered a scaling tool for the IGP.

Summary

After completing this lesson, you should be able to perform the following tasks:

- Explain the need for inter-AS VPNs
- Describe the propagation of VPN routes between two autonomous systems
- Explain how the label switched path in one AS is stitched together with an LSP in the other AS
- Describe the packet forwarding across AS boundaries
- Explain the benefits of using LSP stitching inside a BGP confederation

Lesson Review

Instructions

Answer the following questions:

1. What are the benefits of allowing a VPN to cross AS boundaries?
2. How is the VPN label propagated across an AS boundary?
3. Which router assigns the VPN label that is passed across the AS boundary?
4. How are packets transmitted on the link between the two AS?
5. How can the ASBR forward incoming packets from the other AS onto the correct LSP inside its own AS?
6. What is the benefit of using LSP stitching inside a BGP confederation?

Summary

After completing this lesson, you should be able to perform the following tasks:

- Identify major Virtual Private network topologies, their characteristics and usage scenarios
- Describe the differences between overlay VPN and peer-to-peer VPN
- List major technologies supporting overlay VPNs and peer-to-peer VPNs
- Position MPLS VPN in comparison with other peer-to-peer VPN implementations
- Describe major architectural blocks of MPLS VPN
- Describe MPLS VPN routing model and packet forwarding
- Describe the model for MPLS VPNs to span more than one autonomous system

